Writing Access Policies to Protect Web Services of a Spatial Data Infrastructure Using GeoXACML

Emerson M. A. Xavier
Brazilian Army Geographic Service
Quartel-General do Exército, bl F
Brasilia, Brazil
emerson@dsg.eb.mil.br

Wladimir S. Meyer
Brazilian Army Geographic Service
Quartel-General do Exército, bl F
Brasilia, Brazil
meyer@dsg.eb.mil.br

Abstract

The Brazilian Spatial Data Infrastructure (Infraestrutura Nacional de Dados Espaciais - INDE) was designed to disseminate geospatial data produced by public agencies, except for sensitive data. However, the INDE's Implementation Plan does not specify how to use access policies to available data. This paper proposes an access control system based on GeoXACML as a solution to achieve this purpose. This work also analyzes of the response time impact when using a policy-enabled geospatial Web server. The results validate this propose and out-line several possibilities of constraints. A case study with geodata in OGC Web services is used to show the benefits of the approach.

Keywords: access policies, geospatial web services, SDI, security, GeoXACML

1 Introduction

The Brazilian Spatial Data Infrastructure (INDE) was established under Presidential Decree No. 6666 of November 27, 2008. INDE is the integrated set of technologies, policies, standards and coordination necessary to promote access and dissemination of geospatial data produced by government agencies of the Executive Power [3].

The third article of Decree No. 6666/2008 determines that all geospatial data produced by the federal executive institutions should be disseminated along with their metadata. However, the first paragraph of this article excludes of this requirement the data whose confidentiality is vital to the security of society and state. On the other hand, some geographic data can contain sensitive information of interest to some agencies and should be visible in the INDE for a particular group of people. The data publicity is the rule, and restrict data is the exception, but this case may occurs.

A notable example is the secrecy in the distribution of statistical data through the INDE. Brazilian law stipulates that information collected from the population is sensitive, and should be used exclusively for statistical purposes. For example, a researcher registered in the system who needs statistical data on a larger scale (a census tract) may request this information while most of the other users only have access to data on the scale of a city. These facts point to the need of developing studies to establish an access policy to the geographic information that flows through the INDE's servers.

The INDE's Action Plan determines that the available data must be protected. But does not specify how to materialize and execute policies to access services and data available (who can access what). The producer agency is responsible for developing access policies to its data through sharing agreements and terms of use.

In this context, the issue is to establish a technological basis that allows materialize, effectively, access policies to available data in a Spatial Data Infrastructure (SDI).

A brief survey of information technology systems for spatial data on Web, under access policies perspective, presents the following facts:

- Open Geospatial Consortium (OGC) publishes open standards that enable the creation of state-of-the-art web-based Geographic Information Systems (GIS) solutions [1];
- The consortium Organization for the Advancement of Structured Information Standards (OASIS) has developed a policy language for access control called Extensible Access Control Markup Language (XACML) [7]; and
- Geographic XACML (GeoXACML) is an extension of XACML that incorporates the spatial data types and some spatial operations following the semantics of OGC open specifications [6].

Taking these facts as work assumptions, we propose that the GeoXACML standard configures a solution to materialize and execute access policies for the geospatial data available in a SDI. The aim is to enable the institutions involved in infrastructure building to publish its data (including sensitive) without compromising the principles of information security. Considering the performance as a representative requirement for geospatial services quality [5], this work also investigates the computational impact when access policies are enabled in common SDI services.

The remainder of this paper is structured as follows. Section 2 presents the concepts and theories that support the work. Section 3 presents the access policies described in GeoXACML. The experiment that validates the proposed solution is in Section 4. Finally, Section 5 brings a discussion of the results and work conclusions.

2 Background

Extensible Access Control Markup Language (XACML) is a specification of the OASIS consortium to establish an access control system for Service Oriented Architectures (SOA) [6]. The XACML defines an XML-based language for encoding access policies and an associated semantics [7]. This section brings a brief analysis of the XACML specification and its extension, the GeoXACML.

There are three elements of higher level in XACML: *PolicySet*, *Policy* and *Rule* [7]. A *PolicySet* has a set of elements of type *PolicySet* or *Policy*, along with the procedure for combining the results of each assessment. A *Policy* element contains a set of *Rules* and the procedure for combining the results of each rule. A *Rule* has a boolean expression that returns an authorization decision (permit or deny) if true. The specification enables to compose various types of policies and rules in several ways.

