

Fine-grained Continuous Usage Control of Service based Grids – The GridTrust Approach

Syed Naqvi¹, Philippe Massonet¹, Benjamin Aziz², Alvaro Arenas², Fabio Martinelli³, Paolo Mori³, Lorenzo Blasi⁴, Giovanni Cortese⁵

¹ Centre of Excellence in Information and Communication Technologies (CETIC), Belgium
{syed.naqvi, philippe.massonet}@cetic.be

² e-Science Centre, STFC Rutherford Appleton Laboratory, United Kingdom
{b.aziz, a.e.arenas}@rl.ac.uk

³ CNR Institute of Informatics and Telematics, Italy
{fabio.martinelli, paolo.mori}@iit.cnr.it

⁴ Hewlett Packard Italiana S.r.l., Italy
lorenzo.blasi@hp.com

⁵ Interplay Software S.r.l., Italy
g.cortese@ipsoft.it

Abstract. Access control techniques designed for single domain infrastructures, where users are known by domain administrators, provide considerable liberty in the usage of resources. This paradigm is not suitable for highly scalable and decentralised systems such as Grids and service oriented architectures (SOA), where resources are shared between domains, and users come from remote domains. One approach is to provide policy-driven autonomic solutions that operate a continuous monitoring of the usage of resources by users. This paper presents the services and tools offered by the GridTrust Security Framework (GSF). GSF addresses three layers of the next generation of grid (NGG) architecture: the Grid application layer, the Grid service middleware layer, and the Grid foundation layer. The framework is composed of security and trust services and tools provided at the middleware and Grid foundation middleware layers. Various business case studies are being developed to validate the GridTrust results.

Keywords: Grid technology, usage control, service security, trust infrastructure

1 Introduction

The *Service-based Grids* [Foster] have the ability to provide scalable and low cost service based infrastructures for both business and scientific purposes. They provide language- and platform-independent techniques for describing, discovering, invoking and orchestrating collections of distributed computational services, thus facilitating the development of complex wide-area applications [Gounaris]. However, from the security point of view, they introduce important challenges because the pool of

resources and users are dynamic and managed by different administrative domains. Current access control technology in Grids only provides coarse grained security – i.e. user enjoys unlimited privilege of using the resource once he got access to it. The GridTrust consortium argues that coarse grained access control leaves Grids inherently vulnerable, and that not only the access to a resource needs to be controlled, but also the usage that is made of the resource. This paper presents the GridTrust framework that introduces fine grained and continuous usage control in Grids, and provides the necessary services, tools and methods to deploy it in service-based (OGSA compliant) Grids.

2 Current State of the art in Grid Security

The native authorization system of Globus, the gridmap one, is too simple to satisfy the requirements of a cooperative distributed environment such as the Grid one. Hence, this section describes some attempts to enhance Globus security by integrating external authorization systems.

The Community Authorization Service, CAS, has been proposed by the Globus team [Pearlman]. CAS is a service that stores a database of VO policies that determine the actions that each Grid user can perform as member of the VO. A Grid user that wants to access a Grid service, requests to the CAS service a credential to access this service. The CAS returns a credential embedding a CAS policy assertion, and this credential is presented by the Grid user to the service he wants to exploit. This approach requires that Grid services are able to understand and enforce the policies embedded in the CAS credential, and these policies are coarse-grained, because they only define which of the local services can be accessed by the Grid user.

An alternative solution, described in [Thompson], integrates the Akenti authorization system in Globus. Akenti is an authorization system exploiting X.509 certificates for user identity and distributed digitally signed authorization policy certificates for access decisions. Once the user has been authenticated, the system retrieves the policies for each resource referred in the user request, and matches them with the user's credentials, that include the attributes assigned by the VO to that user. This solution is a pure pull model in which the user capabilities are collected after his authentication.

The solution presented in [Stell], instead, exploits PERMIS, which is a role-based access control infrastructure using X.509 certificates to define users' roles. All access control decisions are driven by an authorization policy that is stored in a X.509 certificate too. PERMIS supports classical hierarchical RBAC, in which roles are assigned to users and privileges on resources are paired with roles.

