

ICENI Virtual Organisation Management

Asif Saleem Marko Krznarić Steven Newhouse John Darlington

London e-Science Centre, Imperial College London, South Kensington Campus, London SW7 2AZ, UK
Email: iceni@imperial.ac.uk

Abstract

Recent advances in high performance distributed computing have led to the emergence of computational and data grids. The resources within a single organisation are being exposed to other users within a 'virtual organisation' (VO). The VO encompasses a dynamic set of distributed resources, a distributed user base and a distributed management infrastructure. Due to this dynamic nature of VOs, there is a need for an infrastructure to facilitate the management of the constituent users and resources. We have developed an easy to use and secure management infrastructure - the Virtual Organisation Management (VOM) Portal - that provides user authentication and authorisation, resource access control and usage logging services based on existing web and grid standards.

1 Introduction

The emergence of 'Computational Grids' have made provision of high performance and distributed computing power, accessible to the scientific community. These systems have made physically distributed computational, storage, software and networking resources integration possible. These resources can be owned by different physical organisations which are cooperating to build computational communities named 'virtual organisations' (VO) [4]. In VOs, user's and resource's availability is highly dynamic and it is difficult to predict the resource needs and availability. In addition, some sort of access control and usage policies need to be defined and enforced. This is why an infrastructure to manage these VOs is needed. VOM provides user registration using grid certificates, resource access control through grid-map file management (grid-map file maps a user's grid identity to the local grid environment - thereby controlling access to Globus) and resource usage accounting and reporting functionalities.

VOM provides the above services through a portal for remote VO management, a grid service to download and upload information into the VOM database, and client tools to interact with the service through Grid Security Infrastructure (GSI) authenticated network connections.

Within VOM, we have incorporated the latest recommendation of a minimal Usage Record (UR) from the UR-WG of GGF [1] and provided a prototype implementation of the Resource Usage Service (RUS) Grid Service being standardised by RUS-WG within GGF.

The rest of this paper includes sections containing further discussion of the services provided by VOM, details of the design and implementation, some related and existing work and how it relates to VOM and finally details of the work being done currently and it's future directions.

2 Discussion

2.1 User roles and responsibilities

The management of a VO is done by assigning roles to registered users. Whenever users log in to the portal they get access to a restricted set of functionalities based on the role assigned to them by the VO administrator. This is done by first authenticating and then authorising users using their digital certificates.

First users have to register with the VO, once their request is approved by the VO manager, it is forwarded to the managers of the resources. Finally, users can view the resources they have been approved to access and accounting information about their resource usage.

The resource manager's responsibility is to approve user's access to the resources they manage by assigning account names to them. He/She also needs to install clients for resource usage logging and grid-map file management on their resources.

A VO Manager is responsible for enrolling users and resources (alongwith assigning managers to them) into the VO. He/She also allocates users to the resources and view the overall resource usage of the VO.

2.2 User registration

Enrolment into a VO, a collaboration such as an e-Science project that has participants and resources from many real organisations, is split into two phases: authentication and authorisation. Authentication, proof of your identity, is encapsulated within a Globus compatible X.509 public key certificate. One source of such certificates is the UK e-Science Certification Authority being run by the UK Grid Support Centre [2]. Authorisation to join a VO is down to the policy of that VO. The VOM portal provides a secure web-based user registration functionality. Users register with the VO by using their grid certificate, embedded within a standard desktop web browser. This mechanism is used to prove their identity during all interactions with the portal. Once the user is registered, it is upto the administrator of that particular VO to approve or reject the user's registration request based on the information provided.

2.3 Resource access control

The VOM portal also provides the facility to automate the work flow in creating accounts on distributed resources. Each resource needs to install a Java based client, which connects to the web service hosted on the VOM Server, to update the resource's grid-map file.

Once the VO manager has approved a user's registration, local resource managers are notified of the request. It is upto a local resource manager to approve or reject the account creation request on that particular resource. Various account creation strategies [6] can be employed to map the grid identity to a local account. The local resource manager can either decide to create a local user account, or map to a pool of existing accounts, or use account templates [7] [9]. After the creation of this account, the local username is entered into VOM for inclusion in the next update of the grid-map file.

This way grid-map files for all the resources can be managed centrally from a portal. Similarly, a client can connect to multiple VOs and update the grid-map file accordingly.

2.4 Resource usage accounting and reporting

VOM provides facility for resource usage logging, which can be uploaded into the central database hosted at the VOM Server. Users, resource administrators and VO managers can view their usage records restricted based on their access privileges.

