

Policy-driven Negotiation for Authorization in the Grid

Ionut Constandache
Duke University, USA
ionut@cs.duke.edu

Daniel Olmedilla
L3S Research Center, Germany
olmedilla@L3S.de

Frank Siebenlist
Argonne National Laboratory, USA
franks@mcs.anl.gov

Abstract

In many Grid services deployments, the clients and servers reside in different administrative domains. Hence, there is a requirement both to discover each other's authorization policy, in order to be able to present the right assertions that allow access, and to reveal as little as possible of the access policy details to unauthorized parties. This paper describes a mechanism where the client and servers are semantically annotated with policies that protect their resources. These annotations specify both constraints and capabilities that are used during a negotiation to reason about and communicate the need to see certain credentials from the other party and to determine whether requested credentials can be obtained and revealed. The result of the negotiation is a state where both parties have satisfied their policy constraints for a subsequent interaction or where such interaction is disallowed by either or both. Furthermore, we present an implementation of a prototype, based on the PEERTRUST policy language, and a reasoning engine that is integrated in the Web services runtime component of the Globus Toolkit. The negotiation process is facilitated through the implementation of WSRF-compliant service interfaces for protocol message exchanges.

1 Introduction

Organizations join in collaborations for the benefit of sharing resources for intensive computations, data storage and query. Each organization is defined by its own administrative domain, while the overlaying collaboration, known also as the virtual organization (VO), is defined by a mutual agreement. Grid toolkits provide the middleware to securely share resources in different domains. Essentially such an infrastructure should provide to clients seamless access to all resources they are authorized to use. However, in current Grids, users have to hardcode into their job programs the details of gathering and providing appropriate credentials in order for access to be granted at each domain. Because authentication along with single sign-on can

be automated by client delegation of X.509 Proxy Certificates to the job being submitted, most authorization mechanisms are being simplified to an identity based scheme, so that they can be accommodated on the same infrastructure as authentication. As a consequence, due to the potential large number of users and different certification authorities, identity based authorization leads to scalability problems.

To address these issues, we exploit annotation of Grid services with machine understandable languages and bind rule-based policies to resources in the context of authentication and authorization. These policies specify access control requirements that must be satisfied by a requester before access is authorized. In addition, we claim that authorization cannot rely on parties' identities alone, because such authorization does not scale. Authorization is, therefore, no longer a one-time mechanism, but an iterative process, in which the entities involved must be able to negotiate and incrementally increase their level of trust based on other party's properties. The application of negotiation-driven authorization mechanisms to Grid environments overcomes current authorization limitations. Furthermore, policy based negotiations provide scalability, advanced access control, and automatic credential fetching. These features enable a complete set of new scenarios in the context of authorization in which user involvement is dramatically reduced (from administrative and development points of view) in favor of automated interactions.

This paper extends the work performed in [6], identifies current limitations in Grid authorization, suggests the application of policy languages to protect access to resources, and demonstrates the power and advantages of this approach. In addition, we propose an architecture in which policy languages can be easily integrated into the Globus Toolkit 4.0, allowing advanced access control and automatic credential fetching. We also describe how this architecture has been implemented by using the PEERTRUST policy language for policy specification.

2 Grid-specific Authentication and Authorization

The Globus Toolkit [13] offers a collection of software that implements different protocols, specifications, standards, and interfaces to meet the requirements identified by the Open Grid Services Architecture (OGSA) [29] for supporting a Grid infrastructure. Since its inception in the late 1990s, the Globus Toolkit has addressed issues such as resource discovery, remote execution monitoring and management, data movement and replication, and security in service-oriented distributed environments. In this section we provide a short introduction to the Globus Toolkit 4.0 (GT4.0), focusing on its integrated Grid Security Infrastructure (GSI) [15].

GSI provides integrity protection, confidentiality, and authentication for sensitive information passed over the network, as well as the facilities to securely traverse the different organizations that are part of a collaboration. The fundamental security mechanism used by the GSI infrastructure relies on public key cryptography and the associated PKIX X.509-based Public Key Infrastructure (PKI) [16]. GSI employs X.509 identity certificate to bind a public key to a unique name (called Distinguished Name, or DN) identifying the private key holder. A trusted certification authority (CA) guarantees through its signing policy that the two belong together and attests to this by signing the certificate.

GSI supports a delegation mechanism through the use of X.509 Proxy Certificates [33]. The user can delegate his rights to agents, intermediaries, or resources such that they can interact with other resources on the user's behalf. The mechanism consists of the agent generating a new key pair and the user signing and issuing a proxy certificate that holds the agent's public key. A proxy certificate can be signed by an end entity certificate or by another proxy certificate. During authentication the whole chain of (proxy) certificates is exchanged between the parties, allowing them to obtain the user identity information from the end entity certificate at the root of the chain.¹

GT4.0 provides an authorization framework for enforcing authorization on both service and client side. On the service side, authorization mechanisms rely on a configured chain of policy decision points (PDPs) to determine whether authorization should be granted or denied for a client invocation. The decision regarding authorization is based on the conjunction of all PDP decisions; that is, authorization is granted only if all PDPs have returned a permit decision. The PDP's decision logic can be implemented based on the client's DN, the resource accessed, and the operation invoked. On the client side, service authorization options are *self* (identity of the service and client are expected to be

¹The user's identity information is authenticated if that user's end entity certificate is signed by a trusted CA.

