Virtualizing Resources: Customer-Oriented Cross-Domain Monitoring for Service Grids

Timo Baur
Munich Network Management Team
Leibniz Supercomputing Center
Department for Communication Networks
85748 Garching, Germany
timo.baur@lrz.de

Sameh Saad
Munich Network Management Team
University of Federal Armed Forces Munich
Information System Laboratory (IIS)
85577 Neubiberg, Germany
saad@informatik.unibw-muenchen.de

***Abstract*— After Computational Grids have lain the foundations for the functionalities and acceptance of Grid infrastructures, Service Grids may now set a new paradigm in their development and construction. We present an organizational model using different roles for each service layer, and discuss the possibility to integrate customer-oriented service chains, located in exclusively delegated virtual service domains, and delivered by multiple providers. This short paper briefly introduces an approach to realizing dynamic customer-oriented resource monitoring for Grids, and describes an ongoing effort to compose virtual monitoring instruments for individual Grid users.**

I. Introduction

Service Grids set a new paradigm in the development and construction of Grid infrastructures. While Computational Grids focused mainly on increasing computational power by providing dependable, consistent, pervasive and inexpensive access to high-end computational capabilities [3], Service Grids now provide a comprehensive base for a distributed, decentralized and loosely coupled [9] provisioning of a much bigger variety of services, functionalities and underlying resources. Meanwhile, a customer-oriented Grid monitoring is still mostly based on tight relations between the customer and a local Grid specialist at a resource providers site.

With increasing numbers of users and resource providers, a monitoring of the resource states and service parameters for failures, configuration, accounting, performance and security should be handled in a provider-independent way, tailored to the needs of individual users. Current systems inadequately support the distinction between a customers and a providers point of view. Thus, methods to monitor only the delegated services about the clusters at the resource providers are required.

As a prerequisite, section II applies a conceptual framework to the services of a large-scale heterogenous Grid infrastructure that considers the roles of customers, users and providers as well as the involved provisioning domains. It outlines a loosely coupled and dynamically composed service provisioning chain from the resources to the users. The dependency of services and resources to virtual organizations (VOs) is sketched in section III and methods to provide and describe customer-oriented views are described. A generic method of cumulative virtualization and grouping (or aggregation) of service functionality, resource state and information is presented. Section IV introduces ongoing work at a prototype implementation. Finally, we conclude with related work and results.

II. Conceptual Framework

The D-Grid initiative, one of the largest Grid projects in Germany integrates the Grid infrastructures of multiple interorganizational scientific communities. Every community brings along its own applications based on diverse Grid technologies, provided by different resource providers at multiple sites. These communities are organized as VOs that use own resources or resources from different resource providers in the Grid. When Meteorologists in the Collaborative Climate Community (C-3) build earth system models from satelite data on their "Grid Infrastructure Data and Processing Grid" which is a part of the D-Grid, they develop own applications on top of a shared Grid service infrastructure. This infrastructure enables them to use shared functionalities like job execution and storage. It also supports management, monitoring and distribution of the functionalities and states of tasks, data and resources. As multiple organizations use the same Grid infrastructure, a structured and defined monitoring process, provided as a service, is necessary, that is able to measure, transport and present the information for specific customers in different communities.

Understanding a Grid as a service, four layers of subservice provisioning can be perceived: Application services, realizing the communities applications. Grid services, enabling a Grids functionalities by encapsulating and composing the states of multiple tasks and data from heterogenous middleware services. Middleware services, providing appropriate instrumentation and encapsulation of resources and enabling their management and monitoring. And low level resource services like e.g. local job scheduling systems, cluster-monitoring, or local or regional network management systems.

A. Organizational Model

The mode of management and cooperation of all services on top of a Grid middleware is peer-to-peer. But, often overlooked, this is accompanied by hierarchical dependencies of the administrative and functional areas, in which the services are located. An organizational model of these areas can be achieved by describing roles and actors for the involved entities, according to organizational aspects, and group them into domains [7]. Administrative domains embed and structure administrative dependencies (see figure 1). They are orthogonal to functional domains, created by functional and technical dependencies.

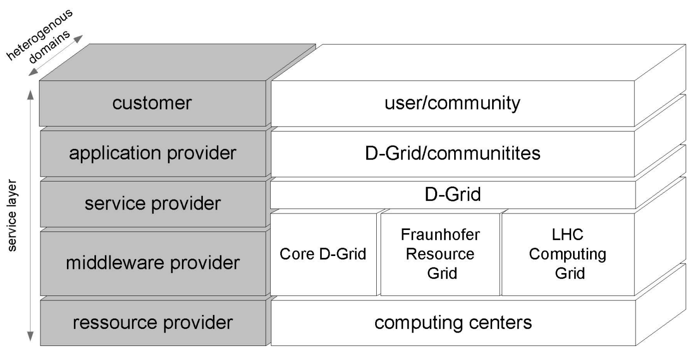

Fig. 1. roles of services and examples for administrative domains

As mentioned above, the roles around every service layer can be provider, customer or user. Analogous to the involved entities resource, middleware, Grid and application service, different providers and customers can be perceived and provider roles like resource provider, middleware provider, Grid service provider and application provider can be defined. The same can be applied for the roles of the customers and users. Every entity can inhibit more than one role and implement more than one layer.

