

# A Proposal of Policy-Based System Architecture for Grid Services Management

Edgar Magaña, Epifanio Salamanca, and Joan Serrat

Universitat Politècnica de Catalunya  
Network Management Group  
Jordi Girona, 1-3 D4-213  
Barcelona, Spain  
{emagana,epi}@nmg.upc.es, serrat@tsc.upc.es

**Abstract.** Communication technologies are evolving along with Grid Computing. The trend of Open Grid Services Architecture (OGSA) is that “everything” is represented as a service, which will be provided to users without any kind of distinctions related to network technology, operative platform, administrative domain, etc. As a result of this trend, there has been a significant management requirement increase in order to cope with the yet elusive rapid and autonomous service creation, deployment, activation and management. This paper describes a proposal to cope with the above mentioned management requirements. Our policy-based management approach enables the establishment, configuration and administration of services on a Grid network that uses the Globus Architecture for Reservation and Allocation (GARA) as an interface for reservation and provisioning of grid resources.

## 1 Introduction

The Grid concept was introduced to denote a proposed form of distributed computing that involves coordinating and sharing computing, application, data, storage, or network resources across dynamic and geographically dispersed organizations. Grid technologies promise to change the way organizations tackle complex computational problems. Nevertheless, Grid computing is an evolving area of computing, where standards and technology are still being developed to enable this new paradigm. The sharing of resources inherent in Grid networks introduces challenging resource management problems due to the fact that many applications need to meet stringent end-to-end performance requirements across multiple computational resources (memory, processor, storage, etc). Furthermore, they should also guarantee fault tolerance and network level Quality-of-Service (QoS) parameters as well as arbitrate conflicting demands.

In order to cope with these resource management challenges, there is a first solution developed for Globus that is called Globus Resource Management Architecture (GRMA) [1]. This solution addresses the wide QoS problem of providing dedicated access to collections of computers in heterogeneous distributed system. But this architecture does not tackle two important issues, namely the

reservation in advance and the resources heterogeneity. Thereby a new solution for resources management was needed. To this aim the Globus Architecture for Reservation and Allocation (GARA) was proposed. GARA extends GRMA in two major ways: The first one introduces the generic resource object, which encompasses network flows, memory blocks, disk blocks and processes. The second one introduces the reservation as a first class entity in the resource manager architecture [1]. Academy and industry have shown special interest on this technology during the last years, thereby several research projects are working on new proposals for improving the management of the grid such as MANTRIP [8] or RNTL e-Toile project [9]. Yang et al [10] presented a management middleware that involves active networks technology and policy-based management. People from Globus are still working on better solutions for management, such as the work presented in [11] related to obtain QoS using the GARA system.

We claim that there is a need for a major entity in charge of providing high-level management to allow quick and autonomous deployment, activation and reservation of services and resources. Additional relevant properties for such entity are fault tolerance and interdomain management support. This paper proposes the fulfillment of the above mentioned goals by means of a policy-based architecture that exploits the facilities offered by the Globus Architecture for Reservation and Allocation (GARA).

The paper is structured as follows. Section 2 analyses the most important management system requirements for the emerging grid services. Section 3 is an overview of the GARA system, which is a fundamental background to understand Section 4 that is related to present the policy-based management proposal and its capability to cater with the requirements mentioned in the second section. Section 5 presents an example scenario. Finally we conclude the paper with some remarks in Section 6.

## 2 Management System Requirements for Grid Services

Grid Computing is affected by continuous innovations and updates due to the huge interest generated among the industry and the scientific community. For this reason a new grid generation is emerging, where schema conversion technologies via common meta-models and ontologies are required to allow data to be moved between storage resources and be shared between tools with different data format requirements. Grid services that automate this process and build on-line libraries for conversion and mapping tools will be of use in a variety of scientific, social and business domains.

At the same time, there are other innovations with a significant impact on Grid Computing, like the creation of mechanisms to handle a large number of heterogeneous resources (computing, storage, networks, services, applications, etc.), the adequate use of data from a wide variety of sources and technical domains, the necessity of designing and implementing new techniques to allow for the interoperability between different data sources is constantly growing.

One of the key challenges will be the co-ordination and orchestration of resources to solve a particular problem or perform a business process. The current generation of grid architecture relies heavily on the program designer or the user to express their requirements in terms of resource usage. Such requirements are usually hardcoded in a program using low-level primitives. But the grid needs to handle resources in a more dynamic way. In other words, grid applications will require the co-ordination and orchestration of grid elements at run time.

Two of the most important grid features are management autonomy and fault-tolerance. This requires redundant resources (computation, data base, network), and also an autonomous and self-regulatory model that ensures the proper working of the management architecture by itself. Maintenance and Scalability are crucial; solutions should be able to support Grids consisting of billions of nodes. This requirement immediately stresses the importance of finding an appropriate trade-off between control systems exerted by users and administrators.

