Definition and Implementation of a SAML-XACML Profile for Authorization Interoperability across Grid Middleware in OSG and EGEE

Gabriele Garzoglio[1] (garzoglio@fnal.gov), Ian Alderman[9] (alderman@cs.wisc.edu), Mine Altunay[1] (maltunay@fnal.gov), Rachana Ananthakrishnan[8] (ranantha@mcs.anl.gov), Joe Bester[8] (bester@mcs.anl.gov), Keith Chadwick[1] (chadwick@fnal.gov), Vincenzo Ciaschini[7] (vincenzo.ciaschini@cnaf.infn.it), Yuri Demchenko[4] (demch@science.uva.nl), Andrea Ferraro[7] (andrea.ferraro@cnaf.infn.it), Alberto Forti[7] (alberto.forti@cnaf.infn.it), David Groep[2] (davidg@nikhef.nl), Ted Hesselroth[1] (tdh@fnal.gov), John Hover[3] (jhover@bnl.gov), Oscar Koeroo[2] (okoeroo@nikhef.nl), Chad La Joie[5] (chad.lajoie@switch.ch), Tanya Levshina[1] (tlevshin@fnal.gov), Zach Miller[9] (zmiller@cs.wisc.edu), Jay Packard[3] (jpackard@bnl.gov), Håkon Sagehaug[6] (hakon.sagehaug@bccs.uib.no), Valery Sergeev[1] (vsergeev@fnal.gov), Igor Sfiligoi[1] (sfiligoi@fnal.gov), Neha Sharma[1] (neha@fnal.gov), Frank Siebenlist[8] (franks@mcs.anl.gov), Valerio Venturi[7] (valerio.venturi@cnaf.infn.it), John Weigand[1] (weigand@fnal.gov)

[1] *Fermilab, Batavia, IL, USA*
[2] *NIKHEF, Amsterdam, The Netherlands*
[3] *Brookhaven National Laboratory, Upton, NY, USA*
[4] *University of Amsterdam, Amsterdam, The Netherlands*
[5] *SWITCH, Zürich, Switzerland*
[6] *BCCS, Bergen, Norway*
[7] *INFN CNAF, Bologna, Italy*
[8] *Argonne National Laboratory, Argonne, IL, USA*
[9] *University of Wisconsin, Madison, WI, USA*

Keywords: authorization, interoperability, SAML-XACML, OSG, EGEE

Abstract: In order to ensure interoperability between middleware and authorization infrastructures used in the Open Science Grid (OSG) and the Enabling Grids for E-sciencE (EGEE) projects, an Authorization Interoperability activity was initiated in 2006. The interoperability goal was met in two phases: first, agreeing on a common authorization query interface and protocol with an associated profile that ensures standardized use of attributes and obligations; and second, implementing, testing, and deploying, on OSG and EGEE, middleware that supports the interoperability protocol and profile. The activity has involved people from OSG, EGEE, the Globus Toolkit project, and the Condor project. This paper presents a summary of the agreed-upon protocol, profile and the software components involved.

1. Introduction

The Open Science Grid (OSG) [1] and the Enabling Grids for E-sciencE (EGEE) [2] are two major projects devoted to promoting science through the use of distributed, Grid computing. Despite the fact that the two projects are mostly independent and operate hardware resources in different parts of the world, a substantial part of the software stack is shared between the two.[1]

Both OSG and EGEE base their authentication infrastructure on Public Key Infrastructure (PKI), leveraging X.509 end-entity and proxy certificates [6,7] for single sign-on and delegation. Initially, both Grids based their authorization infrastructures on policies local to resources. With time, however, they extended their infrastructures to centralize the authorization policies at the level of individual sites. In addition, both Grids extended their infrastructures to include role-based access to resources, based on a user's virtual organization (VO) membership.

While a similar security model based on a user's VO membership was successfully maintained between the two Grids, the mechanisms to centralize authorization policies risked divergence. Drawbacks of such divergence include duplication of work and the requirement that middleware common to both Grids support different authorization plug-ins, depending on the Grid on which it was deployed.

The Authorization Interoperability activity was formed in 2006 to address this issue. The

[1] Examples of shared software products are the disk cache Storage Resource Manager (SRM) [3] and the gLExec identity switching tool [4,5].

collaboration defined a common protocol for use by OSG/EGEE and an identity attribute profile for authorization call-out to site-central policy decision services. In lock step, independent libraries in both C and Java have been implemented according to the agreed protocol and profile and are being used for cross-implementation interoperability.