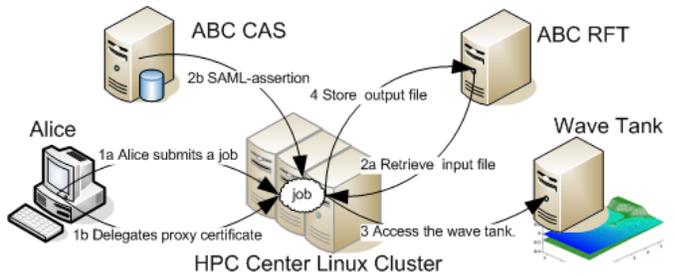


Figure 1. Simple Grid Scenario

the same), *host* (a certain host name is expected of the service accessed), or *identity* (a certain identity of the service is expected where identity refers to the DN present in service certificate).

On the service side, a number of PDP implementations for different authorization options are part of the GT4.0 distribution. These PDPs are configured either at the service level or for the entire GT4.0 container running the service. The options include *self*, *host*, and *identity* as well as *gridmap* (the identity of the client must be mapped to a local user account on the resource in a gridmapfile), *samlCallout* (a SAML authorization callout to an external OGSA Authorization compliant service), and *userName* (username- and password-based authorization).

3 Example Scenario: Interaction between Jobs and Services in GT4.0

Let us consider the following illustrative scenario depicted in Figure 1. Alice, a scientist at University "ABC" (ABC), would like to obtain data regarding oceanic water waves for developing signaling instruments for high waves hazards avoidance. Fortunately, the university and the local Navy Institute have an agreement that allows ABC members to use any of the institute's scientific instruments as long as they are not already in use and ABC has not exceeded its maximum number of monthly allocated hours for institute's resources usage. Both the university and the Navy Institute support the Globus Toolkit 4.0, which provides the common interoperable middleware layer to collect data from a wave tank available at the institute site.

Alice's client program submits a job over the Grid infrastructure to ABC HPC Center Linux cluster. There the job will retrieve the input data required to setup the wave tank, perform CPU-intensive computations, refine the output data and store the results. However, first, the Linux cluster has to authenticate and authorize the submitted job. Alice's client program authenticates using Alice's X.509 End Entity Certificate signed by ABC Certification Authority trusted at the

Linux cluster site. Further, authorization is granted as a grid-mapfile entry maps Alice’s DN to a local user account allowing the job to start with the privileges of the associated local user.

Since the job needs additional resources, a new X509 Proxy Certificate is delegated by Alice to ABC Linux cluster. Using this proxy certificate, the job can authenticate and gain access to other Grid resources. For example, the job will contact ABC Reliable File Transfer Service (RFT) [3] and will retrieve the input data after authentication and authorization are fulfilled based on Alice’s DN at the root of the certificate chain.

Alice is aware of the Navy Institute policy requiring proof of university affiliation before allowing access, so she has coded into the job the retrieval of a SAML statement from the ABC Community Authorization Service(CAS) [22] attesting her involvement with the university. After the required authentication and identity-based authorization with ABC CAS, the job-processing application obtains the assertion and creates a new X.509 Proxy Certificate embedding it. This proxy certificate will ensure access at the Navy Institute wave tank site.²

When the data is ready, the Linux cluster contacts the ABC RFT service, and after another round of authentication and authorization verification, the final results are made available to Alice at the university file server.

4 Current Limitations and Assumptions

During the interactions depicted in the previous scenario, authorization decisions were based on user identity. The user already had an account at the remote locations with a grid-mapfile entry mapping her X.509 certificate’s subject’s DN to her local account. Another assumption was that the ABC CA was trusted at each location the job needed access to (ABC resources and Navy Institute wave tank). There was no intermediary step for the job and resources to argue about trusted authorities. The trusted CAs of each site were expected to be known by the client, who was further expected to provide an appropriate certificate. In these conditions, mapping client identities to local accounts raises serious scalability problems because of the large number of potential users. Even more unrealistic is the requirement of having a single trusted Certification Authority when the Grid spans multiple organizations each having its own authentication infrastructure.

Managing identity-based access control lists requires site administrators to keep track of all possible clients allowed to request services. This is a complicated task because the user mapping entries are the result of the user’s associations with certain organizations and projects. Once such relations

²Further we assume that the wave tank is not in use and ABC has not exceeded its allocated hours, so that all Navy Institute policies are satisfied.

cease to exist, the resource administrators have to remove all these entries to ensure that access is no longer allowed. On the other hand, as soon as new users are required to deploy the Grid infrastructure, new mapping entries have to be entered in all relevant access lists.

Keeping track of all the different credentials, in relation to the resources and the domains where they are needed, is a burden on the client side as well, because of the difficulty in predicting the job’s needs prior to its execution. When the job is expected to present a required and missing credential, the user may not be available. Alternatively, providing the job with all the user’s credentials might imply disclosing sensitive information, a situation that the user may not be comfortable with. The issue for the user or job is to provide just the right credentials to the Grid service it tries to access—no more, no less.

In all the previously described interactions, the authorization decision was a one-step process, without the client asking about the resource access requirements and without any negotiation for granting authorization. The client was simply expected to be aware of and be able to satisfy the service requirements. This approach seems infeasible for large-scale Grid deployments, where each of the computational resources may be under a different authority and where entities with no previous interactions may have to communicate.

