Monitoring and Accounting within the ATLAS Production System

J. Kennedy, Ludwig-Maximilians-Universität Munich ,Germany

Abstract

The presented monitoring framework builds on the experience gained during the ATLAS Data Challenge 2 and Rome physics workshop productions. During these previous productions several independent monitoring tools were created. Although these tools were created to some degree in isolation they provided a good degree of complementary functionality and are taken as a basis for the current framework. One of the main design goals of the current framework is to abstract the monitoring away from the central database of jobs, thus reducing the impact which the monitoring has on the production itself. Furthermore, the framework is aimed towards providing a common monitoring environment which may be seen as a high level source of information covering the 3 grid flavours used for ATLAS productions. The functionality of the framework is described with attention being paid to design considerations and implementation. The experience gained during the project is presented along with an outlook towards future developments.

ATLAS PRODUCTION

The ATLAS experiment at CERN will begin data taking in 2007 when the LHC is commissioned. From then on data is expected to be recorded at a rate of 2 PetaBytes per year with an additional 1 PetaByte of simulated data produced per year.

This huge data volume means that it is impossible to perform all the processing and analysis at CERN. ATLAS aims to utilise distributed computing resources in collaborating countries from all around the world.

To aid with the development of this distributed computing infrastructure a series of data challenges were started in 2002. The data challenges allow us to evaluate the ATLAS computing model, the full software suite, the data model and also to ensure the correctness of technical choices made for ATLAS computing.

PRODUCTION SYSTEM

The production system[1] is designed to provide a common framework in which any grid flavour may be integrated. The production system is formed from several individual elements which when plugged together provide the required functionality for the submission, tracking, recovery and validation of jobs. The individual elements of the production system may be summarised as follows.

- Common database for production jobs
- Common Supervisor run by all facilities/managers
- Executors developed by grid middleware experts
- Data Management system to allow intergrid data transfer and file cataloging

Jobs are defined in the central production database and are picked up by the supervisor element. The supervisor then uses an executor as an interface to the chosen grid. During the previous Data Challenge 2 and associated Rome production a peak rate of about 10,000 verified jobs per day was observed. The current phase of ATLAS MC production,Computing Systems Comissioning(CSC), is expected to surpass this performace by an order of magnitude within the year. Considering the scale of the planned production it is clear that monitoring tools are an essential part of the production system framework. The monitoring framework described here and the snapshots shown in the figures was developed during the pre-CSC production system development/testing phase.

MONITORING

Monitoring the ongoing ATLAS production is an important and diverse task. The idea behind the current monitoring framework is to implement the functionality found to be of most use in previous production and to act as a testbed for monitoring/accounting developments. Furthermore the current system attempts to abstract the monitoring away from the central produciton database in an aim to lower the impact that monitoring production has on the actual production itself.

The monitoring can be split into two seperate aspects. Firstly monitoring the production via interfaces to the production database and secondly monitoring the grid resources by using pre existing grid monitoring tools.

The majority of the functionality provided thus far is based on the former method.

Some areas of interest may be summarised as follows,

- Datasets - Identify which datasets are being processed and the success rate of previous sets
- Jobs - Identify where they are running and correlate any errors
- Accounting statistics - Ensure some simple global statistics are available at the click of a button

- Grid Status - Live time status of grid sites, avaliable resources, jobs running, storage space, installed SW, etc.

One major step forward in the evolution of the production system was the introduction of standard error codes for the executor elements. The codes are not complete and executors still return several of their personal error codes but work is ongoing to consolidate the error codes into a single standard set. The benefits of this error treatment are already being seen and more improvements are planned.

DATASET MONITORING

The Datasets monitoring is primarily used to identify which datasets are available and which are currently being processed. A cron job is used to gather the information on datasets from the production database and this information is then stored in a comma seperated variables (CSV) file for later parsing. Daily copies of this CSV file are kept to allow a seven day history which may be used to track dataset progress.

Once the dataset information is gathered it can be access via a simple web interface. Currently this interface allows users to select datasets in different states of prossessing and to perform simple string matching against the dataset name. The progress on datasets over the last week can also be obtained, to allow users/submitters to identify stalling datasets. The results from a simple search of the datasets can be seen in figure 1.

JOB MONITORING

The monitoring of actual production jobs and the correlation of their errors is probrably the major part of the current monitoring functionality. The ability to track the jobs as they sit in queueing states,execute and fail and also to correlate the errors on failure provides us with the ability to respond faster to problems. Views of the latest batch of jobs are available for the following criteria.

- As a function of Site
- As a function of Executor
- As a function of Task(Dataset)

The information about jobs during the last 24hrs is collected from the central production database and stored in a 24hr snapshot database. This collection occurs each hour so a 24hr snapshot of jobs is available with a maximum latency of 1hr. Storing information this way allows future queries on the latest information to be more efficient.

The interface to the job information is provided by a web page which groups together some simple one-click buttons and also some more advanced configuarable search form.

The information about recent jobs is displayed in two formats

- Tables of pending,running,finished,failed jobs
- Tables of correlated errors

Each table is computed as a function of site/executor/dataset for the previous 24 and 1 hour periods. Figure 2, for example, show the job distribution as a function of the contributing sites.

In addition the executor name can be used to select this same information for a given executor. This provides an executor with the ability to track their jobs seeing which sites are perticularly good/bad for them and also to correlate the errors at the sites to determine if a specific failure is associated with a specific site.

A more advanced and configurable interface to the 24hr snapshot table provides the ability to perform a more complex query by matching against fields such as the site-name,executorname,dataset name, etc.

ACCOUNTING

Basic accounting has been added to the monitoring project in the form of a daily statistics table which records the following information for each day.

- Number of Finished jobs
- CPU consumption of Finished jobs
- Number of Failed jobs
- CPU consumption of Failed jobs
- Number of Pending Jobs
- Number of Running Jobs

The Accounting information may be accessed via a web page interface. This interface allows a user to specify a date range and to gain information on a daily basis or summed and also split via grid(executor) flavour or sum over all grids. Although the information is only available on an executor basis it does allow for quick accounting queries which would otherwise have a larger impact on the production database.

OVERVIEWS

Witihn the overviews section we aim to provide views of the whole production and not just a snapshot of the latest jobs or condensed information. Currently only site information for a given date range available. The Overview queries are queries which present views of the production which require some amount of configurability and a query over the whole of the central production database. Such queries are obviously more expensive than snapshots and something we originally set put to avoid. However the information/functionality they provide is seen as vital for the production.

Site results can be queried over a given date range. It is possible to use wildcards to query a group of sites, for example figure 3 shows the results for production at German sites in the first month of 2006. In addition to displaying the efficiency of the selected site the associated errors within the given time period are correlated.

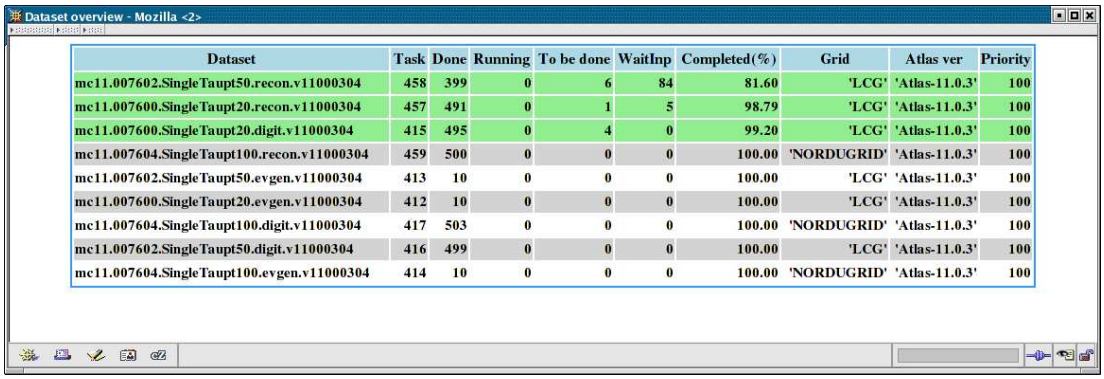

Dataset	Task	Done	Running	To be done	WaitInp	Completed(%)	Grid	Atlas ver	Priority
mc11.007602.SingleTaupt50.recon.v11000304	458	399	0	6	84	81.60	'LCG'	'Atlas-11.0.3'	100
mc11.007600.SingleTaupt20.recon.v11000304	457	491	0	1	5	98.79	'LCG'	'Atlas-11.0.3'	100
mc11.007600.SingleTaupt20.digit.v11000304	415	495	0	4	0	99.20	'LCG'	'Atlas-11.0.3'	100
mc11.007604.SingleTaupt100.recon.v11000304	459	500	0	0	0	100.00	'NORDUGRID'	'Atlas-11.0.3'	100
mc11.007602.SingleTaupt50.evgen.v11000304	413	10	0	0	0	100.00	'LCG'	'Atlas-11.0.3'	100
mc11.007600.SingleTaupt20.evgen.v11000304	412	10	0	0	0	100.00	'LCG'	'Atlas-11.0.3'	100
mc11.007604.SingleTaupt100.digit.v11000304	417	503	0	0	0	100.00	'NORDUGRID'	'Atlas-11.0.3'	100
mc11.007602.SingleTaupt50.digit.v11000304	416	499	0	0	0	100.00	'LCG'	'Atlas-11.0.3'	100
mc11.007604.SingleTaupt100.evgen.v11000304	414	10	0	0	0	100.00	'NORDUGRID'	'Atlas-11.0.3'	100

Figure 1: An example of a search for SingleTau datasets, green shows sets still being processed

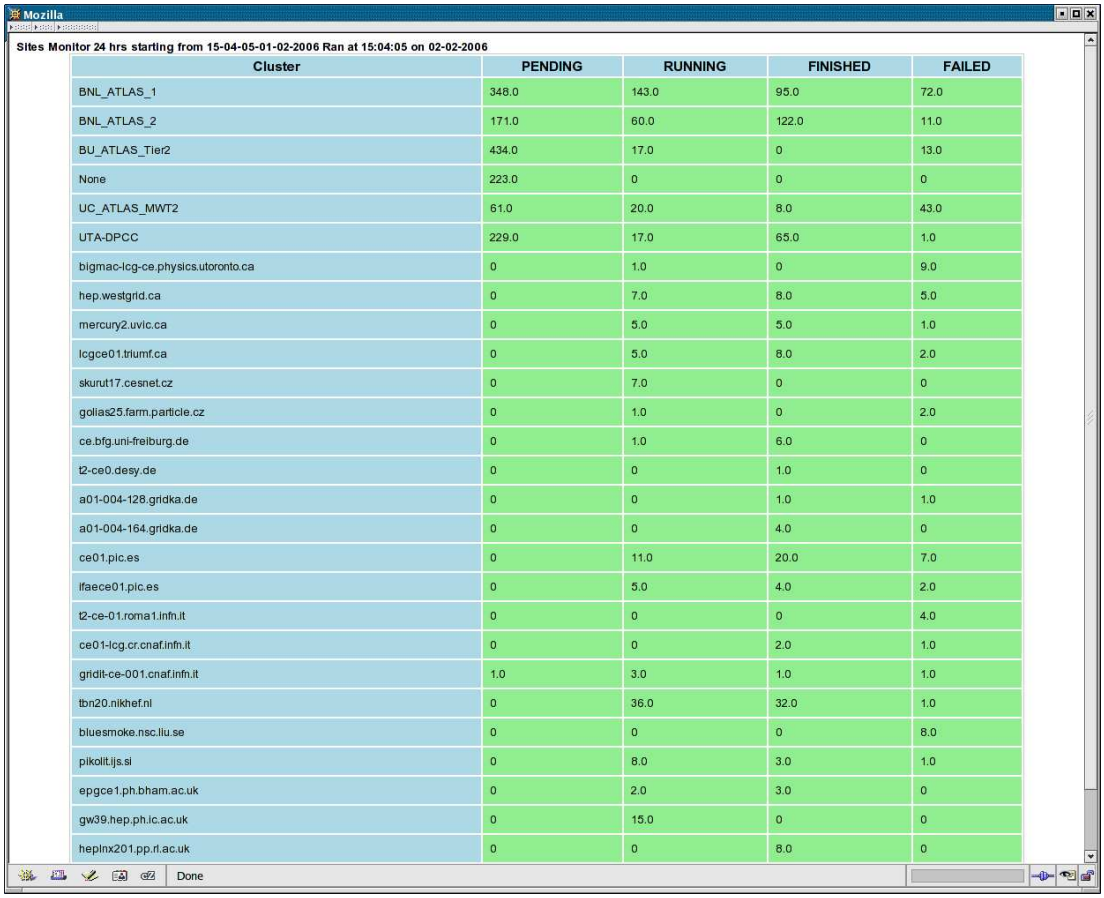

Sites Monitor 24 hrs starting from 15-04-05-01-02-2006 Ran at 15:04:05 on 02-02-2006