All these elements have an associated *Target*, which points to the components involved in its definition. A *Target* is divided into four distinct components that can be combined: *Subject*, *Resource*, *Action* and *Environment*. These components are based on the attributes definition. This means, to identify a *Subject*, for example, it is necessary to have an associated attribute. Having identified the attribute, then the system can get its value and process the access rules.

To illustrate the *Target* operation we can take for example the following access policy: "*John* can *see* (GetMap) *the map of Brazil* during the *morning* (7 am to 12 pm)". In this policy we have identified several target components: "John" is the value of an attribute descriptive for Subject; "GetMap" is the value of an Action attribute; "brazil_map" is the value of a Resource attribute; and "7 am to 12 pm" is the value of an Environment attribute. Figure 1 graphically displays this access policy.

Looking at Figure 1 we can see the importance of attributes for XACML systems. The user "John" was identified by the "subject-id" attribute. The "GetMap" operation was recognized by the "action-id" attribute. The desired feature type is indicated by the value "brazil_map" of the "resource-id" attribute. Finally, the transaction time is obtained by the "current-time" attribute, which is undergone to a temporal predicate.

GeoXACML is an OGC specification that extends the XACML to define a type of geometrical data and new functions to handle such data [6]. The authors say that this specification can be used to build interoperable access control systems for geospatial applications, especially for a SDI.

This specification provides examples of use for protecting Web services defined by OGC. However, this document does not include the names and possible values of the attributes associated with services.

3 Vocabulary for Access Policies with GeoXACML

This work proposes a solution that allows an effective use of GeoXACML to protect geographic Web services into a SDI. This solution includes a vocabulary that associates OGC Web services requests to (Geo)XACML components. The service specifications used in this work are Web Map Service (WMS), Web Feature Service (WFS) and Web Coverage Service (WCS). The requests to these services can be encoded in both Keyword Value Pair (KVP) and XML formats.

The attributes are essential to the semantics of XACML. Thus, this paper proposes a vocabulary to associate the parameters of OGC services requests (such as WMS, WFS or WCS) to interpretable attributes in XACML policies and rules. Two general guidelines were used to define the vocabulary: 1) take advantage of the attributes defined in XACML; 2) map the parameters of KVP requests (whenever possible, or XML in other cases) of OGC Web services into XACML attributes.

The data types in this vocabulary are defined in the XML Schema specification [2]. Just the "geometry" type is described in the GeoXACML specification [6].

3.1 Attributes in XACML Domain

All attributes that identify a Subject or Environment can be used in this solution. The Resource identification depends on which service is requested. Table 1 shows the semantics proposed for the attribute "urn:oasis:names:tc:xacml:1.0:resource: resource-id" for each OGC service.

Figure 1: Graphic example for an access policy.

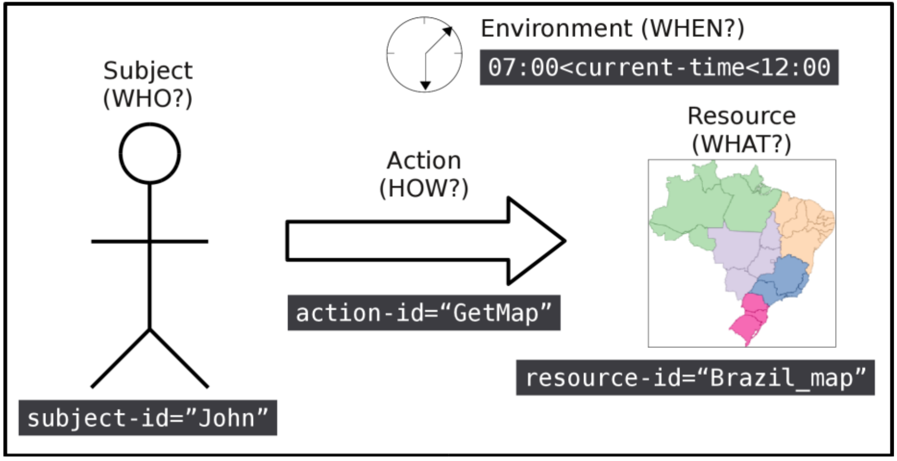

Table 1: “resource-id” attribute description.