The activity had resonance with major middleware providers for both Grids, namely, the Globus Toolkit and the Condor [14] groups. Being active participants in the activity, these groups have started providing middleware that natively supports our common authorization protocol. This support greatly simplifies the process of deploying such middleware on both OSG and EGEE.

This paper is organized as follows. Section 2 presents work related to the Authorization Interoperability activity. Section 3 describes the OSG and EGEE security models. Section 4 summarizes the principal elements of the common Authorization Interoperability profile. Section 5 discusses how the infrastructures implemented the common profile. Section 6 discusses future work, and Section 7 presents a summary of the paper.

2. Related Work

The Authorization Interoperability activity has produced a call-out protocol and attribute profile from resource gateways to policy decision services. The activity limited its scope to the EGEE and OSG security model, whereby identities are described via X509 certificates and identity attributes via VOMS [9] attribute certificates. It also targeted a limited set of authorization systems for implementation, namely, the Grid User Mapping Service (GUMS) [13] for OSG and the Site Central Authorization Service (SCAS) [27] for EGEE.

The Open Grid Services Architecture (OGSA) Authorization Working Group (WG) of the Open Grid Forum (OGF) [24] is addressing the same problem in a more general context. The objective of the OGSA Authorization WG is to define the specifications needed to allow for pluggable and interoperable authorization components from multiple authorization domains in the OGSA framework. There are a number of authorization standards and working systems in the Grid today (Akenti, Cardea, CAS, PERMIS, etc.), in addition to VOMS and XACML; the OGSA-Authorization specifications aim at allowing these solutions to be used interchangeably with middleware that requires authorization functionality. The OGSA-Authorization group leverages authorization work that is ongoing in the Web services world (e.g., SAML, XACML, the WS Security specification suite) and defines specifications for how these should be used for Grid services.

When the Authorization Interoperability activity started, the specification of the OGSA-WG did not address all use cases of interest to our collaboration. Instead of working through OGF's WG process, we tried to focus on our immediate requirements and associated solution as a more efficient path to pragmatic results. However, having three members of the OGSA-Authorization WG participating in our activity, we are regularly feeding our experiences and detailed requirements back into the WG standardization effort.

3. The OSG and EGEE Security Model

OSG and EGEE have a common security model, based on PKI, using X.509 end-entity and proxy certificates. Certificates are used to mutually authenticate every request for service. Integrity and confidentiality of the communication are supported at both the transport and message layer, using the standard Transport Layer Security (TLS) protocol [25].

Resources are made available on the Grid for user communities, also called Virtual Organizations (VOs). Access to resources is granted to users on the basis of their membership in a VO, rather than on the user's personal attributes.

In this common security model, VOs organize their internal membership structure according to hierarchical groups (e.g., /myVO / Production). Members of a group can have special roles for that group (e.g. /myVo / Production / Role = SimulationManager). This structure and the relative membership of each user are managed and published via Virtual Organization Management Servers (VOMS).

Conversely, in the same model, resources are grouped according to the administrative boundaries of the sites that make compute facilities available. Access to different resources (storage, computing, worker nodes, etc.) is managed by middleware that acts as a gateway to those resources. To implement access policy enforcement, gateways of both Grids obtain the user's membership information and the VO organizational structure from each VO's VOMS. These common sources of attributes, together with a well-defined access protocol and asserted unique identities, lay the foundation for interoperations across collaborating Grids.

Each Grid organization makes available to its member sites lists of member VOs and VO preferred resource access policies. Note, however, that the individual sites have the ultimate authority on what VOs, VO groups, and VO members are supported and what privileges are granted to them with respect to the site's resources. Typically, privileges are determined by membership in VO groups and roles, like relative priority in a batch system or read/write access to storage areas. Those attributes that univocally identify users, like the user's X509 Distinguished Name, are used for some VOs to enable operating system-level

protection of concurrently running processes from different users on the same machine.

Users interact with resource gateways on behalf of VOs and VO groups with a certain group role. Before every interaction, users are responsible for including this information with their credentials. The information is expressed in terms of VO membership attributes, or Fully Qualified Attribute Names (FQANs), and is encapsulated in an Attribute Certificate (AC). This AC is requested from the VO's VOMS, which acts as an attribute authority and digitally signs it for future validations.

