



An Evaluation of GridSAM for the Engineering Task Force

Mark Hewitt, North East Regional e-Science Centre
Terry Harmer, Belfast e-Science Centre

23-08-2007

Abstract:

This report focuses on the GridSAM system developed by the London e-Science Centre and distributed by Open Middleware Infrastructure Institute (OMII) at Southampton. This system manages job submission to a variety of resources via a web service interface. The main purpose of this evaluation is to examine the strengths and weaknesses of the system and identify the issues that need to be considered before deploying it in a production environment. GridSAM appears to be an extremely useful tool for large-scale Grid service providers, if only because it provides a standardised interface to the many different platforms and middleware which can exist on the likes of the National Grid Service (NGS). It provides an alternative to the portal style approach, and uses web service standards to implement a standardised interface. However, portals could be written to allow submission via GridSAM. Ongoing work to modularise and simplify the implementation will mean that GridSAM will become an invaluable part of the NGS infrastructure. The Belfast e-Science Centre is currently providing a production deployment of the GridSAM instances for all of the NGS resources. In addition, other GridSAM instances are being deployed at other sites to provide local job submission to NGS resources, which should provide a good testing ground for GridSAM.

UK e-Science Technical Report Series

Report UKeS-2007-06

Available from http://www.nesc.ac.uk/technical_papers/UKeS-2007-06.pdf

Copyright © 2007 The University of Edinburgh. All rights reserved.

1	INTRODUCTION.....	2
2	GENERAL INFORMATION.....	2
2.1	PROVIDER	2
2.2	LICENSING	2
2.3	SUPPORTED PLATFORMS	2
2.4	SUPPORT	3
3	SYSTEMS MANAGEMENT.....	3
3.1	DOCUMENTATION FOR SYSTEM MANAGERS	3
3.2	SERVER DEPLOYMENT	3
3.3	CLIENT DEPLOYMENT	3
3.4	ACCOUNT MANAGEMENT	4
3.5	RELIABILITY	4
3.6	DISTRIBUTED MANAGEMENT	4
3.7	EXTERNAL SERVICES	4
3.8	SCALABILITY	4
4	USER EXPERIENCE.....	4
4.1	DOCUMENTATION FOR USERS	4
4.2	JOINING THE GRID.....	5
4.3	LEGACY APPLICATION INTEGRATION	5
4.4	MIGRATION BETWEEN PLATFORMS	5
4.5	USABILITY	5
4.6	SECURITY FROM THE USER’S PERSPECTIVE	5
4.7	VERIFICATION.....	5
5	DEVELOPER EXPERIENCE	5
5.1	DOCUMENTATION FOR DEVELOPERS.....	5
5.2	LANGUAGES AND TOOL SUPPORT	6
5.3	HOW EASY IS TO DEVELOP NEW SERVICES THAT EXPAND THE GRID’S CORE FUNCTIONALITY?	6
6	TECHNICAL	6
6.1	ARCHITECTURE.....	6
6.2	STANDARDS & SPECIFICATIONS.....	6
6.3	SECURITY.....	6
6.4	INDUSTRIAL SUPPORT	6
6.5	CAPABILITY & FUNCTIONALITY	6
7	BELFAST E-SCIENCE NGS GRIDSAM DEPLOYMENTS	6
7.1	MULTIPLE GRIDSAM INSTANCES IN A CONTAINER.....	7
8	CONCLUSIONS	9
	APPENDIX I	10

1 Introduction

The UK's Engineering Task Force (ETF) is evaluating several Grid middleware solutions in order to take a view on their readiness for deployment on the resources within the National Grid Service (NGS) and those of the wider UK e-Science community. This report focuses on the GridSAM system developed by the London e-Science Centre and distributed by Open Middleware Infrastructure Institute (OMII) at Southampton. This system manages job submission to a variety of resources via a web service interface.

The main purpose of this evaluation is to examine the strengths and weaknesses of the system and identify the issues that need to be considered before deploying it in a production environment.

2 General Information

The GridSAM system is intended both as an add-on for the OMII suite of middleware as well as a standalone system. GridSAM enables compute jobs to be submitted to a connected compute resource using standardised web service interface. The compute job is defined by a job description that conforms to the JSDL job description language standard.

GridSAM allows jobs to be submitted to various resources such as OMII, Globus and Condor. GridSAM's role is to provide a standardised client interface for job submission and data file management.

Its actions are along similar lines to GRAM from the Globus project; however it makes exclusive use of web services technologies.

2.1 Provider

The software is developed by the London e-Science Centre (LeSC) and is distributed as part of the OMII managed programme.

