GRIDVIEW : A GRID MONITORING AND VISUALIZATION TOOL

Rajesh Kalmady*, Digamber Sonvane*, Kislay Bhatt*, Phool Chand*,
Computer Division, Bhabha Atomic Research Centre, Mumbai, India

James Casey#, Zdenek Sekera#, IT Department, CERN, Geneva, Switzerland

Abstract

The LHC Computing Grid (LCG) connects together hundreds of sites consisting of thousands of components such as computing resources, storage resources, network infrastructure and so on. Various Grid Operation Centres (GOCs) and Regional Operation Centres (ROCs) are setup to monitor the status and operations of the grid. This paper describes Gridview, a Grid Monitoring and Visualization Tool being developed for use primarily at GOCs and ROCs. It can also be used by Site Administrators and Network Administrators at various sites to view metrics for their site and by the VO Administrators to get a brief of resource availability/usage for their virtual organizations. The objective of this tool is to fetch grid status and fault data from different sensors and monitoring tools at various sites, archive it into a central database, analyze, summarize it and display it in a graphical form. It is intended to serve as a dashboard (central interface) for status and fault information of the entire grid.. The architecture of the tool is very flexible and new data sources can be easily added in the system. The first version of Gridview is deployed and was used extensively for online monitoring of data transfers among grid sites during LCG Service Challenge 3 (SC3) throughput tests. The paper discusses the architecture, current implementation and future enhancements to this tool. It summarizes the architectural and functional requirements of a monitoring tool for the grid infrastructure.

INTRODUCTION

The computing infrastructure for the LHC (Large Hadron Collider) is being implemented as a worldwide grid called LCG (LHC Computing Grid) [1] spanning three continents. A 3-tier structure of computing centres is being setup to cater to the storage and analysis of data of the order of several petabytes and computational requirements of thousands of users. The LCG connects together hundreds of sites consisting of thousands of components such as computing resources, storage resources, network infrastructure and so on. Jobs belonging to the four experiments will analyze the data being generated in the LHC which in turn will lead to a large scale movement of data across the grid.

*{rajesh, sonvane, kislay, phool}@barc.ernet.in
#{James.Casey, Zdenek.Sekera}@cern.ch

Various performance and functional metrics are monitored by LCG Information Providers and different monitoring tools at various sites of the Grid. These performance and functional metrics include parameters such as availability of compute resources, storage resources, status of jobs, status of various services, network throughput, file transfer logs and so on. In order to visualize the grid as a whole, there is a need for a tool which will display a dashboard for the entire grid indicating its status, collect the monitored data from the various systems into a common database, analyze, summarize and display it in a graphical form. Gridview [2] is a monitoring and visualization tool being developed for this purpose.

The objective of this tool is to provide a high level view of various Grid Resources and functional aspects of the LHC Computing Grid (LCG). The tool should display a dashboard for the entire grid and provide a gist of various metrics monitored by different monitoring tools at various sites. It should also analyze the monitored data for determining usage, behavior and performance of the Grid, detecting and notifying fault situations and user-defined events. The tool is targeted primarily for use in various Grid Operation Centres (GOCs) and Regional Operation Centres (ROCs) to monitor status of the entire grid. It will also be used by Site Administrators and Network Administrators at various sites to view metrics for their site and by the VO Administrators to get a brief of resource availability/usage for their virtual organizations.

GRIDVIEW ARCHITECTURE

The tool is based on the concept of loosely coupled components with independent sensors, transport, archival, analysis and visualization components. The sensors are LCG information providers and other monitoring tools, the transport mechanism used is Relational Grid Monitoring Architecture (R-GMA) [3], and Gridview provides the central archival, analysis and visualization functionality. The architecture of the tool is very flexible and new data sources can be easily added in the system.