The Virtual Organization Membership Service (VOMS) [Alfieri] is another advanced authorization system for Globus. In VOMS a VO has a hierarchical structure with groups and subgroups; a user in a VO is characterized by a set of attributes, 3-tuples of the form *group, role, capability*. The combined values of all these 3-tuples form a unique attribute, the Fully Qualified Attribute Name (FQAN). A user contacts one or more VOMS server in order to obtain the authorization information granted by a VO to him. To access a Grid service the user creates a proxy

certificate containing the information received from the VOMS Servers. To perform the authorization process the information is extracted from the user's proxy and combined with the local policy. The resource providers periodically query VOMS database to generate a list of VO users and map them to local accounts.

However, none of the previous systems performs fine-grained controls, i.e. controls the actions performed during the access. Moreover these models define static rights, because they depend on credentials that can be modified only by administrative actions, and these rights are evaluated only before granting the access, and no further controls are executed while the access is in progress.

Recently, Sandhu et al [Sandhu04] defined a conceptual model, called usage control (UCON), based on the concepts of mutable attributes and continuity of policy enforcement. In [Sandhu06] they propose the adoption of UCON in collaborative computing systems, such as the Grid. A preliminary attempt to adopt this model in Grid has been made in [Martinelli05], and this is the model that we adopt in the GridTrust framework [Martinelli07].

3 The GridTrust Framework

One of the outcomes of the GridTrust framework is the development of a set of online trust and security services and a set of policy modeling, analysis and transformation tools. The objective of developing these tools is to facilitate the development of rigorous Grid-based applications with enhanced security.

3.1 Virtual Organisation (VO) Model

In order to support rapid formation of VOs, we use the concept of virtual breeding environment (VBE) [Camarihna-Matos]. A VBE can be defined as an association of organisations adhering to common operating principles and infrastructure with the main objective of participating in potential VOs. We have adopted the view that organisations participating in a VO are selected from a VBE, as illustrated in Figure 1. Such organisations may provide resources/services (ovals), and include users that utilise VO resources (small squares).

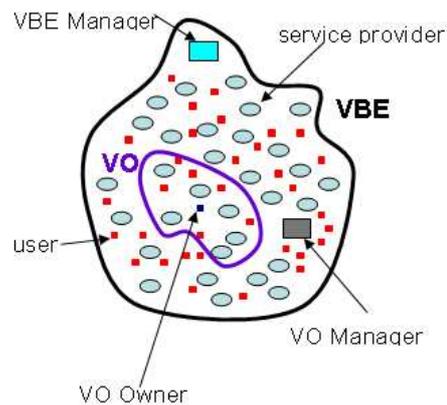


Fig. 1. Organisations and users in VBE

Organisations pre-register to a VBE via a VBE Manager component, including description of the resources they are willing to share in a Grid and the list of potential users belonging to the organisation. When a user requests to create a VO, s/he

assumes the role of VO Owner and contacts a VO Manager with the description of needed resources. The VO Manager is in charge of selecting potential providers and setting up the VO to operation.

3.2 Framework Services

We describe here the different trust and security services that have been developed under the GridTrust framework.

- **The VBE Manager (VBEM)** has the main functionality of a service registry, where service providers register their services and other GSF services can retrieve them given abstract service descriptions. Each Virtual Organization (VO) is created within a specific Virtual Breeding Environment (VBE); and a VBE may contain several different VOs.
- **The VO Manager Service (VOM)** coordinates all the other security services and is the single point of access for users and service providers participating in the VO. The VO Manager is responsible for handling several functionalities. These include VO creation, populating VOs with services required by VO owners to achieve their goals, updating VO policies, evolving the VO by allowing its service providers to subcontract part of their services to other service providers and finally, terminating the VO.
- **The Policy and Profile Manager (PPM)** keeps all the knowledge bases needed by GSF services, namely: VBE and VOs users, with security preferences and their trust and reputation credentials; VOs with their owner and security policies; service providers with their services and the fine-grained security policies regulating access and usage of the services.
- **The Secure Resource Broker (SRB)** is called by VOM with a list of services, needed by the VO Owner to form its VO, and the associated security requirements. It returns the list of providers offering the requested services and also satisfying all the specified security requirements. One of those requirements is the reputation of a service in a VBE.
- **The Trust and Reputation Service (TR)** keeps track of the past and current behavior of VO owners, users and service providers and transforms it into trust and reputation credentials that can be considered by other users, service providers and GSF services when making decisions.
- **The Continuous Usage Control Service (C-UCON)** is an implementation of the UCON policy framework [Sandhu04], where it is deployed on each service provider and is responsible for the evaluation and runtime enforcement of policies about resource usage in VOs. It interacts with the TR service to get the current reputation of users and it also reports feedback to the TR service about users