Current solutions for accomplishing these grid management requirements have limitations, particularly related to high-bandwidth, network re-configuration, fault-tolerance, reliability, scalability, flexibility and persistence. Our approach goes one step further offering the possibility to handle the grid management requirements above mentioned by means of an extension of the Policy-Based System designed and implemented as part of the IST FAIN project [3].

### **3 Globus Architecture for Reservation and Allocation (GARA)**

The most well known technology available to reserve and allocate low-level resources such as memory, processing or storage is the Globus Architecture for Reservation and Allocation (GARA) [2]. As mentioned in the introduction, GARA enlarges the Globus resource management architecture basically in two ways: the first one is to generate an architecture with the ability of co-reserving resources and the second one is to generate a generic resource object in order to manage heterogeneous resources.

The GARA architecture allows applications appropriate access to end-to-end Quality of Service (QoS). To do so, it provides a mechanism for making QoS reservation for different types of resources, including disks, computers and networks. GARA provides management for separately administered resources. In Figure 1 we show the GARA's architecture, which consists of three main components, the information service, local resource managers and co-allocation/reservation agents. The information service allows applications to discover resource properties such as current and future availability. Local resource managers have been implemented for a variety of resource types, this explains the use of term "resource manager" rather than the more specific "bandwidth broker", although each one implements reservation, control and monitoring operations for a specific resource. The reservation and allocation agents compute the resource requirements and send the request to the Globus Resource Allocation

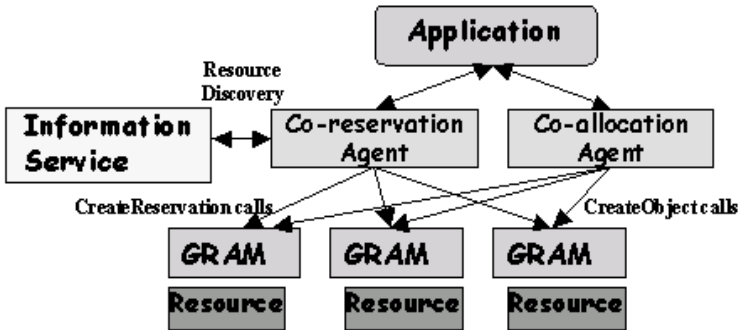


Fig. 1. PBM system with GARA

Manager (GRAM). The GRAM takes the request, authenticates it and if successful forwards it to the local scheduler in order to allocate or serve the resource and finally returns the job handle to GARA. The support for heterogeneous resources is provided by a generic resource object that encapsulates network flows, memory blocks, disk blocks, and even processes.

The support for advanced reservation is allowed by separating the reservation from allocation. For immediate resource allocation, the requested allocation is performed at the time of reservation, but in the case of advance reservation only a reservation handle is returned and the resources need to be reserved at the service start time. A new entity that is called Co-reservation Agent provides this advanced functionality. Its function is similar to the co-allocation agent one, except that after calculating the resource requirement for advanced reservation, it does not allocate but simply reserves the resources. GARA provides a simple, uniform interface for making advance reservations for any type of supported resource. The API provides functions for creating, modifying, binding, claiming, canceling, and monitoring reservations. Although it is as uniform as possible, different resources require different parameters. To accommodate these different needs within a single API, the *create function call* accepts a Resource Specification Language (RSL) string to specify the parameters for single reservation.

## 4 Policy-Based Management Proposal

Policy-Based Networking has attracted significant industry interest in recent years. Presently, it is promoted by the Distributed Management Task Force (DMTF) [5] and is standardized within the Internet Engineering Task Force (IETF) Policy working group [4]. Policy-Based Network Management (PBNM) offers a more flexible, customizable management solution allowing each network node to be configured on the fly, for a specific application tailored for a consumer.

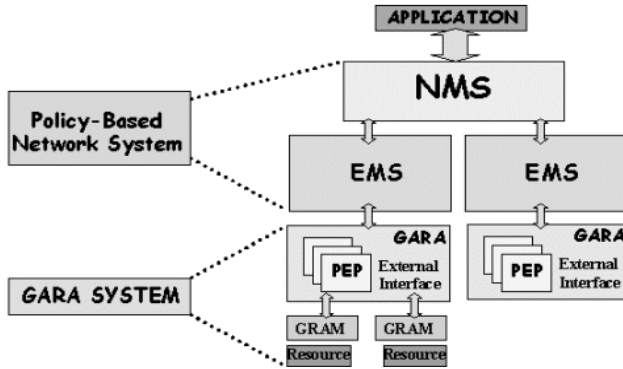


Fig. 2. PBM system with GARA

The Policy-Based Management (PBM)<sup>1</sup> system, presented in this paper, is designed as a hierarchically distributed architecture. Formerly designed to manage Active Networks [3], extensions in terms of management capabilities, have laid an efficient and robust PBM system, which supports co-existence of different management strategies thus facilitating customisation, interoperability with different vendor's equipment, and dynamic extensibility of its functionality to support new services and their deployment as well as new resources on the network.