In this model, users always push this AC with all embedded attributes necessary for authorization to resources. In other words, resources never directly pull attributes from VO or institutional repositories on behalf of the user. When AC-enhanced credentials are pushed to a resource, the resource gateway typically validates the AC and its issuer, extracts all user attributes, and conveys them to a repository of authorization policies, or Policy Decision Point (PDP), central to the site. In turn, the PDP replies with an authorization decision and a set of privilege constraints, also called Obligations. The gateway acts as a Policy Enforcement Point (PEP) and enforces the PDP decision. Figure 1 shows a diagram of the security model.

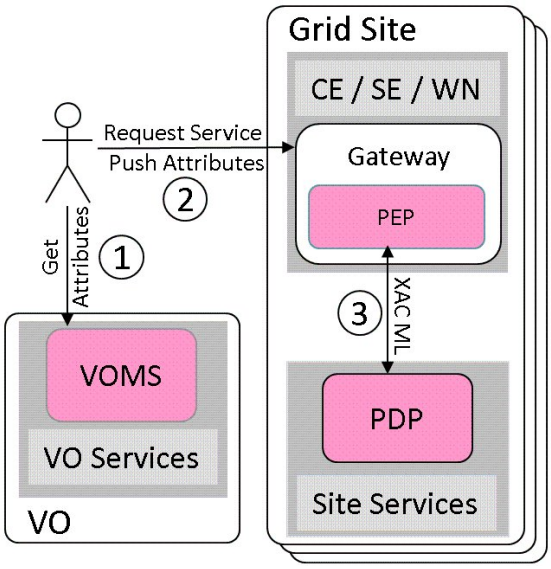

Fig. 1: Diagram of the OSG and EGEE Security Model. User attributes are obtained from VOMS (1). Service requests are issued while pushing user credentials and attributes to a resource gateway at a grid site (2). The resource gateway acts as a Policy Enforcement Point (PEP) and contacts the site-central Policy Decision Points (PDP) using the common XACML Authorization Interoperability protocol (3).

The Authorization Interoperability activity focused on standardizing a protocol for PEP-to-PDP communication. Despite the commonality of the security model, this activity was a fundamental step to allow the deployment of resource gateway implementations on OSG or EGEE, without the need for Grid-specific authorization plug-ins. The common protocol allows Grid developer groups associated with EGEE or OSG to reuse a common implementation of the security call-out libraries, thus reducing maintenance and eliminating duplication of work.

4. The Authorization Interoperability Protocol

In the EGEE and OSG security model, authorization is based on user's X509 identity attributes and VO membership attributes. These attributes are all pushed by the user to the resource gateway through the use of X509 end-entity and attribute certificates. The resource gateway queries a centralized authorization service using the Authorization Interoperability protocol, which is based on the SAML v2.0 profile of XACML v2.0 [16,17]. This query protocol allows all the user's attributes to be encapsulated in the request query. The names and types of those request-attributes are defined in a common profile [18]. The profile also provides an abstraction for what resource types and what actions are considered within the authorization model.

4.1. SAML Profile of XACML

The Extensible Access Control Markup Language (XACML) is a standard defined by the Organization for the Advancement of Structured Information Standards (OASIS). It is a core XML schema for representing authorization and entitlement policies. The Security Assertion Markup Language (SAML), developed also by OASIS, is an XML-based framework for communicating user authentication, entitlement, and attribute information. As of version 2, the two standards can be used together to allow for greater power and flexibility.

In the context of this paper, the SAML profile of XACML is used to convey detailed request-reply message information about the subject, resource, and action from the service gateway (PEP) to the centralized authorization service (PDP) and to convey back the authorization decision and optional local-account mapping information.

In XACML terminology, the message sent from PEP to PDP is referred to as a "Request Context," while the message going from PDP to PEP is referred to as a "Response."

Note that for our work we have standardized only on the authorization query interface as specified by the SAML-XACML profile. We do not mandate use of XACML-compliant policy evaluators at the PDPs.

4.2. Attribute Namespace

A SAML profile can accommodate several profile extensions. To avoid conflicts with other extensions, the Authorization Interoperability profile uses its own prefix "http://authz-interop.org/xacml" in URL format and the associated "x-urn:authz-interop:xacml" in URN format.