violating UCON policies.

Each of these services can be invoked only by mean of the API it exports, hiding all the implementation details on how the service is implemented. The framework is modular so it allows the possibility of adding future new security services if needed.

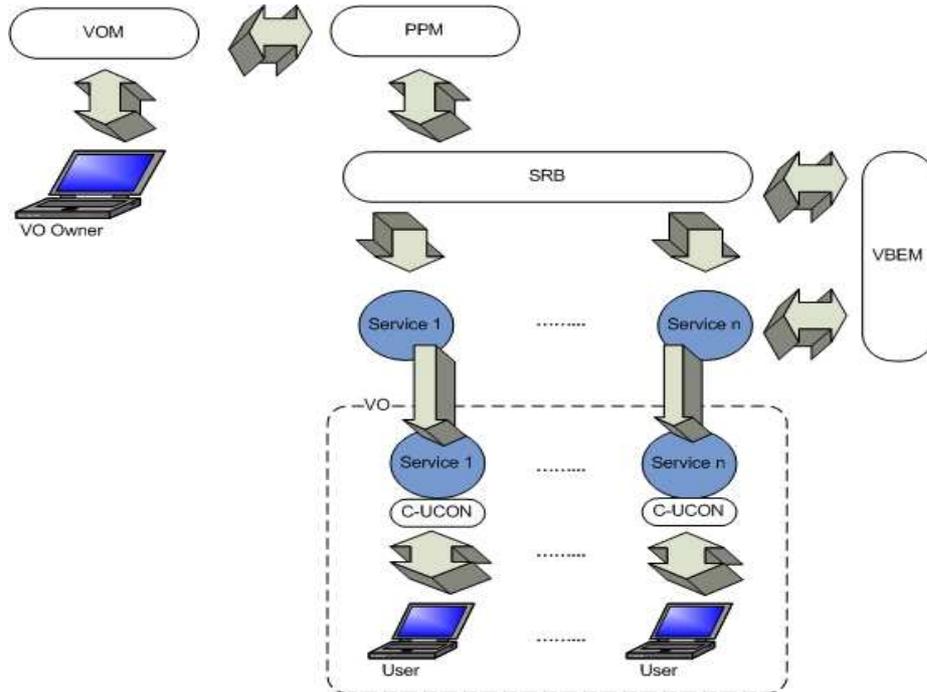


Fig. 2. Establishing a VO using GridTrust Services.

Figure 2 shows the different interactions among the GridTrust services when establishing and running a VO. This is done over several phases. In the first phase, a user (henceforth called the VO owner) requests from the VOM service the creation of a VO. In the next phase, the owner registers with the PPM service (through the VOM) the list of VO users and their security profiles. Then the owner requests from the VOM to search for suitable business services fulfilling its workflow by including the abstract description of each business service and any security requirements it must satisfy (e.g. minimum reputation level). The VOM utilizes the SRB service in its search and once the right candidates are reported back to the owner, the owner informs the VOM of its selection and SRB negotiates and schedules the selected candidates. The VO is now fully operational. Each of the computational resources underlying the business services is protected by a local instance of the C-UCON service, which monitors the users' behavior on those resources.

3.3 Policy Tools

These include a set of tools that are being developed to aid designers and analysts for policy writing during the design phases of the application development.