The various functional components of Gridview (Fig. 1) include

i) R-GMA Archiver Module
ii) Data Analysis and Summarization Module
iii) Presentation Module
iv) GUI & Visualization Module

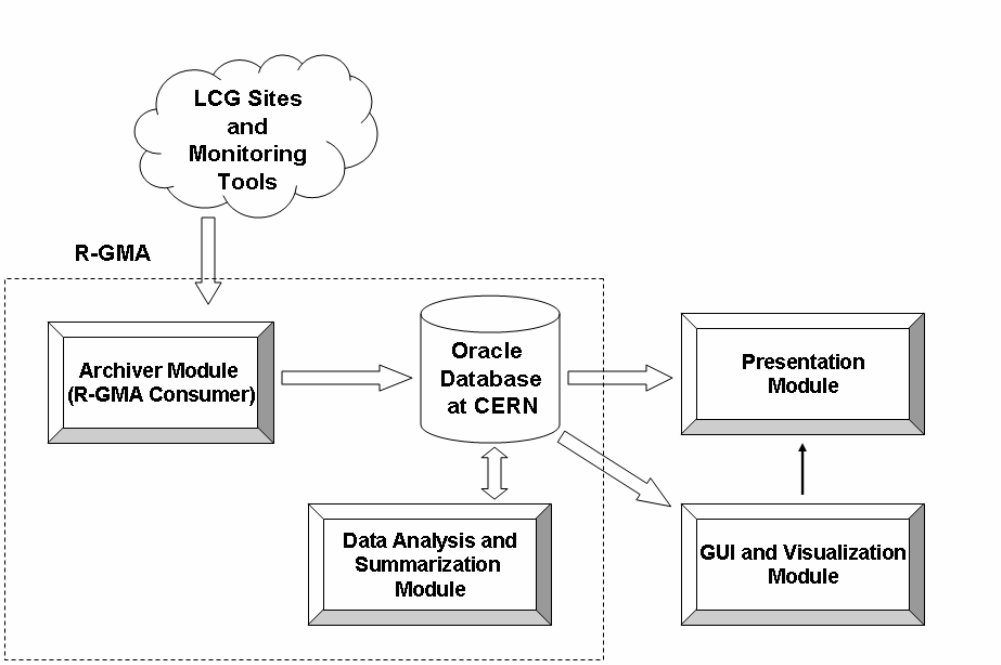

Figure 1: Block Schematic of Gridview Architecture

R-GMA Archiver Module

The information providers and monitoring tools at various sites publish information to R-GMA (Relational Grid Monitoring Architecture). Gridview's archiver module fetches the data from R-GMA and archives it in the central Oracle database at CERN. It is implemented in Java.

Data Analysis and Summarization Module

This module analyzes the archived data and determines usage, behavior and performance of the Grid, detects and notifies fault situations and user-defined events. It stores the derived information in the central database. It is implemented in Java and PHP.

Presentation Module

This module presents current and history information regarding grid status and functionality using conventional graphs (Line graphs, Histograms, Pie Charts etc) and HTML tables. It is implemented using PHP, Javascript, HTML.

GUI and Visualization Module

This is the user interface for GRIDVIEW. It will display various menus and a dashboard for the entire grid. The dashboard will pictorially show various regions/grid sites and facilitate zooming into regions/sites to have a detailed view of local resources such as CEs, SEs, RBs, etc. The user can graphically select the interested grid resources and metrics to view. It will also display status summary and fault information in pictorial form using 3-D graphics. This will be implemented using server side java application and java applet, the two communicating using Web Services.

GRIDVIEW INFORMATION PROVIDERS

The various sensors and other monitoring tools that generate data for Gridview are described below. Most of these tools already publish data to R-GMA and others may do it in near future.

LCG-2 Information Providers

The information providers running at each site publish information about site and its resources (CEs, SEs, Network) according to the predefined GLUE schema.

Gridftp Logs

These are file transfer logs from GSIFTP protocol. They provide details like initiating host, source node, destination node, file name, size, user, data transfer rate, time taken etc. A set of scripts parses these logs at every site and publishes the information to R-GMA.

RB Job State

It is a set of scripts that periodically collect information about the status (submitted, ready, waiting etc) of all the jobs from the Logging and Bookkeeping Service (LB) at various resource brokers (RB) and publish it to R-GMA.