B. Service Chaining

Today, users often plan and operate a Grid directly with the resource providers. A customer-oriented Grid service monitoring which is independent of the monitoring of individual resource providers shouls solely depend on leased or allocated services and resources of a given VO. A consistent provisioning of monitoring information from different sites of resource providers to independent customers as a compound service is therefore necessary. This can be achieved by first decoupling resource providers and customers using a Grid service layer and considering the complete provisioning process from resources in heterogenous domains to the customers. From a top-down view, for example, the community in the scenario acts as an application service provider for their users. Its Grid application software needs Grid services as subservices and the community as a VO can agree with multiple Grid service providers upon the used Grid services and their service level (SLA). The Grid service providers in turn can agree with middleware providers, which in turn can agree with resource providers about the use and instrumentation of their resources.

Figure 2 shows this manifold provider hierarchy as a service chain (as e.g. described in [4]) for the complete process of provisioning a Grid resource from the Leibniz Supercomputing Centre LRZ to a community, traversing the service layers until a defined basic services on a resource, considering all involved roles.

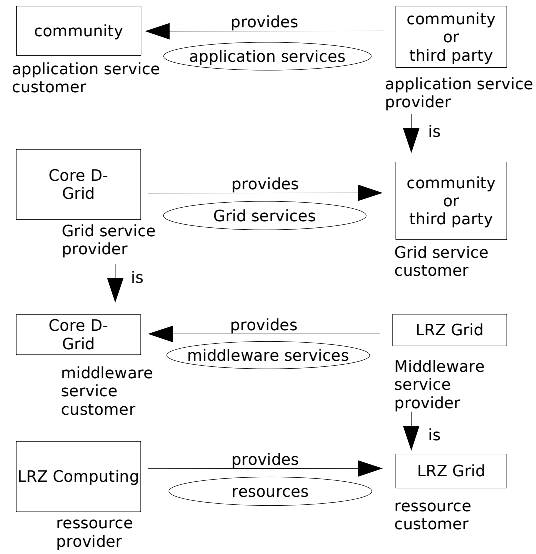

Fig. 2. service chain for the provisioning of an application service

C. Mediation

As in heterogenous domains at first policies and management architectures may be different, monitoring is only possible in a homogenous domain at one layer. Nevertheless, from the view of a customer application at an application layer, a monitoring of all leased services and resources must be possible, even when a compound service is cross-provisioned from heterogenous domains.

Users should be able to use uniform interfaces for the retrieval and manipulation of data from the top layer. The Grid service layer should provide all necessary functionalities without user queries to deeper layers or phone calls to a resource operator. This layer should act as a mediating bridge (as defined in [6]) between domains in subsequent layers to provide the information a customer needs.

This can be achieved by providing proxying or virtualization functionality for underlying services. For example, archive services can be used, that cache the information from the underlying middleware or resource layer. Also, repository or lookup services can act as directories to keep references to all information providers on a Grids service or middleware layer. Using this mechanism, in theory all necessary information between resource and Grid service provider can be mediated. Drill downs to the root cause of a problem can be applied directly at the Grid service layer, which reports the underlying situation.

To ensure an individual providers policies, this mediation process should be policy-mediated. It should furthermore dynamically shape the service provisioning chains by respecting the SLAs and the delegations of a users VO. We call this process VO-mediation. It can be achieved by coupling the proxying or virtualizing service to VO-management mechanisms.

Such a coupling implements a service logic (as specified in [4]) and realizes a tight coupling of a resources lifecycle, a (virtual) service and a VO management, as demanded in [8].

Explicitly coupled services are ”VO-aware”. VO-Mediation can also happen implicitly, e.g. when a dedicated monitoring repository is instantiated for a specific group of people, when access rights are restricted to specific resources and services (e.g. by a gridmap-file) or when an end-to-end-link is established for a specific user.

Given VO-aware and correctly instrumented services, customer-oriented service chains can be established dynamically and controlled according to policies and SLAs of different providers, users and VOs in every layer of a Grid.