Some of the deployed Grids authorize users based on assertions issued by infrastructure services established at the virtual organization level (e.g., CAS). However, Grid service providers may feel uncomfortable relying only on remote assertions, especially in the case of a large Grid environment. As a result, some service providers may use, publish, and enforce policies in the local domain such that the administrators keep fine-grained access control over the resources they provide.³

We argue in the following sections that when Grid deployments support the specification, advertisement, and enforcement of service-level access control policies, together with capabilities for the automatic fetching of credentials, large-scale collection of resources can indeed be supported. Such a mechanism will also enable dynamic negotiation for authorization and access granting based on the properties of both parties.

5 Rule-based Policy Languages

Rule-based policy languages have been used extensively in the context of security policies, trust management, and business rules [5]. The reason researchers have typically

³Even if authorization statements are made by third parties at the VO level, the authorization decision may involve local information and state (e.g., the wave tank should not be in use, and the ABC use of instruments should not have exceeded a certain limit).

selected rule-based languages to express policies is not arbitrary. Their semantics is closer to the way humans think and is especially useful for access control protection where one specifies *what* conditions are to be fulfilled by the requester, without specifying *how*. Furthermore, rule-based policies provide self-described statements that can be exchanged, shared, and reused among parties and that allow one to reason over them, hence enabling interoperability.

Security in distributed environments such as the World Wide Web, peer-to-peer systems, or Grids was traditionally built under the assumption that service providers and consumers are known to each other. In common scenarios, before allowing access to (possibly) sensitive resources, trust relations are established by having clients create an account, store a profile, or requiring administrators to add identities to some kind of access control lists.

Alternatively, through semantic annotations, rule-oriented access control policies, and trust negotiation [34], entities can automatically build a trust relationship, through which they will feel comfortable sharing selected private information. The trust negotiation process evolves through disclosure of credentials and requests for credentials in an iterative and bilateral process. The requests for credentials are in fact authorization policies that have to be satisfied. The trust level reached through a negotiation between one client and a resource is in fact an authorization decision about whether to allow access to the resource. This approach distinguishes itself from identity-based access control because trust is established based on properties of previously unknown parties.

Trust negotiation is triggered when one party requests access to a resource owned by another party. The goal of a trust negotiation is to find a sequence of credentials C_1, \dots, C_k, R , where R is the resource where access is attempted, such that when credential C_i is disclosed, its access control policy has been satisfied by credentials disclosed earlier in the sequence, or to determine that no such credential disclosure sequence exists.

The PEERTRUST language [12, 23] for expressing access control policies is based on definite Horn clauses (the basis for logic programs [19]), namely, rules of the form

$$lit_0 \leftarrow lit_1, \dots, lit_n$$

In the remainder of this section, we concentrate on the syntactic features that are unique to the PEERTRUST language (check [12, 23] for more information) and we will consider only positive authorizations.

References to Other Peers The ability to reason about statements made by other parties is central to trust negotiation. For example, suppose that an online library requires a student id, issued by University ABC, before it allows access to the available online books. One can think of this as a

case of the library *delegating evaluation* of the query “Is the requester a student?” to the University “ABC” (ABC). To express delegation of evaluation to another party, we extend each literal lit_i with an additional *Issuer* argument,

$$lit_i @ \text{Issuer}$$

where *Issuer* specifies the party who is responsible for evaluating lit_i or has the authority to evaluate lit_i . For example, the library’s policy may be expressed as

$$\begin{aligned} &\text{Library:} \\ &\text{allowAccess(Book,X)} \leftarrow \\ &\quad \text{student(X) @ 'ABC'}. \end{aligned}$$

where X represents the requester whose quality of student should be evaluated by ‘ABC’. For clarity, we prefix each rule by the party in whose knowledge base it is included.

The *Issuer* argument can be a nested term containing a sequence of issuers, which are evaluated starting at the outermost layer. For example, it is unlikely that ABC will answer all the requests directly from the library, so a more practical approach is to ask the requester to evaluate the query himself, that is, to disclose his student id:

$$\begin{aligned} &\text{Library:} \\ &\text{allowAccess(Book,X) } \leftarrow \\ &\quad \text{student(X) @ 'ABC' @ X}. \end{aligned}$$

The library can refer to the party who asked a particular query by including a *Requester* argument in literals, so that we now have literals of the form

$$lit_i @ \text{Issuer } \$ \text{Requester}$$

Using the *Issuer* and *Requester* arguments, we can delegate evaluation of literals to other parties and also express interactions and the corresponding negotiation process between parties. Therefore, the final shape of the rule of the online library would look like the following:

$$\begin{aligned} &\text{Library:} \\ &\text{allowAccess(Book) } \$ \text{Req} \leftarrow \\ &\quad \text{student(Req) @ 'ABC' @ Req}. \end{aligned}$$

Local Rules and Signed Rules Each party defines the set of access control policies that apply to its resources, in the form of a set of definite Horn clause rules that may refer to the properties of those resources. These and any other rules that the party defines on its own are its *local* rules. A party may also have copies of rules defined by other parties (credentials), and it may use these rules in its proofs in certain situations. For example, Alice can use a rule that was defined by ABC to prove that she is really a student:

$$\begin{aligned} &\text{Alice:} \\ &\text{student('Alice')} @ \text{'ABC'} \\ &\quad \text{signedBy ['ABC']}. \end{aligned}$$

In this example, the “signedBy” term indicates that the rule has ABC’s digital signature on it. A signed rule has an additional argument that says who issued the rule.