WN Job State

It is a set of scripts that periodically collect job statistics (usage information like wall clock time taken, CPU time used, used real and virtual memory space etc) for all jobs from the Logging and Bookkeeping Service (LB) at various resource brokers (RB) and publish it to R-GMA.

Site Functional Test (SFT) [4]

It daily runs a site testing script on every site and publishes its results. The test script contains several grid functional tests and WN configuration tests. Grid functional tests include Replica Tests using EDG Replica Manager and LCG Tools, checking R-GMA client configuration and functionality and checking list of VOs supported by the site. WN configuration tests include checks for environment variables set, paths and accessibility of experiment software, mount points, installed RPMs, software versions, replica manager configuration and BrokerInfo accessibility.

GIIS Monitor (GStat)

GIIS Monitor collects information published by site GIIS. It includes total CPUs, free CPUs, running jobs and waiting jobs for CEs and used and available storage space for SEs. It also carries out certain tests on sites periodically and publishes their results. The tests include GIIS sanity and performance (response time) checks, RGMA service checks, and checks for services running on CEs and SEs.

LCG-2 Certificate Lifetime

It is a set of scripts that periodically execute an OpenSSL test on Computing Elements (CEs) and Storage Elements (SEs) at various sites to check whether the host certificate is valid.

LCG-2 Job Submission Tests

This is a test run from the GOC system on an hourly basis. It tests both basic globus job submission and RB (Resource Broker) based job submission.

CURRENT IMPLEMENTATION

The current implementation of Gridview features File Transfer Monitoring, Job Status Monitoring, and a prototype Grid Dashboard.

File Transfer Monitoring

The Gridftp logs at various sites are published to R-GMA. Our R-GMA archiver module fetches this data and archives it to database. Our data analysis and summarization module generates periodic (hourly/daily) summaries like averaged throughput and aggregate data movement across the sites. It also generates VO-wise and host-wise detailed summaries. The presentation module presents these summaries and various reports in graphical form using histograms (See Figure 2) and pie charts. In this way, the current version of Gridview summarizes and visualizes live status and history of data transfers across the entire grid. This feature of Gridview is already in production and is being extensively used during the Service Challenge 3 (SC3) throughput tests in July 2005 and January 2006.

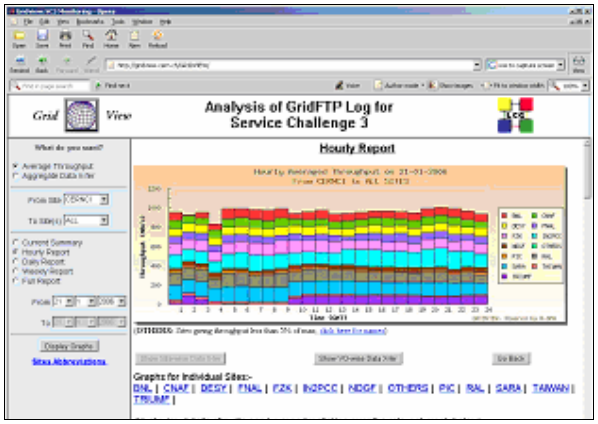

Figure 2: File Transfer Monitoring using Gridview

Job Status Monitoring

The Job Status logs at various Resource Brokers (RB) are published to R-GMA. Our R-GMA archiver module fetches this data and archives it to database. Our data analysis and summarization module generates periodic (hourly/daily) status summaries like total number of jobs in different states at different grid sites, their VO-wise and RB-wise distribution, metrics like site-wise Job failure rate, resource utilization by different VOs etc. The presentation module presents these status summaries and derived metrics in graphical form using histograms and pie charts. In this way, the current version of Gridview summarizes and visualizes live status and history of jobs, their failure rate and resource utilization at various grid sites.