2.2 Licensing

All software produced by OMII (rather than simply repackaged) is produced under a modified BSD licence which allows distribution of the source code and binary forms of the software. (For the text of the modified licence, see Appendix I)

2.3 Supported platforms

The server can run on any Linux platform and the client is currently available for both Windows and Linux. As GridSAM uses a web service interface it should be possible for any architecture to communicate with a GridSAM server deployment. Both machines were used for both client and server. OMII and Condor were running on the same machine at Newcastle, GT4 was running on the same machine at Belfast.

Newcastle Configuration

1 x 1GHz single processor with 1GB RAM.

Belfast Configuration

32 x 2 GHz SunFire Blade System with 1 GB RAM.

GridSAM version 1.1 was installed at both sites. There have been further releases since the evaluation took place but these are not considered in this document.

2.4 Support

GridSAM is available free of charge from OMII. There is currently no commercial support of the package, although support is available from the developers at OMII. However from experience this is limited to the developers only but available to any type of user, be that administrator, installer or user.

3 Systems Management

3.1 Documentation for System Managers

Installation documentation was somewhat fragmented. Full installation requires the user to first install the full OMII stack, and here the documentation is of a good quality and takes you step by step through the process. The GridSAM installation guide is a small part of this documentation, and while it provides sufficient detail to install a GridSAM service the documentation is not extensive.

Some of the commands, such as those required to start and stop the system, referred to an older version of the OMII software.

The installations were carried out using binaries downloaded from the OMII website. Installation from source is possible, but was not attempted in this evaluation. The OMII software itself requires SUSE 9.0, so this was the OS used to avoid any potential compatibility issues.

3.2 Server Deployment

GridSAM was installed at Newcastle together with an OMII stack install, hence it required the same environment as the OMII software itself. The recommended platform was SuSE 9.0, which was used in Newcastle.

The developers had already anticipated Newcastle's problems with port 8080, whereby no application can access this port on campus due to their web-cache policy, so the software runs on port 18080 by default.

The installation process is relatively straightforward when GridSAM is installed at the same time as the OMII stack. However it is rather less simple when a standalone installation is being attempted. There is a general lack of indication as to when the server is running and when it is not.

At BeSC, GridSAM instances were deployed that managed job submission to NGS resources at Leeds, Oxford and Manchester, and other resources at Cardiff, Cambridge and Glasgow. These were hosted under RedHat ES. As above, the installation process was relatively straightforward, and once one service was deployed new services could easily be added. The services were deployed into the OMII container and directly into Tomcat without significant problems.

3.3 Client Deployment

There are two GridSAM clients with the distribution, one for Linux and one for Windows. They both consist of applications which will communicate with the GridSAM server via its web service endpoint. There are various tools for submitting jobs and tracking their progress.

It is not necessary to use the client applications, as these simply provide a more user-friendly way of submitting requests to the GridSAM server.

3.4 Account Management

The job submission required a user certificate in the client. A user's certificate is communicated to the resource being used and enables access control and accounting to be applied at the resource. No accounting control was applied within the hosted GridSAM instance. However, it may be possible to generate such information from the GridSAM or OMII log files themselves.

3.5 Reliability

The software has not been tested under high loads or failure conditions. However, the deployed instances at BeSC have been field deployed for several months and have been used for a significant number of job submissions without a problem. They have not required any maintenance in their deployment and have not crashed in that period.

3.6 Distributed management

There appears to be little backwards compatibility between versions, as the interface for the recently released GridSAM 2 is different from the GridSAM 1 version we tested. GridSAM is not intended to be deployed on every machine in a Grid, but rather to serve as a gateway through which users can submit jobs to a standardised interface which are then passed on to the middleware stack running at the local institution.

OMII are supporting the GridSAM system as part of their 'managed programme', which means updated versions of GridSAM are distributed with releases of the OMII software itself, as well as being available separately.

3.7 External Services

GridSAM is designed to use a standardised web services interface so, theoretically any external system is capable of submitting a job to the GridSAM instance.

Internally, GridSAM is built in a modular structure with separate DRM connectors which allow GridSAM to communicate to a local job submission system. It is possible to write a new DRMConnector for any new middleware system.

3.8 Scalability

We found during our evaluations that a single GridSAM instance is not particularly scaleable when dealing with a heterogeneous environment. Each instance will listen on one port and then take incoming jobs and direct them to one job management system, whether that be OMII, Condor or just using Fork. There is no inbuilt mechanism to either allow a user to specify where the job should be run, or for the system itself to be able to pick between different resources. The GridSAM team have said that this may well be implemented in later versions of the software.