Guards To guarantee a certain evaluation order for the literals in the body of a rule, we split the body’s literals into a sequence of sets, divided by the symbol “|”. All but the last set are *guards*, and all the literals in one set must evaluate to true before any literals in the next set are evaluated. For example, if the library offers a special price for students interested in a paperback copy of a book, the policy could be expressed as

```
Library:
order(Book) $ Req ←
  student(Req) @ Institution @ Requester |
  applyDiscount(Book).
```

The requester would have first to prove he is a student. Then the library would apply a discount for the requested book.

6 Distributed Policy-Based Negotiation for Grid Services Authorization

This section addresses some of the current Grid Security Infrastructure limitations by describing how PEERTRUST policies⁴ can be integrated to accommodate a large, loosely coupled Grid environment.

In our initial example (see Figure 1) we illustrated how different resources may be used in a Grid environment based on the Globus Toolkit 4.0. The fact that resources are often expensive in terms of cost, maintenance, and management normally results in more complex access policies than the already established conditions of two previously acquainted partners (in our example, ABC and Navy Institute). Jobs may use discovery and scheduling services to locate the best available resources, which make it even more difficult to predict what service instances will be used.

Our solution proposes, on one hand, that each Grid service advertises its authorization requirements through access control policies and, on the other hand, that each client specifies his credential disclosure policies and is able to query for resource’s access policies. When such policies are advertised, services and clients can negotiate, increasing incrementally their trust relationship based on satisfying the other party’s policies. We will change our initial scenario in order to accommodate the new capabilities highlighted above.

Alice’s research group needs the data as soon as possible, so Alice’s client program queries the National Monitoring

⁴Although we use the PEERTRUST language for policy specification in our examples and implementation, all the ideas presented in this paper are extensible to any other policy language with *delegation of authority and negotiation capabilities* [25].

and Discovery Hierarchical Service (NMDHS) for the best Linux cluster in terms of memory size and number of available processors. Alice program will send a request of the form `queryingAllowed()` to the NMDHS. The service policy requires any client to have a credential signed by one of the recognized state institutions (check the `id` policy below), so Alice’s client provides Alice X.509 End Entity Certificate issued by the ABC CA. Since the credential is not protected by any policy, Alice’s client program discloses it, authorization is fulfilled, and the requested information is delivered.

```
NMDHS:
queryingAllowed() $ Req ←
  validCredential(Req).
validCredential(Req) ←
  id(Req,'ABC CA') @ 'ABC CA' @ Req.
validCredential(Req) ←
  id(Req,'Navy Ins. CA') @ 'Navy Ins. CA' @ Req.
...
```

The best available Linux cluster belongs to the Research Center for Aeronautical Sciences (RCAS), so Alice’s client program initiates a new negotiation process for submitting Alice’s job. The RCAS Linux cluster policy restricts access to jobs acting on behalf of members of projects listed in the Ministry of Education (MinEdu) database.

```
RCAS Cluster:
submit(Job) $ Req ←
  actingOnBehalfOf(User,Job),
  member(User,Project) @ 'MinEdu CAS' @ Req.
```

Alice’s client program demonstrates that the job is acting on Alice’s behalf by providing Alice’s certificate, but it does not yet have a credential proving that Alice is a member of a listed project. To solve this situation, Alice’s client forwards a query to the Ministry of Education CAS server (MinEdu CAS) and retrieves an assertion attesting that Alice participates in a project regarding signaling instrumentation. This credential is disclosed to the RCAS Linux cluster, and job submission is granted. Since the job needs to contact other resources, a proxy credential gets delegated on behalf of Alice, to the RCAS Linux Cluster.

When Alice’s job starts its execution, it has to retrieve the input data from the ABC Reliable File Transfer Service (ABC RFT). ABC RFT policies require the job to prove that it is acting on behalf of an ABC staff member.

```
ABC RFT Service:
retrieve(File) $ Req ←
  member(Req,'Staff') @ 'ABC CAS' @ Req |
  check(Rights,Req,File) @ 'ABC CAS'.
store(File) $ Req ←
  member(Req,'Staff') @ 'ABC CAS' @ Req.
```

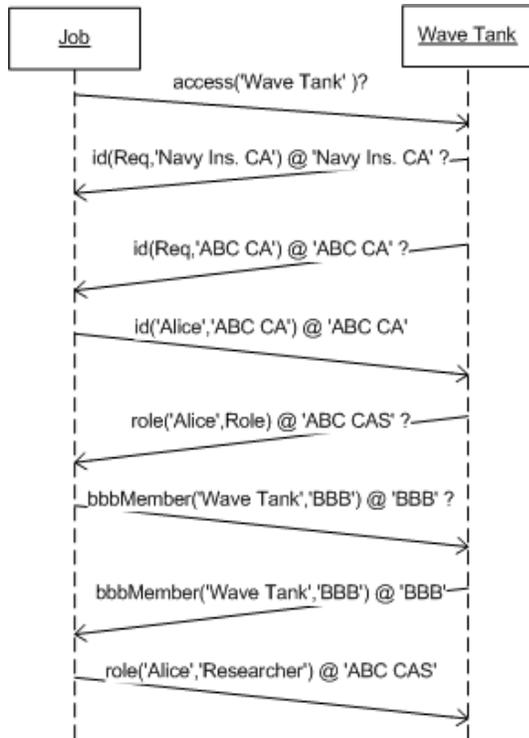


Figure 2. Job Negotiations

Therefore, using its delegated credential, the job authenticates to the ABC CAS, demonstrates it is acting on Alice's behalf, and retrieves a credential attesting that Alice is part of the university staff. Once this credential is released to the ABC RFT service, the service checks by itself, with the same ABC CAS if the required file can be accessed by Alice. ABC RFT acquires an assertion identifying the input file as belonging to Alice, and in this way the file retrieval operation is allowed.