GGF Usage Record WG has specified a minimal set of parameters that constitute a usage record. The

22 parameters include job id, project id, CPU usage, Memory usage, Disk usage among others. VOM uses this recommended set of parameters to store and retrieve usage records.

The RUS client running on a resource connects to the RUS grid service to submit the usage record, each time a job is submitted on that particular resource. A log of all the jobs (in XML format) submitted on a particular resource can also be retrieved by the client. A user can view usage records in tabular (Figure 1) and graphical format (Figure 2) through the web interface. Furthermore, higher level services like QoS monitoring, pricing, billing and payment can be developed on top of the existing architecture.

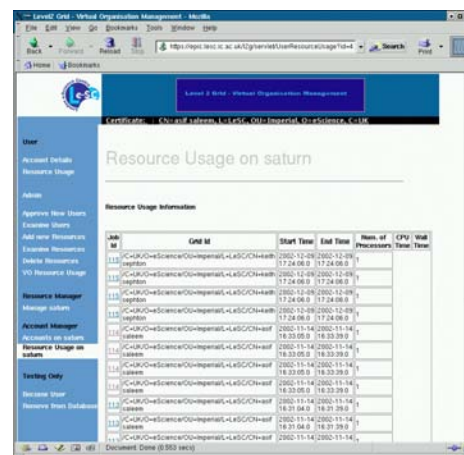


Figure 1: Tabular display of user's resource usage

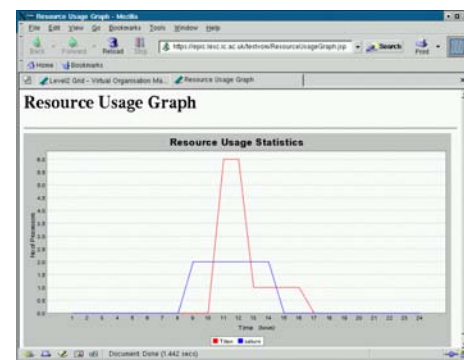


Figure 2: Graphical display of user's resource usage

3 Design and Implementation

3.1 Server Side

VOM portal uses java servlets hosted within a Tomcat servlet container environment. All requests to the VOM portal pass through an Apache server with mod_jk module installed. Apache server forwards

requests to the Tomcat using AJP protocol. We use PostgreSQL database to persist information pertaining to the VO. GT3 Open Grid Services Architecture (OGSA) v1.0 is used for deploying web services in a OGSA customised axis environment. This way we can ensure inter-operability with other grid services and clients can connect to the server using their existing grid credentials.

3.2 Client Side

A client needs to be installed on a resource (currently supports Globus Toolkit version 2.2+) to enable it to access relevant VOM functionality. Clients written in Java communicate with the server using SOAP messages over a GSI connection, which is a secure transport channel. Message level security can also be enabled.

3.2.1 Resource Usage Service (RUS) Client

A RUS client (Figure 3) needs to interact with the local Globus job manager – we have provided extensions for fork and Sun Grid Engine (SGE) job managers. This job manager invokes the RUS client each time a job is submitted so that the usage record generated could be uploaded to the VOM server.

3.2.2 Grid-Map Service Client

Grid-Map Service client (grid-map client) (Figure 4) downloads the latest grid-map file entries available at the VOM server. This process can be automated by setting up a cron job, which runs at specified intervals and updates the resource's grid-map file.

4 Related Work

4.1 EU DataGrid

The EU DataGrid VO management tool [8] uses a Lightweight Directory Access Protocol (LDAP) [10] server to maintain user information. It can be also be extended to store information about applications and resources as in the case of Globus Metacomputing Directory Service (MDS) [3]. LDAP is usually used for read-only type of access and provides a static view of information [5]. Also, it is not the best choice when one wants to view global level data due to inherent inability of LDAP to provide generalised joins [3] as compared to relational database management systems (RDBMS), which can provide relational joins between types. We envisage VOM to not only provide information about users, resources but also provide dynamic information about the current state of the VO e.g. current load on any resource.

This activity is infact complimentary to what MDS already provides.

4.2 GGF Research Groups

Global Grid Forum (GGF) [1] has Usage Record (UR-WG), OGSA Resource Usage Service (RUS-WG) and Grid and Economic Services Architecture (GESAWG) groups working in related areas. Most of their activities are focused on identifying the best practises and standardisation. In VOM, we have incorporated the latest recommendation of a minimal Usage Record from the UR-WG and provided the web service interface of the RUS as recommended by RUS-WG.

5 Current and Future Work

5.1 GT3/OGSA Compatibility

We are working on a new client and update the server to make it compatible with GT3.0 OGSA specifications. This would enable the current grid services to be consumed by other higher level services being developed.