The solution we had to employ to resolve this problem was to deploy several instances of GridSAM, with different endpoints, each of which would submit jobs to different resources and different institutions. However this is clearly not a scaleable solution.

4 User Experience

4.1 Documentation for Users

GridSAM is not intended for end users, and the mechanisms of submission are designed to be hidden behind client software. The documentation is sufficient to install and manage the instances.

4.2 Joining the Grid

Since GridSAM merely provides an interface to an already established grid system, then each system will have its own requirements for joining. However GridSAM can simplify this process by its inbuilt security and certificate mechanisms. GridSAM can of course be configured to only run types of jobs specified by the user, e.g. only jobs from their own VO (Virtual Organisation)

4.3 Legacy Application Integration

GridSAM is designed to submit to any system as long as a DRMConnector can be written for it. Alternative jobs can be run singly using 'fork', however this is not a scaleable solution.

4.4 Migration between platforms

GridSAM is designed to use standardised interfaces for communicating with clients. Integrating it with different middleware applications is somewhat less straightforward, but by no means impossible.

4.5 Usability

This is one area where GridSAM is not as easy as it could be. This is mostly due to the limitations of supporting applications such as Tomcat. Difficulties were experienced in determining the status of the GridSAM server. End users at the client side should not see GridSAM in action, as it should be hidden by a client application: writers of the client applications should see using GridSAM as being just like submitting to any JSDL compliant job execution service.

4.6 Security from the User's Perspective

GridSAM can be deployed to communicate using HTTP and HTTPS. In both cases, however, a user's certificate must be installed in the client application to enable authentication of the user at the resource being used. The issues affecting GridSAM are the issues affecting all deployed web services.

OMII clients require the user to have a certificate and private key in a java keystore, and the password to this stored in a text file. At present this is not permitted under the UK e-Science CA policy, therefore if GridSAM together with OMII were to be deployed on the NGS production service this is something which would need to be looked at very closely.

4.7 Verification

The standard GridSAM client includes a job monitoring tool which will report the status of the job and any errors which occur.

5 Developer Experience

5.1 Documentation for Developers

In attempting to develop an additional component for GridSAM, there is very little (if any) documentation to allow this. All support had to come from the developers.

However, good support is provided by the developers of GridSAM. Newcastle especially received extensive information on the GridSAM architecture, which at the time of writing was not available to the public.

5.2 Languages and Tool Support

GridSAM in its OMII deployment is supplied with a job submission client. However, a deployed GridSAM instance can be used directly using an appropriate (WS-Security enabled) web service call.

5.3 How easy is to develop new services that expand the Grid's core functionality?

It is possible to extend GridSAM to allow it to submit to different middleware applications, but again this is a non-trivial task given the lack of documentation.

6 Technical

GridSAM is packaged as part of an OMII installation and thus relies on OMII for its installation and support. However, BeSC deployed GridSAM within an off-the-shelf Tomcat installation with no problems.

6.1 Architecture

Job submission is a basic resource capability for a grid environment, and a building block for an SOA (Service Oriented Architecture). GridSAM's functionality may be regarded as a basic job submission capability only. There is a need for GridSAM to manage and broker job submission to multiple resources.

6.2 Standards & Specifications

GridSAM implements version 1.0 of JSDL, the web services standard for the submission of jobs.

6.3 Security

The requirement for a client to communicate a user's certificate is a necessary step in usage access and accounting control

6.4 Industrial Support

No industrial support for GridSAM could be identified.

6.5 Capability & Functionality

Job submission is a basic capability that is needed to support usage and management of networked resources. GridSAM provides this basic capabilities and its capabilities (or ones like them) will be needed by NGS and other compute providers. However, the current system is a first step only and it will be required to provide support for multiple resources and the brokering of submission between resources.

7 Belfast e-Science NGS GridSAM Deployments

Belfast e-Science (BeSC) deployed, in Spring 2007, a collection of GridSAM services to provide an interface to the NGS resources. The GridSAM services were hosted in Belfast but provided job submission services to resources in Bristol, Cardiff, Lancaster, Leeds, Manchester, Oxford, RAL and Westminster.

The BeSC hosting environment consists of some 60 servers split across sites in London Docklands, Tadworth and three sites in Belfast. The servers are primarily used as a hosting environment for grid/web services. The preferred hosting environment with the BeSC infrastructure is services hosted within a virtual machine (VM).

Most of the infrastructure within BeSC is 64-bit—this posed an initial problem in that OMII 64-bit deployments were available for Enterprise Edition Operating systems that were not

used with BeSC. However, a 32-bit hosting VM containing OMII services was created to contain GridSAM instances.