Attributes can use either the URL or the URN formats. While we prefer the URL style, we acknowledge that both styles present advantages and disadvantages.

The URL style namespace offers the advantage that, once a domain namespace (DNS) for a host machine is obtained, there is no need for further registrations of the attribute namespace with any standardization body. The uniqueness of the namespace is derived by the uniqueness of the domain name. Moreover, additional services for XML schema resolution and location can be established at the registered domain. For example, both OGF and W3C support direct mapping and resolution of registered XML infoset schemas into URLs.

The URN style namespace offers the advantage of compatibility with standards bodies like OASIS and the Internet Engineering Task Force (IETF). However, using a URN requires the formal registration of the namespace with bodies like the Internet Assigned Number Authority (IANA). To minimize this problem, the Authorization Interoperability profile defines a URN starting with the "x-" prefix [19], for "experimental namespace," that doesn't require registration with IANA.

4.3. XACML Request

In the XACML model, the PEP sends an XACML request with an authorization query to the PDP. The query consists on whether a user or a service (the "subject") is allowed to execute an "action" on a "resource" controlled by the PEP, within the context of certain request conditions or "environment. The request, therefore, contains four attribute sections, or contexts ("subject", "action", "resource", "environment"), that define its scope. We discuss below these four contexts in more detail.

- "**Subject**" - A PEP uses the subject context to declare the entity for which the authorization decision is requested. The subject attributes are used to determine an authorization decision, but the PDP does not need to consider all attributes in the subject section to determine a decision.
- "**Resource**" - The attributes in the Resource context describe the resource targeted for the authorization request. The resource is typically under the control of the PEP, which acts as a gateway to the resource.
- "**Action**" - The attributes in the Action context describe the action that the subject wants to perform on the specified resource.
- "**Environment**" - The attributes in the Environment context convey additional parameters in the authorization request of the subject to perform an action on the specified resource. Sometimes, these attributes specify conditions such as the time of the request; but profiles, like the Authorization Interoperability profile, use it for more complex use cases, as discussed below.

4.4. Authorization Interoperability Request Profile

The Authorization Interoperability group has agreed on a profile for required and optional XACML request attributes, on each of the four XACML request contexts. These attributes encapsulate the access authorization use cases common to the OSG and EGEE models. The following is a short summary of the profile attributes, organized by context. The reader is encouraged to consult the document "An XACML Attribute and Obligation Profile for Authorization Interoperability in Grids"[18] for a full description of the attributes and examples of XACML messages and policies.

Subject: Both OSG and EGEE authorization infrastructures define user and service identities through X.509 end-entity certificates, while the VOMS-issued user's attributes are asserted through X509 attribute certificates. In our profile, the attribute namespace for the subject context is http://authz-interop.org/xacml/subject/. Within this namespace, the following XACML attributes are related to basic X.509 properties:

- subject-x509-id: the Distinguished Name (DN) of the user or service for whom the access authorization is requested.
- subject-x509-issuer: the DN of the entity that signed the user or service end-entity certificate, typically a CA.

The following XACML attributes are related to VOMS attributes:

- vo: the name of the Virtual Organization for which the user is requesting the access authorization.
- voms-signing-subject: the DN of the VOMS service certificate that signed the VOMS attribute certificate.
- voms-signing-issuer: the DN of the entity that signed the VOMS service certificate; this is typically a CA.
- voms-fqan: the list of fully qualified attribute names for the subject; the FQANs express the membership of the subject to VO groups and group roles.

- voms-primary-fqan: the first element of the FQAN list; this FQAN carries particular significance in the OSG and EGEE model: users define on behalf of what VO group or group role they use a resource via this FQAN.

An XACML attribute in the profile is used to define a Condor canonical name:

- subject-condor-canonical-name-id: in the condor system, privileges are associated to users, identified by canonical names. This attributes carries the user canonical name.

In addition, the profile defines a series of optional attributes, which are not discussed in this paper.