Now the job contacts the Navy Institute wave tank (with a query `access('Wave Tank')`) and discovers that it has two options: either to provide an identity certificate signed by ABC CA or one issued by the Navy Institute CA (see Figure 2). By disclosing Alice's delegated certificate, the job demonstrates it is acting on behalf of an entity certified by ABC CA. Further, the wave tank informs Alice's job that it should reveal Alice's role at the university.

```
Wave Tank:
access(Resource) $ Req ←
  checkUser(Req,CA), |
  notInUse(Resource),
  timeLimitNotExceeded(CA).
checkUser(Req,'Navy Ins. CA') ←
  id(Req,'Navy Ins. CA') @ 'Navy Ins. CA' @ Req.
checkUser(Req,'ABC CA') ←
  id(Req,'ABC CA') @ 'ABC CA' @ Req, |
```

```
role(Req,Role) @ 'ABC CAS' @ Req |
Role = 'Researcher'.
bbbMember('Wave Tank','BBB') @ 'BBB'
signedBy ['BBB'].
```

In this case, Alice had provided a policy to the job, stating that Alice's roles are revealed only to entities that have proved their liability by disclosing a Better Business Bureau (BBB) membership credential.

```
Job:
role('Alice',Role) @ 'ABC CAS' $ Req ←
member(Req,'BBB') @ 'BBB'.
```

The wave tank has no policy protecting its BBB credential; therefore, after its disclosure, the job contacts ABC CAS to retrieve Alice's roles. Since Alice is a researcher and assuming that the other local policies are fulfilled, the job is granted access to the wave tank.

The refined and corrected data generated by the Linux cluster is saved by using the ABC RFT Service, which allows file storage if the action is requested by a university staff member (see the RFT policy above). The job has cached this credential (during a previous interaction) and does not need to ask for it again from the ABC CAS. By revealing this proof of Alice's university staff membership, the job is allowed to use the ABC RFT for storing its output files.

This scenario shows that the job has been submitted with only one credential and dynamically negotiated authorization at each resource accessed. This negotiation involved learning from resources themselves which credentials are needed (specified in policies) and where to retrieve them from. Resources were shared with no implied previous interactions and with no further involvement of administrators other than to set policies for access control and credential disclosure.

7 Architecture and Implementation

One of our main design goals was direct integration with Grid services paradigms. To achieve this goal, we devised an extension that is easily pluggable into any GT4.0-compliant Grid service or client. In this section, we present the technical details of our implementation (see also Figure 3). Because this implementation was developed for the GT4.0 Java Container, it is entirely written in Java and therefore requires Java-based Grid services and clients for a straightforward integration.

As we have seen in our discussions regarding the PEERTRUST capabilities, a service might allow different possibilities for granting access (in our example the wave tank was requiring for an identity certificate signed by ABC CA or one issued by Navy Institute CA), and the client may

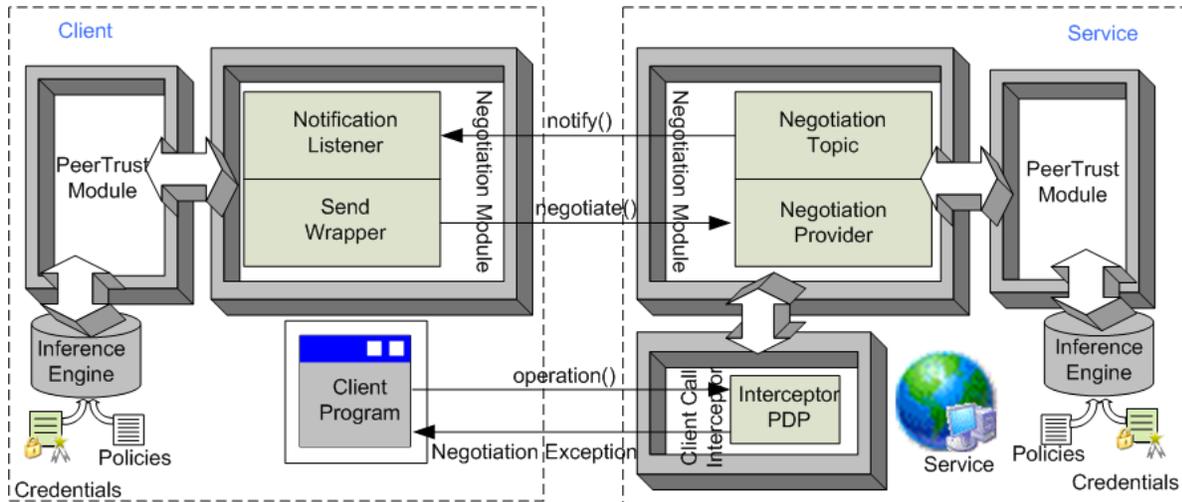


Figure 3. High-Level Architecture

choose which one to follow, according to his own credential protection policies.⁵ Because of the GT4.0 limitation in the PDPs decision composition (no disjunction supported), we have relied on the toolkit PDP support only for protection against unauthorized calls. We have used the PEERTRUST language and implemented its logic outside PDPs to permit policy combinations as both conjunctions and disjunctions. Our Interceptor PDP, part of the *Client Call Interceptor* (see Figure 3) intercepts client operation invocations and allows or denies the calls according to the existence of a previous successfully completed negotiation. If the answer is positive, the PDP grants the Grid service operation invocation. If the answer is negative, the PDP throws a negotiation exception to inform the client that a trust negotiation should be initiated and completed successfully before access is granted. In the latter case, the client Negotiation Module must start a negotiation with the homologous module on the Grid service (with a query of the form `request('OperationName')`). This triggers the negotiation process, in which policies and credentials are exchanged. The entire negotiation process is automated and requires no explicit user intervention.⁶