3.3.1. The Policy Requirements-to-Design Tool

The policy requirements-to-design tool facilitates linking of the security policies expressed in KAOS goal-oriented requirements model [vanLamsweerde] to the operational specifications of those policies. These policies are expressed in a formal process algebraic language POLPA [Martinelli06] [Martinelli07] [Aziz]. The tool eventually builds up a library of policies that comply with the specified UCON requirements.

3.3.2. The Policy Refinement Tool

The aim behind the policy refinement tool is to allow policy designers to write VO-level policies using the stakeholders' alphabet and then refine it, in a correct and automatic manner, to a resource-level policy written using the resources' computational alphabet. The language at both levels is based on POLPA; however, the tool allows the designers to automatically derive low-level policies for VO resources. The tool also composes the refined VO-level policies with existing pre-VO policies at the resource level.

3.3.3. The VO Modeling and Animation Tool

The VO modeling and animation tool generates Grid-based VO models in a formal refinement-based manner and then animates changes in the behavior of VOs when changes in the VO policies take place. We follow a formal approach that is based on the Event-B refinement language, where we envisage that the tool will develop interfaces for the Rodin modeling tool and the Pro-B animation tool.

3.3 Implementation status

The first version of the GridTrust Framework has already been implemented and most of its source is available under the Apache 2.0 license. A public demonstration will be provided during the ICT 2008 conference in Lyon. To access further GridTrust publications and the current software release please refer to www.gridtrust.eu.

4 Validation Scenarios for GridTrust Framework

4.1 Distributed Content Management Case Study

The case study aims at researching and showcasing dynamic access and usage control mechanisms, similar to those outlined in [Sandhu06], for applications implemented in a Grid architecture.

The application outlined by this scenario is a *general purpose, workflow-enabled content management tool*, which supports a distributed organization in the execution of collaborative projects with the following characteristics:

- They aim at the production of some complex, sophisticated ‘digital’ product (e.g. a software system, or some multimedia product).
- They are ‘knowledge-intensive’ and ‘content-intensive’. Workers depend on and need access to several sources of knowledge as well as digital content assets, which they assemble / use to create the product. This need must be supported by appropriate search facilities.
- The production process is structured along some workflow (e.g. a software production process, or a web / content publishing process), and foresees several phases. Policies which control access to these assets may vary according to the phase or state in the project workflow.

The application (a ‘VO’ in GridTrust terminology) offers access to a virtual content management (‘CM VO’) infrastructure, made out of several application servers, where users can: a) create a repository or collaborative ‘workspace’ where content can be stored b) upload content to such workspace c) search and retrieve content. Content managed through the infrastructure includes unstructured documents as well as multimedia content.

In VO Creation phase, the CM VO application discovers and registers application servers providing the actual CM services, thus creating the content management infrastructure.

In VO Usage phase, users access CM services while the GridTrust usage control infrastructure enforces appropriate access and usage control. The case study addresses two perspectives related to usage control: resource usage and collaboration. Overall, it aims at covering several of the types of usage control policies mentioned in [Sandhu06]

4.1.1. Resource usage

The CM VO allows on-demand provisioning of a content management infrastructure (See [Alfresco Cluster in the Cloud] for a similar scenario).

Users of the VO create a workspace, where they can store and share content with their partners, using the VO resources. The GridTrust infrastructure must ensure that users use VO resources in a fair and controlled way. To guarantee availability of resources and performance to all users of the CM VO, thresholds on the usage of resources must be enforced. For example:

- User can create and own at a given point in time in the CM VO only ‘max_spaces’ spaces and ‘max_content’ content items, occupying ‘max_disk_space’ i.e. the sum of the space requested to host the content objects owned by the user
- User can only perform a given number of queries in a time interval (e.g. in a minute/ hour/ day...). Queries can require non-trivial system resources, especially if they match a large number of content objects.
- Users can download only a given number of contents in a given time period

Note that access/ usage control may be needed both at VO level and node level.