Resource: In our profile, we define only resources of particular interest to our community. The resource targeted by the request is expressed by using the OASIS attribute name "resource-id." The following values, prefixed with http://authz-interop.org/xacml/resource-type/, are possible:

- CE: a computing element is a gateway to a cluster of computing resources; typically, a CE controls access to a computing cluster.
- WN: a worker node is a machine that is part of a computing cluster. This resource is generally controlled by a local batch system and may not be directly accessible by services on the Grid or the Grid authorization infrastructure. Both EGEE and OSG, however, adopt pilot-based workload management systems, like GlideinWMS [20], Panda [21], and DIRAC [22], that allow the registration of a worker node to a VO-specific pool of grid resources; this registration is achieved by submitting to the CE a "pilot" job, which is then responsible for the execution of user jobs. In these cases, access to a WN can be centrally controlled by the site authorization system.
- SE: a storage element controls access to files and storage pools.

Other attributes that characterize Grid resources are defined within the namespace http://authz-interop.org/xacml/resource/. These attributes carry information such as the Domain Name of the resource or the DN and issuer of the X.509 host certificates that defines the resource identity.

Action: In our profile, the action is expressed by using the OASIS attribute name "action-id." We defined an enumeration of possible actions for "action-id," each used in specific OSG and EGEE use cases. The following actions, prefixed with http://authz-interop.org/xacml/action-type/, are possible:

- queue: the subject requests authorization to interact with the job queue of the specified computing resource. This action is used in conjunction with the CE resource type, typically when requesting authorization to submit a job to the batch system queue controlled by a CE.
- execute-now: the subject requests authorization to execute immediately a job at the specified computing resource This action is used in conjunction with the CE (computing element) or WN (worker node) resource types, to execute a job at the computing element resource gateway machine or at a worker node.
- access: the subject requests authorization to access a specified storage resource. The scope of the request is implementation-dependent: the request can specify access to a single file, a list of files, or a remote/local storage pool. By design, this action generalizes finer-grained types of access, like read access, write access, and file system administrative access. Such fine-grained access control is delegated to the authorization layer of storage services.

Since both EGEE and OSG use Globus to control access to computing resources, the profile defines one attribute, prefixed with http://authz-interop.org/xacml/action/, to convey the detail of the Globus request:

- rsl-string: the Globus Resource Specification Language string.

Environment: As mentioned in the Resource section, both OSG and EGEE support direct management of jobs to worker nodes via pilot-based workload management systems. Our profile uses the environment context to convey to the PDP the identity of the pilot job. These attributes use the namespace http://authz-interop.org/xacml/environment/pilot-job/ and have the same attribute name as the attributes of the subject context.

4.5. XACML Response

In the XACML model, after processing the PEP's XACML request for access authorization, the PDP returns an XACML response to the PEP.

The principal element of an XACML response is an Authorization Decision Statement. This contains the actual decision: "Permit," "Deny," "Indeterminate," or "NotApplicable." In theory, the PEP can query a set of PDPs and combine the results according to different policies. The combined result must evaluate to "Permit" before the PEP will allow access.

As a part of the Authorization Decision Statement, the PDP can return a "Permit" with stated conditions under which the access should be granted, known as "Obligations." These Obligations typically identify privilege restrictions for the resource access.

4.6. Authorization Interoperability Response Profile

The Authorization Interoperability profile defines a set of Obligations to restrict the privileges granted for common OSG/EGEE use cases

accessing a computing or storage resource. For resources such as compute nodes, these access rights are defined by and enforced with the privileges of a specific local POSIX account. Additionally, for storage, access privileges can be limited, for example, to a specific subset of the file system. The reader is encouraged to consult the profile document [18] for a full description of the attributes and examples of XACML messages and policies.

The profile uses the namespace http://authz-interop.org/xacml/obligation/ for its obligations and the namespace http://authz-interop.org/xacml/attribute/ for the attributes related to these obligations. The following obligations are defined in the profile:

- uidgid: this requires the PEP to grant access to the resource with the privileges of the specified local Unix ID and Group ID.
- secondary-gids: this requires that the PEP grants the privileges associated to the specified secondary Group IDs.
- username: this requires the PEP to grant access to the resource with the privileges of the specified local Username.
- afs-token: this obligation conveys an AFS token, which the PEP must put in the environment of processes accessing the resource.
- root-and-home-paths: when accessing a storage element, this restricts access to a specified portion of the file system
- storage-access-priority: when accessing a storage resource, this conveys the priority of the request relatively to other requests.
- access-permissions: when accessing a storage resource, this informs the underlying storage system to restrict access to read-only or read-write mode.