III. Virtual Instrumentation

When a service chain for the monitoring workflow is established vertically across provisioning layers or domains to a customers virtual service, a repeated instrumentation of the monitored resource is applied involving a communication system. The instrumentation steps in every layer establish associations between the specific underlying resource and its representation. This applies a virtualization process as outlined in [5] where virtualization is described as a set of transformation processes during which associations between virtualized entities and ’underlyings’ are established and changed. A virtualization layer internally maintains or indirects these associations. As it acts as a virtually homogenous domain that exchanges the information, it could also be called a virtual domain.

By defining metrics to measure the state of resources and the quality of services (QoS) and by specifying service access points, underlying services are transformed across layers into virtual services in a virtual domain which enables them to share common characteristics. They can also change their virtual states interdependently, if some of their underlying resources or services are connected.

We thus define a virtual domain *VD* as a set of functionality or information provided at the service access points of virtual services that map (or indirect to) the functionalities and information available at the service access points of a subset of related services S_1,...,S_n in other domains D_1,...,D_n by the use of an underlying communication system of stateful transmitters.

Examples for such virtual domains are virtual environments, virtual machines, Grid middlewares, Grid services or management objects.

To realize a customer-specific multi-layer virtualization, the virtualization process must be divided into two more granular processes. The first process, service virtualization, maps a specific resource state or service functionality as well as specific QoS-parameters into such a virtual domain, facilitating its creation and usage. In a second step, service organization, this functionality and information is grouped or composed for a next iteration of the process, choosing a specific set of virtual services.

The process of service virtualization implies tunnel couplings between domains D_1,...D_n by the use of a domain *VD* as a tunnel. As long as a mapping to the original information or functionality is preserved, the QoS-parameters as well as the functionality of the original services can be referenced or proxied. The second process of service organization applies a grouping, delegation, composition or aggregation of the resulting set of virtualized services to choose a specific functionality or QoS-parameters for a customer and his given user. In this way, a shaping of the virtualization process and the resulting virtual domains is achieved.

If this is done according to the requirements and SLAs of a VO as customer, the previously mentioned coupling of a VO-management implementation realizes a VO-aware service logic for VO-mediation.

A sequential and repeated application of this process through a multi-layered service chain then results in dynamic and dedicated customer-oriented services. Applying the process at least twice, e.g. by first instrumenting resources with services and then services with virtual services, a two-phase hierarchical mapping is realized, that can be called a virtual instrument (as also, e.g., in [1]). Figure 3 shows how this process can be applied to create two different customer specific monitoring environments.

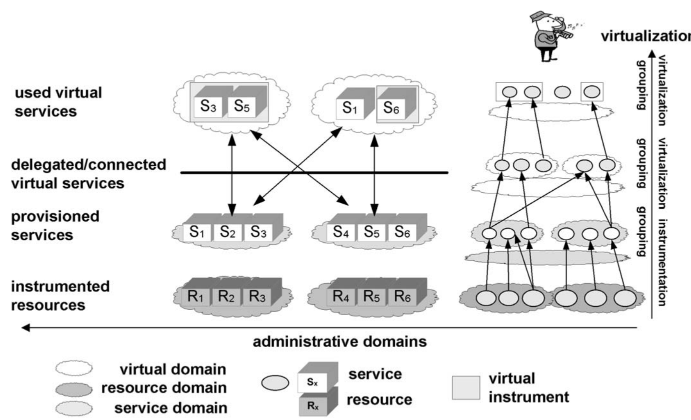

Fig. 3. paths between resources, virtual domains, services and users

As an example, two Grid users monitor two supercomputing resources at two different sites. These are connected by the same virtual middleware domain with information about nodes R_x in both clusters (service virtualization). Nodes are grouped or delegated (service organization) according to the VOs SLAs and the required usage functionality of the user. Given two VOs have a portfolio of nodes S_3,S_5 (VO A) and S_1,S_6 (VO B) provided at a VO-authenticated middleware access point and referenced in two VO-specific or VO-aware lookup-services. The user in VO A may then connect his local graphical user interface (his own virtual domain), which receives and replicates the information about the services (service virtualization) and choose the nodes of interest S_3 and S_5 for his monitoring task (service organization) while another user in the VO B is monitoring S_6.

The QoS-parameters of the used services and the status information of the used resources are propagated along the service chain to the virtual service, where a customer-specific aggregation of the parameters can take place. The quality of the virtual services (QoVS) is thereby determined by the quality of the virtual domain or environment of a given VO as well as by the qualities (QoS) of the set of subservices constituting it. These in turn are repeatedly determined by the Quality of their underlying subservices in the service chain until the quality of a basic service, device or resource is defined at the resource providers sites.