The *Negotiation Module* is responsible for the negotiation management. It handles the communication with other parties using *Grid connectors* (the Notification Listener and Send Wrapper interface on the client side and Negotiation Topic and Negotiation Provide on the service side - all implemented with Globus Toolkit 4.0) and interfaces with the

⁵For instance, the client might choose the one where less sensitive information is revealed or might choose a path with first required credentials available locally.

⁶Although we refer most of the time to fully automated negotiations, it would of course be possible to monitor and require confirmation from the user for all or some steps of the negotiation.

PEERTRUST *Module*. Whenever a negotiation succeeds, the Negotiation Module caches the identity of the client and the operation granted, in order to avoid the same negotiation in (presumably) successive calls. Ideally, the authorization decision should be invalidated once the shortest-lived credential (between those disclosed) has expired; but since we plan to experiment with a wide variety of credential formats, each requiring an independent parsing, the current implementation relies only on the expiration time of the client proxy certificate (the one used by the client for authentication).

On the service side, the Grid Connectors are composed of two modules: the *Negotiation Provider* and the *Negotiation Topic*. These modules are plugged into the Grid service, enhancing its functionality with policy-based negotiation capabilities. The client can use the additional Grid service's operations, implemented in the *Negotiation Provider Module*, for pushing policy requirements and credentials to the server side.

Negotiations are essentially asynchronous. A request might, for instance, trigger a new policy negotiation with a third party and so on, a situation that makes it difficult to predict when the request will finish its evaluation. Moreover, a client might choose to continue its usual execution without waiting for the answer from the server (e.g., in order to start in parallel several negotiations with different entities) and might simply prefer to be informed asynchronously when each negotiation is finished.⁷

Asynchronous client notifications can be implemented through GT 4.0 support for WS-Base Notification [1] and WS-Topics [31], which standardize asynchronous communications between consumers and providers. We associate

⁷Note that a server might act as a client if it starts a new negotiation with a third party.

each negotiation (triggered by the client request for a service operation) with a topic of interest (*Negotiation Topic*) through which messages can be delivered from the service to the client side. The client subscribes to this topic via a subscribe operation exposed by the notification producer (in our case the Grid service). On the client side, the notification consumer (in Figure 3 the *Notification Listener*) exposes a notify function that the Grid service uses to inform the client about its own requirements in terms of authorization.

A Grid service describes its supported operations and their associated parameters through a WSDL file. We have defined a special WSDL file that describes all the interfaces needed to perform our negotiations. Any Grid service can use our negotiation capabilities by including these definitions in its own WSDL description. Doing so, allows a Grid service to expose additional operations that are used to push messages to the Grid service during the trust negotiation process. These operations are implemented externally to the application-specific Grid service code in our Negotiation Provider Module. This module is plugged into a Grid service through the specification of a “provider” (our defined *TrustNegotiationProvider*) in the Grid service deployment descriptor (WSDD file). “Providers” are used by GT4.0 to confer extra functionalities to Grid services. GT4.0 support for notifications is also plugged into a Grid service by means of two providers (delivered with GT4.0) configured in the same WSDD file. We emphasize that in order to confer the negotiation capabilities to any Grid service we require small modifications to its descriptors (configuration) files, but no modification to the application code itself.

For the client side, we have developed an API with a jar file to facilitate easy integration of trust negotiation capabilities with the client code. Through this API, the client needs to set the Grid service address and must implement an interface (*Send Wrapper*) to communicate with the service with which trust is negotiated.

So far we have described how the policy-based negotiations can be integrated into GT4.0, but we have not described how the reasoning is performed. Queries for enforcement of policies are sent to the *PEERTRUST Module*, which controls the access to the *Inference Engine*. The *Inference Engine* is built on the Minerva Prolog engine⁸ and reasons on policies and credentials configured locally or retrieved during the negotiations undertaken. The Inference Engine answers to the *PEERTRUST* module queries indicating whether a request conforms to local policies (therefore allowing or disallowing the disclosure of credentials requested) or whether extra conditions must first be satisfied by the requester. The evaluation of a query allows one to include a proof in any answer returned to other parties. The

⁸Minerva Prolog (<http://www.ifcomputer.com/MINERVA/>) provides a java-based prolog engine that allows for an easy integration.

proof contains the policies and credentials (possibly third-party credentials) required to attest the other party’s request.

Because our implementation uses internally an abstract representation of credentials and policies, it can use different kinds of formats (certificates, signed assertions, signed RDF statements, etc). Currently we support X.509 v3 certificates, which can carry holder attributes in the extensions. In addition, we have integrated the use of proxy certificates that carry SAML assertions retrieved from a Community Authorization Service.