Below is the sequence of steps followed in BeSC to deploy GridSAM for NGS Resources using GridSAM 2.0.0 using OMII 3.2.0 with Tomcat 5.0.25.

1. Creating a GridSAM Virtual Machine.

A single GridSAM instance VM was constructed and proved reliable in test deployment.

OS: SuSe 10.1. RAM: 512 Mbytes VM Size: 1 Gbyte.

2. Creating a Multiple Instance GridSAM Virtual Machine.

The single GridSAM VM was deployed with second service¹. In initial tests both services proved reliable. However, the Java VM quickly ran out of memory and GridSAM crashed. Further testing of the single instance VM demonstrated that its JVM size would increase steadily and eventually crash with a memory error also.²

OS: SuSe 10.1. RAM: 512 Mbytes VM Size: 1 Gbyte. Java VM size: 512 Mbytes

3. Creating a host OS GridSAM

The VM approach provides a neat and easily managed solution to deployment. However, it was clear that it would not work for a large number of GridSAM instances. Thus, a 64-bit machine was converted to run a 32-bit OS to provide a hosting environment for GridSAM.

In this OS a collection of GridSAM instances were deployed in a single OMII/Tomcat container. The environment provided a stable environment for running up to 4 GridSAM instances which ran approximately 2 weeks and crashed with a JVM out-of-memory error. This 4-instance behaviour was predicable in several weeks of testing. With more than four instances the JVM would crash within 4 hours for 5 instances and 5 minutes for 8 instances.

OS: SuSe 10.1. RAM: 4 Gbytes. VM Size: 1 Gbyte

4. The Final Configuration.

Given the experience outlined above, the final BeSC deployment for NGS resources consisted of 2 Tomcat containers serving two ports with each containing 4 GridSAM instances. As described above, these operate without problem for approximately 2 weeks without problem but will then crash with out-memory-problems. To prevent this, the hosting server is automatically rebooted once a week in a scheduled 15 minute maintenance period.

OS: SuSe 10.1. RAM: 4 Gbyte. VM Size: 1 Gbyte

7.1 Multiple GridSAM Instances in a Container.

The standard OMII container must be configured carefully to permit multiple GridSAM instances to run correctly.

1. move the files listed below from
../OMII/jakarta-tomcat-5.0.25/webapps/gridsam/WEB-INF/lib to
../OMII/jakarta-tomcat-5.0.25/shared/lib
bc*.jar

¹ See later for instructions on deploying multiple instances.

² BeSC believe there is a memory leak with GridSAM or its supporting libraries, but have no conclusive proof of this.

cryptix*.jar

puretls*.jar

2. make a copy of the

../OMII/jakarta-tomcat-5.0.25/webapps/gridsam/

folder e.g.

../OMII/jakarta-tomcat-5.0.25/webapps/gridsam2/

3. modify the following files to configure the new service:

.../gridsam2/WEB-INF/classes/jobmanager.xml

or which ever configuration file your jobmanager.xml points to configure the remote job manager. (This is outlined in the OMII GridSAM instructions.)

.../gridsam2/WEB-INF/classes/database.xml

to configure database for GridSAM. (This is outlined in the OMII GridSAM instructions.)

...gridsam2/WEB-INF/classes/config/wse/service.properties

to set the service path. (This is outlined in the OMII GridSAM instructions.)

...gridsam2/WEB-INF/classes/config/wse/services.properties

to set the service path. (This is outlined in the OMII GridSAM instructions.)

4. Restart tomcat. (not OMII)

8 Conclusions

GridSAM appears to be an extremely useful tool for large-scale Grid service providers, if only because it provides a standardised interface to the many different platforms and middleware which can exist on the likes of the National Grid Service (NGS).

GridSAM provides an alternative to the portal style approach and uses web service standards to implement a standardised interface. However, portals could be written to allow submission via GridSAM.

Ongoing work to modularise and simplify the implementation will mean that GridSAM will become an invaluable part of the NGS infrastructure.

BeSC is currently providing a production deployment of the GridSAM instances for all of the NGS resources. In addition, other GridSAM instances are being deployed at other sites to provide local job submission to NGS resources which should provide a good testing ground for GridSAM.

Appendix I

The OMII (Open Middleware Infrastructure Institute) distribution comprises open source software from different sources with different licences. All of the software developed by OMII is provided under the Modified BSD licence as given below. Licences for all other software may be found in the top level 'licenses' directory. Copyright (C) 2004-6, OMII, University of Southampton. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of the OMII, University of Southampton nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.