Grid Dashboard (GUI & Visualization Module)

The prototype Grid Dashboard pictorially represents grid sites, their status summaries and fault conditions on the world map. It has different views representing CPU status, storage status, job status and fault status. Quantities are represented using cylinders at appropriate site locations on the map with the height of the cylinder indicating total quantity and different colored portions indicating break-up (See Figure 3). It uses 3-D graphics

and facilitates zooming and rotation to avoid cylinders from some sites blocking the view of others. The current

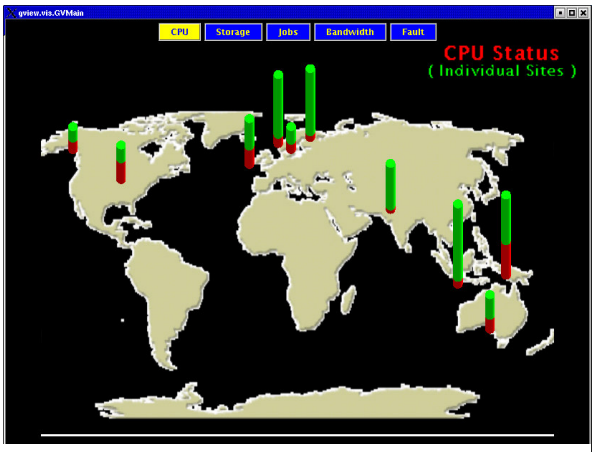

Figure 3: GUI and Visualization Module of Gridview

prototype dashboard implementation works with simulated data, the mechanisms for collecting real status and fault information for grid sites are yet to be in place.

ON-GOING AND FUTURE WORK

Service Availability Monitoring

Gridview is currently being interfaced with Site Functional Tests and GStat for monitoring the availability of various services in the Grid like CE, SE, RB, site BDII, R-GMA, LFC, Myproxy, VOMS etc. Sites will be rated according to their average resource availability and acceptable thresholds for resources like CPUs, storage, network bandwidth etc. Service availability metrics like uptime, failure rate, MTTR (Mean Time To Repair) will be generated and plotted for various services. Different colour coding will be used to represent service status of a site/resource on the dashboard, eg. Green to represent up, Yellow to represent degraded, Red for down and so on. It will provide a service level view for the top management indicating quality of service for sites and whether they are meeting their MoU targets.

Monitoring FTS Statistics

Gridview will also be extended to archive and visualize file transfer statistics like successful transfers, failure rates and so on for various FTS (gLite File Transfer Service) [5] channels across grid sites.

Enhancement of Grid Dashboard

Grid Dashboard, the principle interface and display of Gridview will be enhanced to represent the LCG structure and its hierarchy such as various regions, sites within them and the service nodes in each site. User will have the consolidated view of the entire grid as well as the capability to drill down to a particular resource to view its relevant metrics. Menus will facilitate selection of different views for different user categories (like ROC admins, site admins, VO admins, LCG Managers etc) and metrics of their interests.

CONCLUSIONS

Over the last several months, Gridview has proved to be a useful tool in monitoring the performance of the LCG. During the SC3 tests in July 2005, it was the chief tool in analysing the results of the throughput tests. This has once again been repeated in the rerun of the SC3 throughput tests in January 2006.

The job status monitoring module will be useful to VO admins and site admins to keep track of the workflow statistics for their VOs and sites respectively.

R-GMA is used as the transport mechanism for flow of monitoring information from the sites to Gridview's central database at CERN. Gridview is now one of the main consumers of R-GMA tuples and the Gridview team has contributed to the development of R-GMA in the form of finding and reporting problems and bugs.

REFERENCES

1. LHC Computing Grid (LCG) http://lcg.web.cern.ch/LCG/
2. Gridview Link http://gridview.cern.ch/
3. Relational Grid Monitoring Architecture (R-GMA) http://www.r-gma.org/
4. Site Functional Tests (SFT) http://goc.grid.sinica.edu.tw/gocwiki/Site_Functional_Tests
5. gLite File Transfer Service (FTS) http://egee-jra1-dm.web.cern.ch/egee-jra1-dm/FTS/