5.2 Policy Management

We are working on support for VO, Resource and User level policy specification and enforcement [11]. At the moment, it is implicit but in order to support large scale VOs it is essential to make it explicit so that users, resources providers and VO managers know what services to provide and expect in return.

6 Conclusion

The VOM portal has now been deployed across the UK e-Science Grid to provide user management and accounting capability (<http://www.grid-support.ac.uk/l2g/>). Users, resource administrators and VO managers can view their own, or their resource's, usage records through a simple access control model. We plan to integrate VOM with the ICENI middleware to provide features such as transient and dynamic grid services registration, querying and instantiation through the web. Also, VOM can be extended to monitoring the usage of other higher-level services.

7 Acknowledgements

The authors want to thank members of the UK Grid Engineering Task Force for testing and evaluating the software. The work has been carried out as part

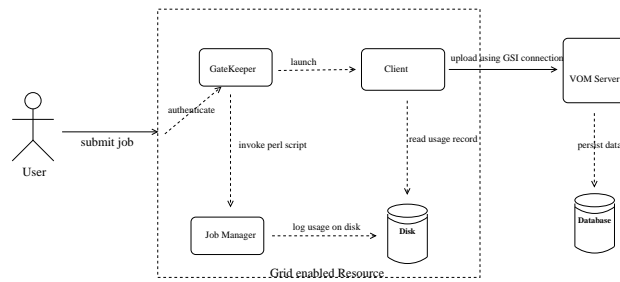


Figure 3: Resource Usage Service Client

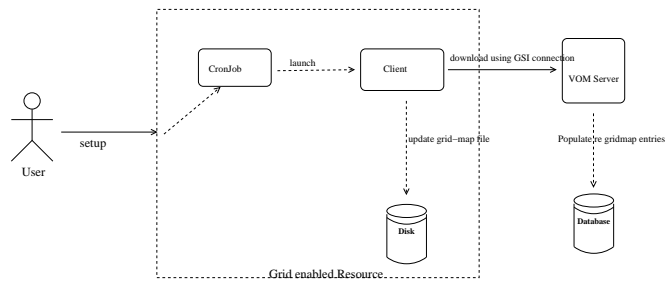


Figure 4: Grid-map Service Client

of OSCAR-G project at London e-Science Centre supported by the Department of Trade and Industry, CompuSys and Intel.

References

- [1] Global grid forum. http://www.ggf.org/L_WG/wg.htm.
- [2] Grid support centre. <http://www.grid-support.ac.uk/>.
- [3] I. Foster, K. Czajkowski, S. Fitzgerald, and C. Kesselman. Grid information services for distributed resource sharing. In *Proceedings of the Tenth IEEE International Symposium on High-Performance Distributed Computing (HPDC-10)*, 2001.
- [4] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the grid: Enabling scalable virtual organisations. *International Journal of Supercomputing Applications*, 2001.
- [5] I. Foster and G. von Laszewski. Usage of ldap in globus. ftp://ftp.globus.org/pub/globus/papers/ldap_in_globus.pdf.
- [6] Thomas J. Hacker and Brian D. Athey. A methodology for account management in grid computing environments. In Craig A. Lee, editor, *Proceedings of Grid Computing - GRID 2001, Second International Workshop, Denver, CO, USA, November 12, 2001*, volume 2242 of *Lecture Notes in Computer Science*, pages 133–144. Springer, 2001.
- [7] Laura F. McGinnis, William Thigpen, and Thomas J. Hacker. Accounting and accountability for distributed and grid systems. In *2nd IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2002), 22-24 May 2002, Berlin, Germany*, pages 284–285. IEEE Computer Society, 2002.
- [8] J. A. Templon and D. Groep. Vo server information. <http://marianne.in2p3.fr/datagrid/documentation/ldap-doc.pdf>, 2002.
- [9] William Thigpen, Thomas J. Hacker, Brian D. Athey, and Laura F. McGinnis. Distributed accounting on the grid. In *Proceedings of the 6th Joint Conference on Information Science, March 8-13, 2002, Research Triangle Park, North Carolina, USA*, pages 1147–1150. JCIS/Association for Intelligent Machinery, 2002.
- [10] M. Wahl, T. Howes, and S. Kille. Lightweight directory access protocol (v3). IETF RFC 2251.
- [11] Glen Wasson and Marty Humphrey. Toward explicit policy management for virtual organizations. In *IEEE Workshop on Policies for Distributed Systems and Networks (POLICY '03)*, 2003.