In comparison with previous traditional network management approaches, PBM offers a more flexible and customisable management solution allowing each managed element (i.e., switch, router, scientific instruments, etc) to be configured on the fly to support customer tailored services. The PBM architecture is able to address a number of domain issues that can be translated to fixed configuration settings. For example, QoS issues often needs complex interactions between relevant network management components. These complex interactions can be easily implemented in the current PBM architecture. Moreover, according to the policy framework, policies are defined or modified by an administrative tool and the intervention of the administrator is just required for special activities. On the other hand is possible to define new types of policies and to extend the system as much as necessary to manage new target elements, without extending the core information model. The proposed PBM system consists of two levels: the Network Management System (NMS) and the Element Management System (EMS), as we show in Figure 2. The NMS is the entry point of the management architecture; therefore it offers the necessary GUIs for controlling the entire architecture. In fact, it is the receiver of policies, resulting from Service Level Agreements between various categories of users. The enforcement of these SLAs requires reconfiguration of the network, which is carried out by means of policies sent to the NMS. These policies, called network-level policies, are processed by the NMS Policy Decision Points (PDPs), which decide when

<sup>1</sup> From here and forward, we will express the Policy-Based Management System just as PBM for simplifying the notation.

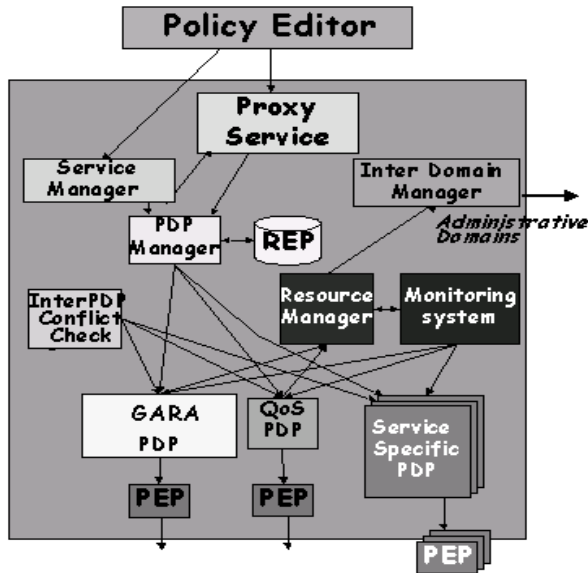


Fig. 3. Policy-Based Management Components

policies can be enforced. When enforced, they are delivered to the NMS Policy Enforcement Points (PEPs) that map them to element level policies, which are, sent to the EMSs. EMS PDPs follow a similar procedure at the element level. Finally, the node PEPs execute the enforcement actions on the managed target elements. The PBM system interaction between components is shown in Figure 3, whilst an in depth description is provided in [7]. The flexibility and scalability of our PBM system proposal makes easy the integration process with GARA. With this goal in mind we have added element-level PEPs within the GARA system, as shown in Figure 3, in order to enforce the element-level policies sent by the EMS PDPs which are derived from policies sent by the NMS. The use of open interfaces in GARA allows all PEPs across the network to share the same view of grid resources, making them platform independent. The new PEPs will contain the necessary code to adapt the parameters from high-level policies to resource-level commands.

## 5 Case Example

The functionality of the entire system is explained in this section. We propose a configuration service scenario because it is one of the simplest cases and therefore it will ease the understanding of the functionalities and specially the management framework advantages. Figure 4 shows the above mentioned example where we distinguish the two main entries to the management framework. The first one comes directly from the application and the second one is from the network administrator Policy GUI. For this example we consider the second one only. Let's