5. Implementation

The proposed security model has been implemented by both OSG and EGEE and integrated in their different Grid infrastructures. The following sections describe details of the infrastructure implementations.

5.1. OSG Implementation of the Security Model

The OSG implementation of the security model is based mainly on the infrastructure provided by the VO Services project [12]. For those resource gateways that do not natively implement the Authorization Interoperability profile, the project offers a PEP call-out module for computing gateways, called PRIMA [13], and a call-out module and server for storage gateways, called gPlazma [3]. In addition, the project provides a PDP implementation, called GUMS [13].

PRIMA is a plug-in based on the Globus Security Infrastructure (GSI) [8]. It extracts the user's X509 DN and the first FQAN in the list of VO membership attributes, if present, and sends them over the network to GUMS. GUMS returns a mapping to a local POSIX account if the user is authorized. By enforcing the POSIX account policy, the host itself implements privilege restrictions at the resource.

GUMS is implemented as a Web service front-ending an authorization database. The base GUMS configuration typically consists of a list of VOMS servers and associated account mapping rules. On regular intervals, GUMS retrieves the list of user DNs and associated FQANs from all the VOMS servers listed in its configuration and synchronizes its database accordingly. A request from PRIMA triggers a database search, and the associated host-account mapping is returned if the user's DN and optional FQAN are found in the database. In compliance with the Authorization Interoperability profile, GUMS exposes a XACML authorization query interface, and the mapping information is returned as a "username" obligation.

PRIMA and GUMS communicate over a GSI connection with mutual authentication based on X.509 host certificates. Before the Authorization Interoperability activity, the communication protocol was a modified version of the SAML 1.0 Authorization Decision Query and Statement standard [13]. Currently, both PRIMA and GUMS support the new Authorization Interoperability protocol.

A similar mechanism is used in gPlazma for storage authorization. gPlazma uses GSI to extract the user's DN and first FQAN, if present, from the X509 proxy certificate and sends them to GUMS, using the same protocol as PRIMA. After receiving a reply from GUMS, gPlazma augments it with storage-specific attributes and forwards it to the storage system, which enforces the policy decision.

The latest version of the Globus Web services GRAM (WS-GRAM) computing gateway, which is part of the Globus Toolkit 4.2 [26, 28], has been enhanced to support the authorization call-out to GUMS using the Authorization Interoperability protocol and profile. Because of the common protocol, WS-GRAM can also interface to the SCAS PDP.

5.2 EGEE Implementation of the Security Model

The traditional EGEE implementation of the security model extends the GSI security libraries with the LCAS/LCMAPS framework [10,11]. Authentication and authorization are based mostly on FQANs. LCMAPS uses a enhanced grid mapfile format and a local enforcement, which maps the first FQAN in the list to a pool of POSIX accounts. The user's DN is not listed in the mapping file, but different DNs are still guaranteed to be mapped to

different accounts via an internal tracking mechanism.

Recently, EGEE has recognized the need for a centralized authorization service and has started the implementation of a PDP, called the Site Central Authorization Service (SCAS). A SCAS PEP is also being implemented as an LCMAPS plug-in. This plug-in is used by common middleware, such as the pre-Web Services Globus Gatekeeper, GridFTP, and the gLExec identity switching tool [4,5]. The SCAS PEP and PDP communicate via the Authorization Interoperability protocol. The user's Grid credential to local account mapping information is returned via the "uidgid" and "secondary-gids" obligations. Because of the common authorization protocol, gLExec can easily be deployed in both EGEE and OSG with minimal configuration changes.

5.3. XACML Libraries

The Authorization Interoperability activity has developed a set of libraries that implement the Authorization Interoperability protocol. These libraries are used in the implementations of PEPs and PDPs in both OSG and EGEE.

The authorization query is expressed as SAML/XACML messages, which are sent on the wire as SOAP messages over a TLS transportation layer. This protocol is typically implemented as a Web service interface.

For the SAML/XACML message processing, the Java applications leverage the OpenSAML v2.0 libraries [23], which were developed in collaboration with the Internet2 project. The C applications use equivalent processing libraries that were developed for this activity and are now part of the Globus Toolkit. In summary, the results of our activity are now generally available through the Internet2 and Globus open source projects.