4.1.2. Collaboration

The CM VO allows controlled sharing of such content resources among several users and organizations, which require traditional access control mechanisms. It also provides content workflow capabilities, hence should allow restricting or otherwise customizing the access to content / documents to specific users based on context. Types of policies we research in this case study include several dynamic, history-based access and usage control scenarios:

- Status of Shared Objects - access to granted based on the status of a content in a workflow
- Dynamic Separation of duties

Implementation of the application is work-in-progress. The CM VO is being implemented as a portal where VO users register and get services. Individual nodes providing content storage, indexing and query capabilities are implemented as Globus services interfacing to a JSR-170 Content Repository.

4.2 Supply Chain Case Study

The proposed Supply Chain scenario is based on two main ideas. The first is to use an auctioning system exploiting competition between transporters and allowing customers to find the best provider for each task. The second idea is to have route computing services, i.e. computational services providing maps and libraries to execute applications that solve the logistic optimization problem, to allow even SME transporters to optimize their routing. The routing computing service and possibly also the auctioning system are hosted on Grid resources.

Vigor, a pharmaceutical company, receives an order from a hospital. Vigor's warehouse has enough supply of the required goods, so only a transporter is needed to ship the order and satisfy the customer. Vigor's procurement system creates a Request for Quotation for the required transportation task specifying source, destination, expected arrival time, volume, weight and type of the goods and sends it to the Auctioning Service, thus creating an auction for it.

Celer, a transportation company, gets notified of the RfQ and its Automatic Quotation System analyzes the auction terms versus the company's policies and current availability of resources to determine if it's worth bidding or not. After a positive bidding decision the Automatic Quotation System needs to calculate the cost of delivering given the current resource engagement. For this calculation a job is sent for execution to the Grid Routing Service. The computation considers all the pending transportation tasks and time / capacity constraints for the Celer's fleet and optimizes (recalculates) the whole set of routes, one for each vehicle, to compute the incremental cost of executing the required transportation task. From the result an offer is created, with an estimated time for delivery, and a bid sent to the Auctioning Service. Choice of the best offer can be based on price, planned delivery time and transporter's reputation, depending on proponent's requirements.

To give a size to the scenario imagine a small group of 10 producers that create an auction for each of their 50 daily transportation tasks, and a group of 30 transporters that bid on every auction. Even for this small group it is 500 auctions per day (nearly one every minute in working hours), spawning 15.000 jobs of routing optimization every day. If the group of participant actors is not a small one the number of NP-complete problems to be solved in a single day may raise to several millions.

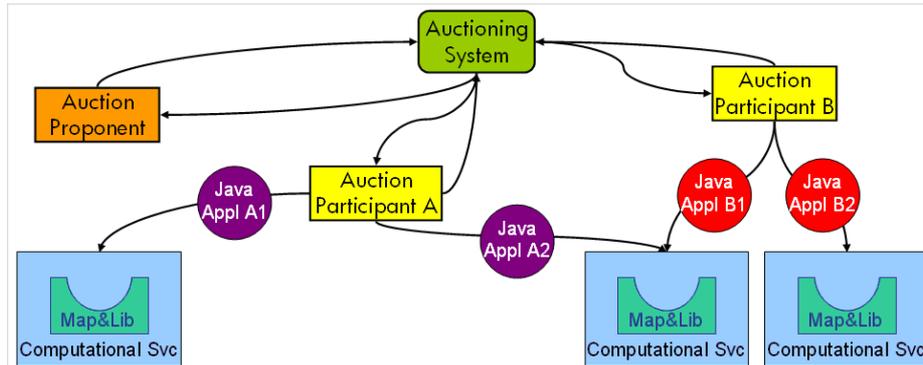


Fig. 3. Supply chain scenario's actors.

The components of the business scenario are the following (see Fig. 3):

- Auctioning system, a custom service running reverse auctions of type First-Price Sealed-Bid, allows producers to propose Requests for Quotations (RfQ))
- Auction Proponent, it's the Producer application (creates auction, receives result, creates Delivery VO)
- Auction Participant, is the i -th Transporter application (notified of an auction, creates Routing VO, invokes routing calculation, sends an auctioning offer)
- Map&Lib, are Routing support services, (maps, map access library, base routing functionality) made available by the Computational Service provider
- Java Appl, is the Routing application (executed on Computational Service) which may be different for each Auction Participant

The resources shared between domains are the Computational Routing Services, each hosting Distance-Time Matrices (maps) and providing sophisticated algorithms for solving the given Operating Research problem (e.g. VRPTW, see [Gambardella]).