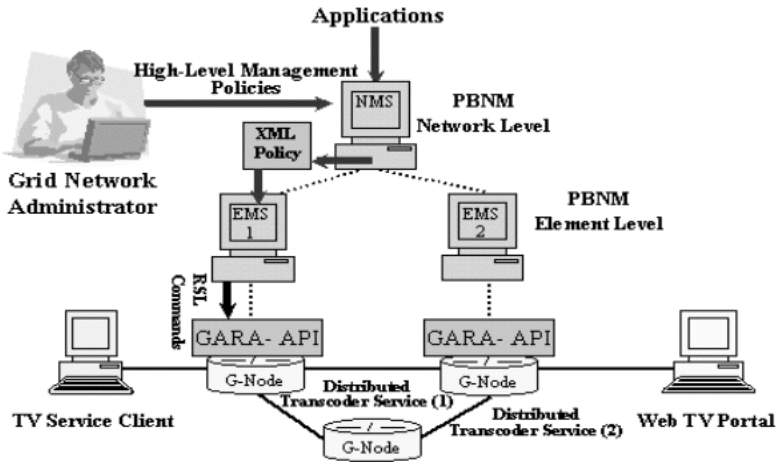


Fig. 4. Example scenario

assume that a TV service client is receiving a data video stream from a Web TV portal (server). As this server offers the video in coded format, clients are using a Distributed Transcoder Service (DTS) on G-Nodes (routers with the Globus Toolkit installed). At a given time, and due to other network requests, the G-Node's resources where our client is attached become unable to continue offering the decoding task with the appropriate quality constraints. The PBM is alerted on that fact by its monitoring system and therefore it starts a resource search task. As soon as the PBM discovers resources available on another G-Node, it generates the necessary resource-level (RSL) commands through GARA's PEPs and contacts with the appropriated interface in order to reserve and allocate the resources for the second DTS (Distributed Transcoder Service (2) in the figure). Once the PBM receives the notification of a correct policy execution, a re-configuration of resources task will be issued on the first G-Node in order to de-allocate all unused resources in that node to be offered to other applications or services. Note that finally the service is running on both G-Nodes without any service interruption as well as without the need of any manual change by the network administrator. The detailed interactions between components inside the PBM system can be found in [6].

## 6 Conclusions

Our proposal consists of an architecture taking advantage of the synergy obtained by coupling policy based management technology and GARA to obtain a full-featured management framework, which fulfills the most important requirements for managing grid services. The scenario illustrates the benefits of the proposed framework in terms of automation, scalability, flexibility and fault tolerance. The proposal is suitable for giving support for dynamic, reconfigurable

on-demand, secure and highly customizable computing storage and networking environments. The next step will be to develop further grid services and the corresponding GARA's PEPs as well as carrying out the appropriate performance testing.

**Acknowledgements.** Part of the work presented in this paper is based on results obtained in the context of the EU IST Project FAIN (IST-10561), partially funded by the Commission of the European Union.

## References

1. K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith, and S. Tuecke. *A resource management architecture for metacomputing systems*. In The 4<sup>th</sup> Workshop on Job Scheduling Strategies for Parallel Processing, 1998.
2. I. Foster, C. Kesselman, C. Lee, R. Lindell, K. Nahrstedt, and A. Roy. *A Distributed Resource Management Architecture that Supports Advance Reservation and Co-Allocation*. In the International Workshop on Quality of Service, June 1999.
3. EU IST FAIN Project Web Site <http://www.ist-fain.org/>
4. IETF Policy Web Site <http://www.ietf.org/html.charters/policy-charter.html>
5. Distributed Management Task Force Web Site <http://www.dmtf.org/>
6. E. Magaña, E. Salamanca, J. Vivero, A. Galis, B. Mathieu, Y. Carlinet, O. Koufopavlou, C. Tsarouchis, C. Kitahara, S. Denazis and J. L. Mañas. *A Policy-Based Management Architecture for Active and Programmable Networks*. IEEE Network Magazine, Special issue on Network Management of Multiservice Multimedia IP-Based Networks, May 2003, Vol. 17 No.3
7. Y. Nikolakis, Edgar Magaña, M. Solarski, A. Tan, E. Salamanca, J. Serrat, C. Brou, and A. Galis. *A Policy-Based Management Architecture for Flexible Service Deployment in Active Networks*. International Working Conference on Active Networks (IWAN03) Kyoto, Japan 2003.
8. EU IST MANTRIP Project Web Site <http://www.solinet-research.com/mantrip/>
9. E-toile Project Web Site <http://www.ens-lyon.fr/LIP/RESO/Projects/Etoile/>
10. K. Yang, A. Galis, C. Todd. *Policy-Based Active Grid Management Architecture*. 10<sup>th</sup> IEEE International Conference on Networks ICON 2002.
11. I. Foster, A. Roy, and V. Sander. *A Quality of Service Srchitecture that Combines Resource Reservation and Application Adaptation*. In 8th International Workshop on Quality of Service (IWQoS '00), June 2000.
12. R. Buyya. *Nimrod-G An Architecture for a Resource Management and Scheduling System in a Global Computational Grid*. 4<sup>th</sup> International Conference on High Performance Computing in Asia Pacific Region (HPCAsia 2000) IEEE CS Press, Los Alamitos, USA 2000.