5.4. Infrastructure Tests and Deployments

With the completion of the implementation of the XACML libraries and their integration with the principal resource gateways in OSG and EGEE, the infrastructure has undergone a series of interoperability tests. The targeted resource gateways were the pre-Web services Globus Gatekeeper, the Web service Globus Gatekeeper v4.2, GridFTP, the SRM/dCache Storage Service, and the gLExec identity-switching tool. Each of these gateways has been tested with both the GUMS- and SCAS-PDPs. Note that minimal changes to the gateway configuration were needed to switch between PDPs.

In addition to internal tests, the infrastructure is undergoing certification tests in both Grids for production deployment, which is scheduled at dozens of resources for early 2009.

5.5. Limitations and Future Operations

The main limitation of the current infrastructure is that authorization call-out modules implement only a common (thus interoperable) subset of the specifications. In particular, the implementations neglect those attributes used to express policies of future interest for OSG and EGEE. These include stricter authorization policies on pilot-based workload management systems, on job execution actions (e.g. "execute-now" vs. "queue"), and on storage access priorities and restrictions. We envision providing support for such policies as the need within either Grid arises.

The operations of the infrastructure involve more considerations than interoperability. We plan to discuss these in detail in a future paper, after enough experience has been gathered with the new interoperable infrastructure. Crucial operational processes already in place for the OSG and EGEE infrastructures define VO membership management (a common source of VO information between OSG and EGEE allows member access to both Grids), management of the desired mapping policy templates from the VO, and management of the implemented mapping policies by the Sites.

6. Future Work

The Authorization Interoperability collaboration envisions work in three main areas:

1) Extending the support of the protocol to additional resource gateways and policy decision points. These include the Site Authorization Service (SAZ) PDP, Globus Reliable File Transfer and Delegation services (PEPs), and the Berkeley Storage Manager (BeStMan) Storage Service (PEP). Working with some of these groups, we have observed that providing reference C or Java implementations of the call-out module was sufficient for them to develop working prototypes.
2) Extending the protocol to include additional use cases. These may include additional obligations, specifically for the storage use case. In general, maintaining authorization systems interoperable between OSG and EGEE is a goal of both Grids.
3) Working in the context of the OGSA-Authorization OGF Working Group, to ensure that all of our current use cases are included in the more general interoperability activity. The expectation is that the future OGF standard will eventually replace our Authorization Interoperability profile.

7. Summary

The goal of the Authorization Interoperability activity is to ensure interoperability between the middleware and authorization infrastructures used in the OSG and EGEE projects. Both Grids have a common security model, whereby users push to resources identity and attribute assertions, based on X509 certificates and VOMS attribute certificates. Both Grids are also moving toward a distributed authorization infrastructure, with site-central PDPs. In this context, authorization interoperability was achieved by defining a common authorization decision query protocol with an associated profile.

The Authorization Interoperability protocol is based on the SAML v2.0 profile of XACML v2.0. A set of attributes and obligations specific to the needs of OSG, EGEE, Globus Execution Service, and Condor are defined in a separate profile. The protocol and profile have been implemented as part of the authorization tools of Globus Toolkit, OSG, and EGEE, while the Condor team is planning to follow suit. Interoperability test suites have helped us to ensure common adherence to the commonly agreed standards across implementations.

The definition of a common protocol is a significant step forward for OSG and EGEE, as it enables better interoperability of services as well as providing software reuse opportunities across our projects.

Acknowledgments: Fermilab is operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy. This work was partially funded by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Dept. of Energy, under Contract DE-AC02-06CH11357.

References

[1] R. Pordes et al., "The Open Science Grid," Journal of Physics: Conference Series 78 , Institute of Physics Publishing, 2007 (15 pp)

[2] "EGEE Home," http://www.eu-egee.org/, Accessed October 2008.

[3] A. S. Rana et al., "Introducing Advanced Fine-grained Security in dCache-SRM for PetaByte-scale Storage Systems on Global Data Grids: gPLAZMA `grid-aware PLuggable AuthoriZation MAnagement System'," Nuclear Science Symposium Conference Record, 2006. IEEE, pp. 632-636, ISBN: 1-4244-0561-0.