This solution raises several security challenges such as selection of services with compatible security policies or continuous control of the execution of unknown applications, among others

A future implementation of the system will allow monitoring the whole delivery phase and verifying transporter's compliance with the offered terms of service (considering the offer as a SLA). The reputation index of a transporter is based on a history of its accomplishments; with the current implementation it lowers if the transporter's behavior is not in line with VO security policies, but using a SLA monitor the reputation can be increased with successful shipments and lowered if the transporter doesn't fully comply with the terms agreed in the SLA.

4.3 Benefits of the GridTrust Framework

SRB is useful to clients as it allows to find Computing Services with compatible policies and granting access to the needed libraries.

UCON benefits Service providers in that continuously controls that the unknown code running on their servers is not violating policies or even executing harmful operations; UCON benefits clients too, because each executing application and data are protected against intrusion of other clients' applications.

TR is most useful to Service Providers by providing a measure of users' reputation, but may be useful to clients too when they want to base their partner's choice on the reputation index in addition to other parameters such as cost or performance.

VOM and VBEM at the end are the essential coordinators of the whole GridTrust Security infrastructure.

5 Discussions

Trust and security are fundamental issues in Grid, because of its collaborative and highly distributed nature. Grid users belong to distinct administrative domains that adopt distinct security mechanisms and have different security policies. These users are typically unknown, and no trust relations may exist among them. Hence, sharing resources in the Grid could be dangerous, because unknown and untrusted users could execute dangerous or even malicious applications on them. Another important issue is that accesses to services could be long-lived, i.e. could last hours or even days, and users' permissions may depend on conditions which are mutable over time.

The authorization systems adopted in Grid so far do not address all these issues. These authorization systems simply decide whether to allow a given user to access a service. No further controls are executed on the actions performed by the applications executed by remote Grid users on the local resource. Otherwise in GridTrust both VO Owners and Service Providers can define fine-grained security policies which are then continuously enforced by the GridTrust C-UCON Service.

The GridTrust project is innovative because it addresses the main security issues of Grid environments by proposing an integrated framework that provides a set of services performing the main Grid interactions in a secure way. These tools allow the Grid participant to create and manage VOs, to select resource providers having certain security requirements, to manage users' reputation, and to execute applications on behalf of remote Grid users while performing a fine-grained and continuous monitoring of computational services according to the UCON model.

Hence, a main advantage of the GridTrust framework is that it is not a simple authorization system, but consists of a set of services enhancing the security of the whole Grid lifecycle, from VO formation to VO dissolution.

Another interesting feature of the GridTrust framework is that all the components have been developed as Globus services and have been integrated in the Globus environment. Hence, the GridTrust framework could be adopted in current Grid nodes built on Globus with minor modifications. For the same reason, the GridTrust components could be easily integrated with other Globus based (security) services.

6 Conclusions and Perspectives

The GridTrust framework addresses the security and trust requirements of service-based Grids. In this paper, we have presented its approach for fine-grained continuous usage control of Grid resources. The Continuous Usage Control service of the GridTrust framework controls the usage of Grid's computational resources by applying fine-grained and history-based access control, and improves state of the art with mutable attributes, obligations and continuous enforcement. The GridTrust framework features fine grained monitoring of the actions performed by applications on the resources. The history of these actions is used in the evaluation of new requests. The access rights are therefore dynamic in GridTrust framework because attributes and conditions may change over time. The two presented use cases confirm GridTrust framework's workability.

Acknowledgements

The research leading to the results presented in this paper has received funding from the European Union's sixth framework programme under grant agreement number FP6-033817.