8 Conclusions, and Related and Further Work

Our main contribution to the Grid authorization solution space is the presented negotiation process in which entities, distributed over the Grid, engage each other for access granting. During this process, entities express their access requirements and enumerate the required credentials with their associated third-party issuers. Since credentials are retrieved dynamically during the negotiations, the access policies fulfillment is established at runtime. Furthermore, services and clients are true peers in the negotiation, since both are able to protect their local resources (e.g., credentials and provided services) and to formulate requests for the required credentials from the other party.

Although negotiations require in average more round-trips and are therefore more expensive than conventional approaches, it is a trade-off with the time and efforts needed for hard-coded configurations and for out-of-band policy information exchange. We have performed an extensive set of simulations in order to estimate how the overhead of our proposed negotiations can be reduced. We have simulated a Grid scenario in which 1100 clients and 100 grid services interact in order to gain and allow access⁹. We assumed a number of 16 messages exchanged on average for each interaction¹⁰ and a service revisitation probability uniformly distributed between 49% and 92%¹¹. Using an authorization decision caching mechanism combined with client credential caching on the service side, we have been able to reduce the number of messages exchanged to 7 (assuming a revisitation probability of 49%) and close to the ideal 2 message interaction (for a revisitation probability of 92%). Caching mechanisms reduced the message overhead with percentages between 50% and 75% and may be an acceptable mechanism to better benefit from the enhancements our authorization negotiations provide.

⁹The ratio 1 to 11 of services to clients is consistent with current Web statistics for server:client ratio [21, 17]. We used web statistics due to lack of data about client and service distribution in real Grids.

¹⁰16 messages would accommodate simple scenarios in which each actor requires at most 4 credentials or interacts with at most 4 credential issuers

¹¹probabilities showed by studies on web surfing behavior [28, 9]

One of the current trends in Grid deployment is to move from identity-based to attribute-based authorization or role-based access control (RBAC). The main reasons for this shift are the scalability issues associated with identity-based policies and the ease of administration through abstraction mechanisms as groups and roles. Projects such as VOMS [2], PERMIS [8], PRIMA [20], and XPOLA [11] have proposed solutions for attribute-based access control. GridShib [32], a recently started project with the goal of integrating Shibboleth [27] and the Globus Toolkit, aims to transparently retrieve attribute assertions from a Shibboleth service for the requester. The Community Authorization Service(CAS) set at the VO level is an authorization service where the client can ask for the access rights needed to invoke a request at a remote service. Despite the increased scalability, however, all these authorization solutions rely on fixed configurations of trust roots, on expected knowledge of required credentials, and on expected knowledge of authorities from which such credentials can be retrieved. Their deployment is therefore inflexible, and policy changes are expensive. In addition, most authorization schemes are server policy centered, and no policy is enforced on the service itself by the client.

The ATNAC framework [24] complements our work by using suspicion levels based on the history of past interactions and the federated knowledge of client behavior over the Grid in order to protect trust negotiations and access control against DoS attacks.

It is encouraging to see that some of the requirements for policy advertisement and capability matching that we have identified and addressed in our negotiation framework are recognized in emerging standards and specifications such as XACML [10], SAML [26], WS-Agreement [14], and WS-Policy [30]. We are following their development closely and hope to leverage in our own negotiation protocols work some of the policy primitives that may result from these efforts.

In the near future, we will investigate, further, the most simple scenario, where only a two-step negotiation is required to query resource access policies and subsequently fulfill them. Additionally, we also plan to refine our simulations by using a Zipf distribution for the selection of clients and servers. Zipf distributions are applicable in many fields of daily life and in our case it would show how some clients may interact more often than others, and how some servers may be preferred over others. In addition, declarative policies may create loops during their distributed evaluation; we plan to integrate existing techniques for loop detection [4] and the avoidance of deadlocks [18].

A further area of research is the possible utilization of our traceable negotiation process for accounting purposes, such as billing and audit. Our infrastructure already supports the recording of the negotiation steps at each resource

involved. The iterative process of requests for and disclosures of credentials may be extended to include the negotiation for pricing and exchange of payment until an agreement is reached.

Lastly, we plan the integration of a more expressive language than PEERTRUST (e.g., PROTUNE [7]) into our framework to allow, for example, the execution of arbitrary operations when certain policies are satisfied (e.g., logging information).

Acknowledgments

This work was partially funded by the European Commission and by the Swiss State Secretariat for Education and Research within the 6th Framework Programme project REVERSE number 506779 (cf. <http://reverse.net>), and by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Dept. of Energy, under Contract W-31-109-Eng-38.

References

- [1] Web Services Base Notification. <http://docs.oasis-open.org/wsn/2004/06/wsn-WS-BaseNotification-1.2-draft-03.pdf>.
- [2] R. Alfieri, R. Cecchini, V. Ciaschini, L. dell'Agnello, Á. Frohner, A. Gianoli, K. Lórentey, and F. Spataro. VOMS: An Authorization System for Virtual Organizations. In *Proceedings of the 1st European Across Grids Conference*, Santiago de Compostela, February 2003.
- [3] W.E. Allcock, I. Foster, and R. Madduri. Reliable Data Transport: A Critical Service for the Grid. In *Building Service Based Grids Workshop, Global Grid Forum 11*, 2004.
- [4] Miguel Alves, Carlos Viegas Damásio, Daniel Olmedilla, and Wolfgang Nejdl. A distributed tabling algorithm for rule based policy systems. In *7th IEEE POLICY 2006*, London, Ontario, Canada, June 2006.
- [5] Grigoris Antoniou, Matteo Baldoni, Piero A. Bonatti, Wolfgang Nejdl, and Daniel Olmedilla. Rule-based policy specification. In Ting Yu and Sushil Jajodia, editors, *Secure Data Management in Decentralized Systems*, volume 33 of *Advances in Information Security*. Springer, 2007.
- [6] Jim Basney, Wolfgang Nejdl, Daniel Olmedilla, Von Welch, and Marianne Winslett. Negotiating Trust on the Grid. In *2nd WWW Workshop on Semantics in P2P and Grid Computing*, New York, USA, may 2004.