Service	Data type	Meaning – parameter (operation)
WMS	string	LAYER (GetMap) or QUERYLAYER (GetFeatureInfo)
WFS	string	TYPENAME parameter (DescribeFeatureType, GetFeature, GetFeatureWithLock, LockFeature and Transaction)
WCS	string	IDENTIFIER(S) parameter (DescribeCoverage and GetCoverage)

The identification of the Action component is directly related to the request performed for each service. The attribute “urn:oasis:names:tc:xacml:1.0:action:action-id” receives the value of REQUEST parameter sent to each service. Table 2 brings the intended meaning.

Table 2: Values for “action-id” attribute.

Service	Data type	Possible values (operations)
WMS	string	GetCapabilities, GetMap or GetFeatureInfo
WFS	string	GetCapabilities, DescribeFeatureType, GetFeature, GetFeatureWithLock, LockFeature or Transaction
WCS	string	GetCapabilities, DescribeCoverage or GetCoverage

The semantics presented in this subsection is sufficient to write access policies associating users (Subject) executing operations (Action) on resources (Resource) at a given time (Environment). An example is “John Doe may see the layer Highways”.

But the XACML semantics can be extended to the OGC domain in order to accept various other parameters as can be seen in the next subsection.

3.2 Attributes in OGC Domain

The semantics of the attributes defined in the XACML specification are commonly used, they are not directed to the geographical domain. The GeoXACML specification also does not define the attributes for the OGC Web services, limited to only set the “geometry” type and its functions. This subsection presents a proposal for new attributes in OGC Web services domain for use in systems based on these standards. All proposed attributes fit in the Action component.

To keep the same format of the XACML standard attributes we propose to use the highest level Universal Resource Name (URN) used by OGC in some of its specifications: “urn:ogc:def”. This URN must be followed by the specification or service abbreviation of the request and its parameters, all lowercase. Figure 2 presents and exemplifies the general structure of a proposed OGC domain attribute.

Figure 2: General structure (a) and examples (b) of attributes in OGC domain.

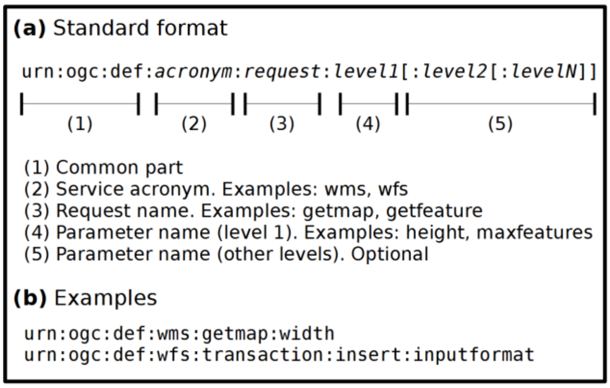

(a) Standard format

```
urn:ogc:def:acronym:request:level1[:level2[:levelN]]
```

(1) Common part
(2) Service acronym. Examples: wms, wfs
(3) Request name. Examples: getmap, getfeature
(4) Parameter name (level 1). Examples: height, maxfeatures
(5) Parameter name (other levels). Optional

(b) Examples

```
urn:ogc:def:wms:getmap:width
urn:ogc:def:wfs:transaction:insert:inputformat
```

The name of the parameters must obey, whenever possible, the parameter name written in KVP format of the corresponding service specification. If this encoding is not available, the parameter name as it appears in the XML schema that describes the request should be used. In the latter case, more than one level in the URN can be used to precisely describe the parameter. As example, Table 3 presents some attributes proposed for a WMS GetMap request.

Table 3: Attributes for a WMS GetMap request.

Attribute identifier	Data type	Meaning
urn:ogc:def:wms:getmap:crs	string	CRS (SRS) parameter
urn:ogc:def:wms:getmap:bbox	geometry	BBOX parameter
urn:ogc:def:wms:getmap:width	integer	WIDTH parameter
urn:ogc:def:wms:getmap:height	integer	HEIGHT parameter
urn:ogc:def:wms:getmap:format	string	FORMAT parameter

This approach allows defining rules that directly use the request parameters, increasing the scope to protect available data through more specific policies. The adoption of these attributes allows writing rules like “John Doe may draw (WMS GetMap) layer Highways if requests images with a maximum of 512 pixels wide (‘width’ parameter)”.

4 Proof of Concept: Securing Geospatial Web Services

This experiment consists of verifying the validity of our hypothesis by examining the use of XACML access policies for different geospatial Web services. We used two computers: a server that contains the spatial data and services and a client that accesses these resources.