References

- [Alfieri] R. Alfieri, R. Cecchini, V. Ciaschini, L. dell Agnello, A. Frohner, A. Gianoli, K. Lorente, F. Spataro: *VOMS: An Authorisation System for Virtual Organizations*. Proceedings of 1st European Across Grid Conference, 2003
- [Alfresco Cluster in the Cloud] <http://ihatecubicle.blogspot.com/2008/05/alfresco-cluster-in-compute-cloud.html>
- [Aziz] B. Aziz, A. Arenas, F. Martinelli, I. Matteucci and P. Mori, *Controlling Usage in Business Processes Workflows through Fine-Grained Security Policies*, In Proceedings of the 5th International Conference on Trust, Privacy & Security in Digital Business (TrustBus 2008), Turin, Italy, 01-05 Sep 2008, Lecture Notes in Computer Science, Vol. 5185, Springer, 2008.
- [Camarinha-Matos] L.M. Camarinha-Matos, and H. Afsarmanesh, H. *Elements of a base VE infrastructure*. Journal of Computers in Industry, 51(2):139–163, 2003. (available at <http://www.uninova.pt/~cam/ev/CiI.PDF>)
- [Foster] Foster I., Kesselman C., Nick J., Tuecke S., *Grid Services for Distributed System Integration*, IEEE Computer, 35(6), pp 37-46, 2002
- [Gambardella] L. M. Gambardella, E. Taillard, G. Agazzi, *MACS-VRPTW: A Multiple Ant Colony System for Vehicle Routing Problems with Time Windows*, in New Ideas

in Optimization, D. Corne, M. Dorigo and F. Glover (eds), 63-76, McGraw-Hill, London, 1999

[Gounaris] Gounaris A., Paton N., Sakellariou R., Fernandes A., Smith J., Watson P., *Modular Adaptive Query Processing for Service-Based Grids* CoreGRID Technical Report TR-0076, March 2007

[Martinelli05] F. Martinelli, P. Mori, A. Vaccarelli: *Towards Continuous Usage Control on Grid Computational Services*. Proceedings of Joint International Conference on Autonomic and Autonomous Systems and International Conference on Networking and Services (ICAS-ICNS 2005), IEEE Computer Society, 2005, 82

[Martinelli06] F. Martinelli, P. Mori and A. Vaccarelli: *Fine Grained Access Control for Computational Services*. Technical Report Number TR-06/2006, Istituto di Informatica e Telematica, Consiglio Nazionale delle Ricerche, Pisa, 2006.

[Martinelli07] F. Martinelli and P. Mori: *A Model for Usage Control in GRID Systems*, in Proceedings of the First International Workshop on Security, Trust and Privacy in Grid Systems (GRID-STP07), 2007.

[Pearlman] L. Pearlman, C. Kesselman, V. Welch, I Foster, S. Tuecke: *The Community Authorization Service: Status and Future*. Proceedings of Computing in High Energy and Nuclear Physics (CHEP03), 2003

[Sandhu04] R. Sandhu, J. Park: *The UCON_{ABC} usage control model*. ACM Transactions on Information and System Security, 7(1), 2004, 128-174

[Sandhu06] Zhang, X., Nakae, M., Covington, M. J., and Sandhu, R. 2006: *A usage-based authorization framework for collaborative computing systems*. In Proceedings of the Eleventh ACM Symposium on Access Control Models and Technologies (Lake Tahoe, California, USA, June 07 - 09, 2006). SACMAT '06. ACM, New York, NY, 180-189. DOI= <http://doi.acm.org/10.1145/1133058.1133084>

[Stell] A. J. Stell, R. O. Sinnott, J. P. Watt: *Comparison of Advanced Authorisation Infrastructures for Grid Computing*. Proceedings of High Performance Computing System and Applications 2005, HPCS, 2005, 195-201

[Thompson] M.R. Thompson, A. Essiari, K. Keahey, V. Welch, S. Lang, B. Liu: *Fine-Grained Authorization for job and resource management using Akenti and the Globus toolkit*. Proceedings of Computing in High Energy and Nuclear Physics (CHEP03), 2003

[vanLamsweerde] van Lamsweerde, *Requirements Engineering in the Year 00: A Research Perspective*, in Proceedings of the 22nd International Conference on Software Engineering, Limerick, Ireland, ACM, pp. 5-19, 2000. (available at <http://www.sis.uncc.edu/~seoklee/teaching/Papers/lamsweerde00requirements.pdf>)