- [7] Piero A. Bonatti and Daniel Olmedilla. Driving and Monitoring Provisional Trust Negotiation with Metapolicies. In *6th IEEE POLICY 2005*, pages 14–23, Stockholm, Sweden, jun 2005.
- [8] D.W. Chadwick and O.Otenko. The PERMIS X.509 Role Based Privilege Management Infrastructure. In *7th ACM Symposium on Access Control Models and Technologies*, 2002.
- [9] A. Cockburn and B. McKenzie. What do web users do? an empirical analysis of web use. *International Journal of Human-Computer Studies*, 54:903–922, 2002.
- [10] OASIS eXtensible Access Control Markup Language (XACML) TC. http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml.
- [11] L. Fang, D. Gannon, and F. Siebenlist. XPOLA - An Extensible Capability-based Authorization Infrastructure for Grids. In *4th Annual PKI R&D Workshop*, 2005.
- [12] Rita Gavrioloie, Wolfgang Nejdl, Daniel Olmedilla, Kent E. Seamons, and Marianne Winslett. No Registration Needed: How to Use Declarative Policies and Negotiation to Access Sensitive Resources on the Semantic Web. In *1st European Semantic Web Symposium (ESWS 2004)*, volume 3053 of *LNCS*, pages 342–356, Heraklion, Crete, Greece, may 2004. Springer.
- [13] Globus Toolkit. <http://www.globus.org>.
- [14] Grid Resource Allocation Agreement Protocol WG. <https://forge.gridforum.org/projects/graap-wg>.
- [15] Grid Security Infrastructure. <http://www.globus.org/security/overview.html>.
- [16] R. Housley, T. Polk, W. Ford, and D. Solo. *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. RFC 3280, 2002.
- [17] Internet world stats homepage. <http://www.internetworldstats.com/stats.htm>, 2006.
- [18] Ninghui Li, Wenliang Du, and Dan Boneh. Oblivious signature-based envelope. In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 182–189, New York, NY, USA, 2003. ACM Press.
- [19] John W. Lloyd. *Foundations of Logic Programming*. Springer, 2nd edition edition, 1987.
- [20] M. Lorch, D. Adams, D. Kafura, M. Koneni, A. Rathi, and S. Shah. The PRIMA System for Privilege Management, Authorization and Enforcement in Grid Environments. In *Proceedings of the 4th Int. Workshop on Grid Computing*, Phoenix, AZ, USA, November 2003.
- [21] Netcraft october 2006 web server survey <http://news.netcraft.com/>, 2006.
- [22] L. Pearlman, C. Kesselman, V. Welch, I. Foster, and S. Tuecke. The Community Authorization Service: Status and Future. In *Proceedings of the Conference for Computing in High Energy and Nuclear Physics*, La Jolla, California, USA, March 2003.
- [23] PeerTrust Project. <http://www.L3S.de/peertrust/>.
- [24] T. Ryutov, L. Zhou, C. Neuman, N. Foukia, T. Leithhead, and K. E. Seamons. Adaptive trust negotiation and access control for grids. In *6th IEEE/ACM International Workshop on Grid Computing*, 2005.
- [25] K. Seamons, M. Winslett, T. Yu, B. Smith, E. Child, J. Jacobson, H. Mills, and L. Yu. Requirements for policy languages for trust negotiation. In *IEEE POLICY'02*, page 68, 2002.
- [26] OASIS Security Services (SAML) TC. http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security.
- [27] Shibboleth Project, Internet2. <http://shibboleth.internet2.edu>.
- [28] L. Tauscher and S. Greenberg. How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems. *International Journal of Human Computer Studies*, 47(1):97–138, 2001.
- [29] The Open Grid Services Architecture OGSA, version 1.0.
- [30] Web Services Policy Framework. <http://www-128.ibm.com/developerworks/webservices/library/specification/ws-polfram/>.
- [31] Web Services Topics. <http://docs.oasis-open.org/wsn/2004/06/wsn-WS-Topics-1.2-draft-01.pdf>.
- [32] V. Welch, T. Barton, K. Keahey, and F. Siebenlist. Attributes, Anonymity, and Access: Shibboleth and Globus Integration to Facilitate Grid Collaboration. In *4th Annual PKI R&D Workshop*, 2005.
- [33] V. Welch, I. Foster, C. Kesselman, O. Mulmo, Laura Pearlman, S. Tuecke, J. Gawor, S. Meder, and F. Siebenlist. X.509 Proxy Certificates for Dynamic Delegation. In *3rd Annual PKI R&D Workshop*, April 2004.
- [34] William H. Winsborough, Kent E. Seamons, and Vicki E. Jones. Automated trust negotiation. DARPA Information Survivability Conference and Exposition, IEEE Press, Jan 2000.