[4] I. Sfiligoi et al., "Addressing the pilot security problem with gLExec," Journal of Physics: Conference Series 119, Institute of Physics Publishing, 2008 (6 pp)

[5] D. Groep et al., "gLExec: Gluing Grid Computing to the Unix World," Journal of Physics: Conference Series 119, Institute of Physics Publishing, 2008 (11 pp)

[6] ITU-T Recommendation X.509 (1997 E): Information Technology - Open Systems Interconnection – The Directory: Authentication Framework, June 1997.

[7] S. Tuecke et al., "Internet X.509 Public Key Infrastructure (PKI) Proxy Certificate Profile," RFC 3820, http://www.ietf.org/rfc/rfc3820.txt

[8] Overview of the Grid Security Infrastructure, http://www.globus.org/security/overview.html, Accessed October 2008.

[9] R Alfieri et al., "From gridmap-file to VOMS: managing authorization in a Grid environment," Future Generation Computer Systems 21 (4) pp549–558 (2005)

[10] R. Alfieri et al., "Managing Dynamic User Communities in a Grid of Autonomous Resources," Proceedings of the Computing in High Energy and Nuclear Physics conference, 24-28 March 2003, La Jolla, California, USA (TUBT005, ePrint cs.DC/0306004)

[11] T. Röblitz et al., "Autonomic Management of Large Clusters and Their Integration into the Grid," Journal of Grid Computing 2 247260 (2004)

[12] VO Services Project Home Page, http://www.fnal.gov/docs/products/voprivilege/, Accessed October 2008.

[13] M. Lorch et. al., "Authorization and account management in the Open Science Grid," Proceedings of the 6th IEEE/ACM International Workshop on Grid Computing, IEEE, 2005 (8pp), ISBN: 0-7803-9492-5

[14] D. Thain, T. Tannenbaum, and M. Livny, "Distributed Computing in Practice: The Condor Experience," Concurrency and Computation: Practice and Experience, Vol. 17, No. 2-4, pages 323-356, February-April, 2005.

[16] "SAML Specifications," http://saml.xml.org/saml-specifications, Accessed October 2008.

[17] "OASIS XACML TC," http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml, Accessed October 2008.

[18] M. Altunay et. al., "An XACML Attribute and Obligation Profile for Authorization Interoperability in Grids," FNAL Doc DB 2685-v1,

Fermilab, 2008 (40 pp), http://cd-docdb.fnal.gov/cgibin/ShowDocument?docid=2685

[19] L. Daigle et al., "URN Namespace Definition Mechanisms," RFC 2611, http://www.ietf.org/rfc/rfc2611.txt

[20] I. Sfiligoi, "Making Science in the Grid World: Using glideins to Maximize Scientific Output," Nuclear Science Symposium Conference Record, 2007. NSS '07. IEEE 2, Honolulu, HI, USA, 2007, pp. 1107-1109, ISBN 978-1-4244-0923-5

[21] T. Maeno et al., "PanDA: distributed production and distributed analysis system for ATLAS", *Journal of Physics: Conference Series* **119** (2008) 062036 (4pp), http://www.iop.org/EJ/abstract/1742-6596/119/6/062036

[22] A. Tsaregorodtsev, V. Garonne, and I. Stokes-Rees, "DIRAC: A Scalable Lightweight Architecture for High Throughput Computing," Fifth IEEE/ACM International Workshop on Grid Computing (GRID'04), 2004, pp. 19-25

[23] Internet2 / OpenSAML: http://opensaml.org, Accessed October 2008.

[24] The OGF OGSA-Authorization Working Group: http://forge.gridforum.org/sf/projects/ogsa-authz , Accessed October 2008

[25] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol - Version 1.2," RFC 5246, http://www.ietf.org/rfc/rfc5246.txt

[26] B. Lang, I. Foster, F. Siebenlist, R. Ananthakrishnan, and T. Freeman, "A Multi-policy Authorization Framework for Grid Security," pp. 269-272, Fifth IEEE International Symposium on Network Computing and Applications (NCA'06), 2006.

[27] The Site Central Authorization Service information page. http://www.nikhef.nl/grid/lcaslcmaps/scas/Accessed October 2008

[28] M. Feller, I. Foster, and S. Martin, "GT4 GRAM: a Functionality and Performance Study," Proceedings of TeraGrid 2007 Conference, Madison, WI.

The following government licenses should be removed before publication: