Report on Monitoring / Accounting tools

V5

I. Blanquer, A. Carrion[(1)]
J. Gomes, G. Borges[(2)]
UPV, Valencia, Spain[(1)]
LIP, Lisbon, Portugal[(2)]

Preface

The aim of this document is to provide information about the characteristics of diverse tools for monitoring or/and accounting in Grid infrastructures. For each tool considered, the following section will outline features such as: main functionality, technical requirements or operational complexity.

List of tools

GStat

GStat (http://gstat-prod.cern.ch/gstat) generates web based reports which allow administrators to trouble shoot IS (Information System) issues. For that purpose, GStat scripts are executed periodically (approximately every 30 minutes) to collect information published by each site. Then, the information is processed in order to detect IS failures and errors. In addition to that, GStat allows accessing information about the sites (services, software and VOs supported at the site) and usage information (statistics on jobs slots, jobs and storage space). GStat also offers historical test results in graphs and tables, being the graphs available for the last 24 hours, 3 days, 3 months or one year.

GStat runs on a single server and the current deployment includes two independent replicas, one in Taiwan (http://goc.grid.sinica.edu.tw/gstat/) and one in Italy (egee017.cnaf.infn.it/gstat). Thanks to the availability of these replicas via Web, the need of performing a private installation is eliminated.

VO added value: GStat has a VO view that provides consolidated information about VO jobs and storage. The data is gathered from the information system and it can be browsed and viewed in statistics graphs (hourly, daily, weekly, monthly and yearly). The tool is helpful to identify information system problems at the sites and in TOP BDIIs. It is also a friendly way to browse the information system.

Grid View

Grid View (http://gridview.cern.ch/) is a high-level monitoring tool that uses different information sources from sites monitoring and logging to consolidate historical and instantaneous views of the performance of a specific view of the infrastructure. The tool was designed by CERN for WLCG and the new interface has been more tailored for use by LHC and other CERN experiment VOs.

It is not intended to be replicated but it is synchronised to CIC and GOCDB portals. It can be used for VO statistics easily. Not all the transactions could be registered, depending on the configuration of the sites.

GridView allows to keep track of the infrastructure reliability and availability metrics. Using this tool, it is possible to view the performance historical data of site' different components and services, and their contributions to the overall site value along different time scales (days, weeks, months and years).

VO added value: It provides statistical information per VO about data transfers, job status and FTS (File Transfer Service). It is also able to present values for services avalability and reliability metrics based on a given set of tests considered critical for a given VO.

Grid MAP

Grid Map (http://gridmap.cern.ch) is a visualisation tool developed in the frame of the CERN open lab that provides a hierarchical view of the resources available in the LCG infrastructure. It provides views for: CPU capacity, running jobs and monitoring test results, using a visualization approach based on area sizes to outline the size of the resource provider and uses a simple colour scale (green-orange-red) to denote operative problems.

This tool is quite lightweight for the client-side and relies on existing monitoring systems. It is based on Apache application container and python.

VO added value: The tool can provide a visualization of the monitoring test results per VO. A VO representative can easily get a quantitative idea of the ammount of working / problematic resources, and their relative number with respect to the overall infrastructure

Nagios

Nagios is a established monitoring tool which allows monitoring and notifications for failures of hosts, networks and Grid services. Nagios is not a Grid monitoring tool 'per se' but it can be used to monitor the operation of Grid services in the same way that it is used to monitor more common network services like HTTP or IMAP. Regarding the monitoring of host resources, Nagios can retrieve data such as: the processor load, the disk and memory usage, the running processes, the log files, etc. It can also offer aggregated information, although it is basically related to the availability status of the sites.

The basic Nagios installation process is not very complicated. In fact, it simply requires the local installation of a server with a set of prerequisites (Unix-like OS, Apache2 and various perl modules depending on the distribution).

Nagios provides a web interfaces in which it can be seen the current network status, notification and problem history, log file, etc. The access to this web interface is managed by a simple authorization scheme (based on Grid certificates) that allows determining what users can see and do.

One of the most interesting features of Nagios is its great extensibility and flexibility thanks to a simple plugin design that allows skilled users to develop their own host and service checks.

VO added value: VOs can easily deploy and establish a monitoring infrastructure, build VO specific nagios tests, and test VO specific needs on their infrastructure.

Ganglia

'Ganglia' (http://ganglia.info/) is monitoring system targeted at high-performance computing systems such as clusters and Grids. While Nagios shows more qualitative information and can issue alarm notifications, Ganglia can provide quantitative information using various metrics such as: CPU utilization, Memory usage, disk usage, network statistics, number of running processes, etc. For each metric, Ganglia shows a historical graph for the past hour, week, month or year. Nevertheless, Ganglia can be integrated with Nagios. In other words, Ganglia provides more fine-grain monitoring status of a resource. One of the Ganglia added values is its capability to offer many different metrics views, and therefore, establish correlation between different quantities in order to optimize resource usage.

Ganglia is very portable and runs on a extensive set of operating systems and processor architectures: Linux (i386, ia64, sparc, alpha, powerpc, m68k, mips, arm, hppa, s390), FreeBSD, NetBSD, OpenBSD, DragonflyBSD, MacOS X, Solaris, AIX, IRIX, Tru64, HPUX and Windows NT/XP/2000/2003/2008). Moreover, the effort necessary for installing, configuring and managing a Ganglia server is not high.

In order to access the application, Ganglia provides a PHP Web Frontend which shows the information gathered via real-time dynamic web pages.

VO added value: Could be used to provide consolidated information about VO resource usage, however development effort will be required to develop the probes for gathering information from grid resources.

Hudson

Hudson, (http://hudson-ci.org/), is a tool for monitoring executions of externally-run jobs (such as Grid jobs), gathering and organizing user-customised tests.

Hudson is much more limited than Nagios but easier to install and configure (via a web GUI). Since it is based in Java, the portability is assured (to run minimally it requires JRE 1.5 or later or a servlet container that supports Servlet 2.4/JSP 2.0 or later).

Also, Hudson has a plugin support which allows the deployment of particular user cases.

VO added value: Simple tool for monitoring resources from the VO perspective, using own services.

SAM

SAM or Service Availability Monitoring (https://lcg-sam.cern.ch:8443/sam/sam.py) is a framework for the monitoring of production and pre-production grid sites. It provides a set of probes which are submitted at regular intervals, and a database (Oracle) that stores test definitions, node and site information, and the test results. In contrast to other monitoring tools, SAM provides monitoring of grid services from a user perspective. The metrics used by SAM are the following: availability of a single node providing a service (CE, RB, etc.), availability

of a whole site (all the services provided by the site), availability of an individual central services per VO and the whole availability of all the central services of a given VO. All this information is offered in different types of time integration: daily, weekly and monthly.

The installation and further support of the SAM framework requires considerable effort and the use of the Oracle database is a limiting factor. The access to the application is provided by a web portal (phyton-based) which displays and permits querying the test results. Only users with valid grid certificates can view these pages.

One of the most interesting features of SAM is the adaptability to many scenarios of usage.

VO added value: Allows performing VO specific tests. However SAM will be likely dropped in favour of r-Nagios.

EGI Metrics portal

The goal of the EGI Metrics portal (link?) is to have a portal which automatically collects different kinds of data and calculates a set of metrics that can help estimating the performance of a Grid project and keep track of its evolution. All the information is collected from different sources like APEL, GOCDB, GGUS and GridView and the results are showed in a single place (the portal). The metrics used by the portal can be splitted into four groups: size metrics (show changes in the available resources of the Grid infrastructure), usage metrics (quantify the resource usage and efficiency), operations metrics (central services availability and reliability) and user support metrics (number of user support tickets created and response times).

The installation of the portal is not difficult and only requires the following: Apache web server, PGP, Jpgraph and Python. The low number of dependencies makes EGI Metrics portal very flexible if modifications should be needed.

VO added value: Consolidated view of the VO activity without requiring specific configurations. Good for reporting and comparing the performance of the users in the VO.

Dashboard

Dashboard (http://dashboard.cern.ch/) collects information from multiple sources and provides monitoring of various computing activities of the Grid infrastructure like: job processing, data management, site status and transfer tests. Due to the high amount of data collected by Dashboards, it uses a backend database based in Oracle. Moreover, Dashboard shows historical information in a web-based visualization system. The dashboard was developed for the LHC experiments.

The installation of Dashboard requires Python and Oracle. The framework is modular and typically is composed of information collectors that populate a data repository, plus APIs to access the data. The framework is also composed of agents that perform regular tasks such as collect data, produce statistics, analyze statistics or send alarms. The dashboard is compatible with the gLite, ARC and VDT grid middleware. The dashboard applications are developed using the framework and can be divided in common and VO specific dashboards. Some examples are:

- Data Transfer: is a VO specific dashboard for data management monitoring.
- Job Monitoring: is a common dashboard aimed at several types of audiences, users, VO managers, support teams, site admins.
- Task monitoring: is a VO specific dashboard that provides a user centric view of user submitted tasks.
- Job Summary: is a VO specific dashboard that provides a view of the VO recent running jobs.
- Historical view: is a VO specific dashboard that provides statistics about VO executed jobs.
- Site Monitoring: is a common dashboard that visualizes information about SAM tests.
- Site Status Board: is a common dashboard for the monitoring of site performance from the VO perspective.

VO added value: The dashboard provides a VO centric view of the infrastructure by collecting information from several sources and presenting it on a VO perspective. The tool is useful for the VO managers and operations to obtain a global view, for the sites supporting the VOs to understand how they are performing, for the end users running tasks in the grid.

Bazaar

Bazaar (http://grid.cyfronet.pl/bazaar) is a resource allocation tool which formalizes the management of resources by using SLAs (the agreement between a Resource Provider and a VO about details of services supplied) in a Grid infrastructure. Thus, it does not have monitor or testing capabilities.

Bazaar has its own web interface (integrated in the EGI Operations portal).

VO added value: Allows VOs to request grid resources and services.

Regional Nagios

The r-Nagios is a grid monitoring framework based on Nagios and developed by WLCG to monitor grid infrastructures. Regional Nagios was conceived to replace the SAM tests under the assumption that each NGI or federation of NGIs would deploy and operate its own grid monitoring system.

Besides Nagios itself, the r-Nagios uses a configuration generator that builds a Nagios configuration from information about the grid topology and the metrics to be measured through probes. The Nagios probes for the grid are performed from a User Interface. For authentication purposes a valid proxy certificate to run the probes is obtained from a MyProxy repository.

The r-Nagios probe results can be used by the operations dashboard and tools such as MyEGEE/MyEGI. The results of the probes are stored in the r-Nagios Metrics Database. The results are published in a Message Bus which enables a Nagios instance installed as remote to listen for this information and display it. This feature can be used to consolidate information from several r-Nagios instances.

MyEGEE/MyEGI is a tool for visualizing the r-Nagios metrics database. The tool is a separate component from r-Nagios but is installed together.

VO added value: The tool can be conceivably used to run VO specific tests.

HammerCloud

HammerCloud Is a distributed stress testing system based on the Ganga job management system. It is used by the CERN ATLAS experiment to test a set of sites by running a large number of real analysis jobs. HammerCloud is now being adopted by other LHC VOs. It is implemented as a web application and uses gLite User Interface commands to submit the jobs. Extensive test reports are produced and presented through a web interface.

VO added value: It is useful to validate benchmark and tune the sites.

Summary of monitoring systems

	Inf. source	Main objective	Requirements	Configuration	Operation costs	URL
Gstat	IS (BDII)	Monitor IS system and statistics	Own server, although already working for all CIC VOs and GOCDB sites.	Tests provided	Zero, considering that it is currently provided for all registered VOs.	http://gstat-prod.cern.ch/gstat http://goc.grid.sinica.edu.tw/gstat/
GridView	Several logs and monitoring sources	To provide a high-level view of different services	Own server, although already working for all CIC VOs and GOCDB sites.	Not flexible	Zero, considering that it is currently provided for all registered VOs.	http://gridview.cern.ch/
GridMap	SAM tests	Monitor availability and overall site status	Web server	Not specified, but inherent to SAM tests	Not described	http://gridmap.cern.ch
Nagios	Own tests	General Monitoring	Own server	General Grid tests available, specific tests can be developed	Medium: Installation and tuning requires a non-trivial knowledge.	https://twiki.cern.ch/twiki/bin/view/EGEE/ExternalROCNagios
Ganglia	Own tests	Fine Grain General Monitoring	Own server	General tests on the resources.	Typical operations of any 24x7 service.	http://ganglia.info/
Hudson	Own tests	General Monitoring	Own server	Based on Java, easy to deploy and customise.	Lower than in nagios. Typical operations of any 24x7 service.	http://hudson-ci.org/
SAM	Own tests	General Monitoring	Own server, Oracle license.	Complex but powerful.	High, mostly intended for infrastructure admins.	https://lcg-sam.cern.ch:8443/sam/sam.py
EGI Metrics portal	Information from monitoring and database systems	Consolidate distributed accounting information	Apache web server, PGP, Jpgraph and Python. It will be available for the CIC VOs.	Flexible tool.	Minimum if no changes are required.	http://www3.egee.cesga.es/gridsite/metrics/CESGA/egee_view.php
Dashboard	Own tests + other sources	General monitoring and accounting information	Own server and and Oracle back-end.	Adaptable to many different scenarios	High: Very complex tool for large VOs needing detailed monitoring.	http://dashboard.cern.ch/
Bazaar	User based	Reservation of resources and follow-on	Own server, although it is already working for all CIC VOs.	Fixed configuration	High at the human operation level for the coordination of user	http://grid.cyfronet.pl/bazaar

					response.	
r-Nagios	Own tests	General Monitoring	Own server and UI	General Grid tests available, specific tests can be developed	High: Installation and tuning requires a non-trivial knowledge.	https://twiki.cern.ch/twiki/bin/view/EGEE/ExternalROCNagios
Hammer-Cloud	Own tests	Job Execution Stress tester	Own server and UI	Requires a real user application adapted for the purpose	High: Installation and tuning requires a non-trivial knowledge.	http://hammercloud.cern.ch/