IV. PROOF-OF-CONCEPT

For the monitoring in the scenario Grid, three metadata archives (based on the monitoring systems of the middleware Globus Toolkit, gLite and UNICORE) have been implemented for the Core D-Grid. A hierarchy of underlying Grid monitoring services has been structured according to administrative dependencies and middleware-related requirements. They act as virtual domains providing middleware-related resource information. For example, the service based on Globus provides a mapping to URIs (Uniform Resource Identificators) of available Grid services. It also caches dynamic information about the status of Globus resources using the GLUE schema. Two of the indexes have been integrated in the World Wide Web where monitoring information is provided to the users and customers of the D-Grid.

It is possible to include more archives for particular VOs and their monitoring of the Grid service layer. The system is able to act as a base for further integration efforts between multiple middlewares and user interfaces. The prototype will be extended to realize VO-aware information services, defining specific formations of resources, services and actors. Figure 4 shows as example three views from a dynamic user interface, viewing delegated services for given users, VOs as well as QoS-Information.

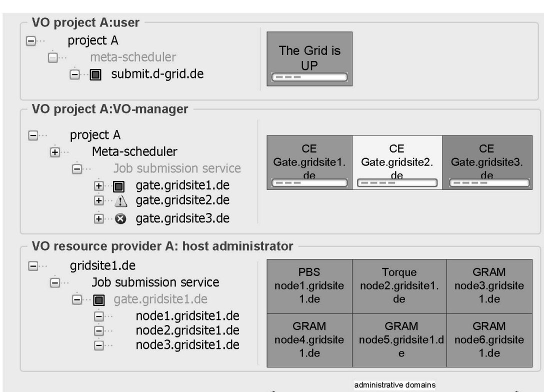

Fig. 4. three examples for VO-sensitive monitoring views

V. RELATED WORK AND RESULTS

There are many areas, where virtual domains have been proposed: the authors in [10] outlined the relation between physical and virtual network resources as a process of abstraction and partitioning. [11] introduces virtual links as an abstraction that relates to a partition of a virtual resource being used by a given virtual network, as well as a virtual network inheritance tree, including links to users. The necessity of an extension of service management, the development of virtualization concepts that consider the dynamicity of organizational structure and services and to automate VO formation workflows was mentioned in [2]. Motivated by this, we have developed a concept of VO-mediated service chains, service virtualization and virtual domains and deducted dependencies of Quality of Service parameters for virtual services. The concept helped us to understand how to monitor heterogenous resources throughout various Grid layers and customer-oriented and how community-dependent views on their status can be achieved.

A shortcoming of the approach is that it does not include event handling. Furthermore, its user-orientation does not consider a direct monitoring at resource layer.

For our next steps, we will work on the realization of this concept as a prototype for Customer Service Management (CSM) which will enable Grid users to manage their delegated resources and services and to visualize their individual monitoring and accounting data.

ACKNOWLEDGEMENT

The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful discussions and valuable comments on previous versions of this paper. Parts of this work have been funded by the German Federal Ministry of Education and Research under contract 01 AK 800 B and by the EC IST-EMANICS Network of Excellence #26854.

REFERENCES

[1] J. Chen. Virtual instrumentation interface for srrc control system, 1995.

[2] G. Dreo Rodosek, H.-G. Hegering, and B. Stiller. Dynamic virtual organizations as enablers for managed invisible grids. In *Proceedings of the 2006 IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Vancouver, Canada, April 2006. .

[3] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the grid: Enabling scalable virtual organizations. In *International Journal of Supercomputer Applications*, volume 15, 2001.

[4] M. Garschhammer, R. Hauck, H. G. Hegering, B. Kempter, I. Radisic, H. Roelle, and H. Schmidt. A case-driven methodology for applying the mnm service model. In *Proceedings of the 2002 8th International IEEE/IFIP Network Operations and Management Symposium*, 2002.

[5] S. Graupner, R. Koenig, V. Machiraju, J. Pruyne, A. Sahai, and A. van Moorsel. Hpl-2003-125: Impact of virtualization on management systems. Technical report, 2003.

[6] Object Management Group. Common object request broker architecture (corba), v 3.0.3, 2004.

[7] H.G. Hegering, S. Abeck, and B. Neumair. *Integriertes Management vernetzter Systeme: Konzepte, Architekturen und deren betrieblicher Einsatz*. dpunkt-Verlag, 1999.

[8] W. Hommel and M. Schiffers. Supporting virtual organization life cycle management by dynamic federated user provisioning. In *Proceedings of the 13th Workshop of the HP OVUA*, 2006.

[9] I. M. Llorente, R. S. Montero, and E. Huedo. A loosely coupled vision for computational grids. In *IEEE Distributed Systems Online, vol. 6 no. 5*, 2005.

[10] W. Szeto, Y. Iraqi, and R. Boutaba. A multi-commodity flow based approach to virtual network resource allocation, 2003.

[11] Daniel Villela, Andrew T. Campbell, and John Vicente. Virtuosity: Programmable resource management for spawning networks. *Computer Networks*, 36(1), 2001.