Cluster	PENDING	RUNNING	FINISHED	FAILED
BNL_ATLAS_1	348.0	143.0	95.0	72.0
BNL_ATLAS_2	171.0	60.0	122.0	11.0
BU_ATLAS_Tier2	434.0	17.0	0	13.0
None	223.0	0	0	0
UC_ATLAS_MWT2	61.0	20.0	8.0	43.0
UTA-DPCC	229.0	17.0	65.0	1.0
bigmac-lcg-ce.physics.utoronto.ca	0	1.0	0	9.0
hep.westgrid.ca	0	7.0	8.0	5.0
mercury2.uvic.ca	0	5.0	5.0	1.0
lcgce01.triumf.ca	0	5.0	8.0	2.0
skurut17.cesnet.cz	0	7.0	0	0
golias25.farm.particle.cz	0	1.0	0	2.0
ce.bfg.uni-freiburg.de	0	1.0	6.0	0
t2-ce0.desy.de	0	0	1.0	0
a01-004-128.gridka.de	0	0	1.0	1.0
a01-004-164.gridka.de	0	0	4.0	0
ce01.pic.es	0	11.0	20.0	7.0
ifaece01.pic.es	0	5.0	4.0	2.0
t2-ce-01.roma1.infn.it	0	0	0	4.0
ce01-lcg.cr.cnaf.infn.it	0	0	2.0	1.0
gridit-ce-001.cnaf.infn.it	1.0	3.0	1.0	1.0
tbn20.nikhef.nl	0	36.0	32.0	1.0
bluesmoke.nsc.liu.se	0	0	0	8.0
pikolit.ijs.si	0	8.0	3.0	1.0
epgce1.ph.bham.ac.uk	0	2.0	3.0	0
gw39.hep.ph.ic.ac.uk	0	15.0	0	0
heplnx201.pp.rl.ac.uk	0	0	8.0	0

Figure 2: Job distribution as a funciton of site

The overviews should aim to provide information and functionality such as,

- Global correlations of errors
- Global collections of job attempt noumbers
- Global views of job states
- A Tunable Accounting Interface

SITE INDEXES

The idea of site indexes is to use grid monitoring tools/information to provide a live time view of the underlying grid on which ATLAS runs its jobs. We have attempted to gather a common set of information from each grid flavour which is useful to provide an overview of the grid status for ATLAS.

EXPERIENCE

This section collects some notes on the experience gained during the development of the monitoring framework.

#	Site	Finished	Failed	Efficiency (%)
1	t2-ce0.desy.de	1203	793	60
2	grid-ce0.desy.de	737	418	64
3	a01-004-128.gridka.de	205	10	95
4	a01-004-164.gridka.de	132		100
5	ce.bfg.uni-freiburg.de	973	353	73
6	grid-ce.desy.de	33	16	67
7	lcg-gridka-ce.fzk.de	2		100
-	Total	3285	1590	67

Site	TRF	WRP_IN	None	WRP_OUT	CE_PBS_WALLTIME	CG_GETOUT_EMPTYOUT
t2-ce0.desy.de	47	291	5	12	2	436
grid-ce0.desy.de	75	259	5	6	2	69
a01-004-128.gridka.de	2	2	0	0	5	1
a01-004-164.gridka.de	0	0	0	0	0	0
ce.bfg.uni-freiburg.de	79	211	0	1	2	56
grid-ce.desy.de	9	5	0	0	1	1
lcg-gridka-ce.fzk.de	0	0	0	0	0	0

Figure 3: Overview of German site statistics and errors for Jan 2006

Caching: The monitoring project has made use of snapshot tables to store either recent data and to store compressed information thus leading to increased speed when querying. We believe that this has so far proved an adequate solution to ward off problems due to overly agressive monitoring.

One-click: As the project progressed it became apparent that users tend to prefere to have simple one-click button functionality to gather basic information. However these simple queries are often greatly helped if some configurable tools exist for more complex query construction.

Overviews: Although initially the monitoring interface was to be restricted to querying the snapshot tables it has become clear that queries of the whole dataset are needed. The usefulness of the gathered information outweighs the cost of such a query. Providing the functionalty for such queries now seems a must but it should be accompanied with a warning about the cost and users should be educated to perform such extensive queries only when really needed.

OUTLOOK

The monitoring framework as presented provides a fairly broad set of functionality and has aided with the startup of the CSC production. However it is by no means complete and should continue to evolve as the production ramps up during the comming year.

One obvious missing ingredient is simple graphics to help display results. Bar/Pie chart displays would allow for quicker and more intuative interpretation of the gathered information and their generation and usage in several areas is currently under investigation. In addition global views of the production and additional functionality to identify pathalogical jobs are also required.

The future of monitoiring is guaranteed in such a high throughput production environment, as such the project will continue to evolve as in line with the production system. By providing a framework and gathering user feedback it is hoped that the current project will aid in the development of the proposed offline shift framework.

ACKNOWLEDGEMENTS

I would like to thank Guido Negri, Marco Mambelli, Oxana Smirnova and Tomasz Wlodek for their contributions to the project.

REFERENCES

[1] G. Poulard, “ATLAS Experience on Large Scale Producions on the Grid”, These proceedings.