The map server used in this experiment is built over the TerraLib library [4]. The Web-GIS has a policy enforcement point (PEP) based on a mediated architecture. The geodata are stored in a PostgreSQL/PostGIS database. The Apache JMeter application is used to measure the performance of server access.

This experiment uses vector geographic data (access via WMS and WFS) and raster (accessed via WCS). We selected four layers of vector data from four topographic maps of a region near the Brasilia city: *Contour_Line*, *Water_Mass*, *Building* and *Road*. The raster data is composed of a mosaic of GeoCover Landsat images covering an area equivalent to the vector data. The study area was divided into 25 cells of 20 by 20 km. Requests for services must comply with this partitioning in order to verify its validity.

4.1 Defining Access Policies

In this experiment we wrote three policies, one for each service to protect. The WMS service has three rules, while the WFS and WCS services have one. We chose simple rules to test different attributes as listed in the Table 4.

Table 4: Policies and rules defined.

Policy target	Rule meaning	Rule ID
WMS	Any user may access the *Contour_Line* layer	1
WMS	Users outside the *Hidrography* group are not allowed to access the *Water_Mass* layer	2
WMS	Any user can see the *Building* layer, except within the Federal District (represented by a polygon)	3
WFS	The *Road* layer may only return a maximum of 250 features on each request	4
WCS	The *GeoCover* layer may not be accessed along the São Bartolomeu River (represented by a line)	5

Each policy was translated to the XACML form, using Resource, Subject and Action attributes. There is no Environment-driven rule, so this element may be ignored. The Table 5 summarizes the defined rules.

Table 5: Policies and rules defined.

Rule ID	Resource	Subject	Action	Rule Effect
1	Contour_ Line	any	any	permit
2	Water_ Mass	!=Hidro graphy	any	deny
3	Building	any	GetMap within a polygon	deny
4	Road	any	GetFeature Features<250	permit
5	GeoCover	any	GetCoverage intersects a line	deny

All accesses were performed by an unauthenticated user to the system. For each rule, 25 requests were executed against the target layer for the desired service. All responses were in accord with the corresponding rule. The rules 3 (layer Building) and 5 (layer GeoCover) relied on spatial predicates to restrict access. Some requests were permitted and other denied when policies enabled as expected.

4.2 Performance Issues

In order to check the influence of using the access policies into the system's performance, were executed two types of request: direct access (without the mediator) and policy activated (mediator with XACML module turned on). This method aims to identify the differences in performance when using a PEP to access the services, since the response-time is a key-aspect in web services.

Figure 3: Response time to direct access and with policies enabled.

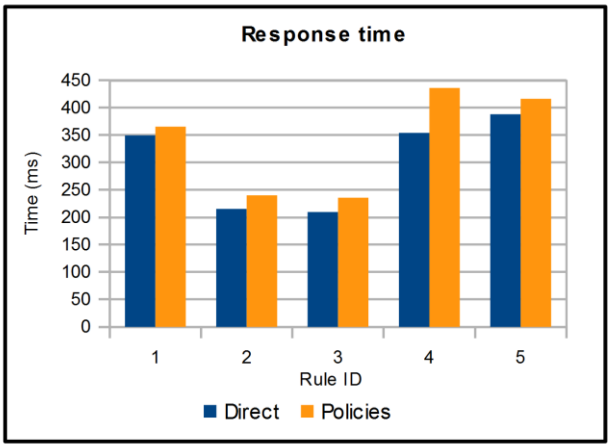

The graph in the Figure 3 indicates that the average response time has grown about 12%, when using policies in a mediated architecture comparing with the direct access. This increase of response time may be caused by the mediator and PEP module.

5 Conclusions

The solution presented here includes a vocabulary that associates the OGC Web services requests to XACML access policies. The experiment points to the validity of this solution to successfully employ various types of restriction on services: subject, resource and action; scalar and geometric predicates.

This article targets to investigate the adoption of XACML and its GeoXACML extension to protect the access to geospatial Web services. This need arose with the implementation of the Brazilian SDI and with the needs, of some producers, to protect part of their geospatial data. The main contribution of this paper is to present a vocabulary for mapping service requests into access policies attributes. Another contribution is to investigate the computational impact when access policies are adopted in these services.

The results from the experiment indicate that the system used as PEP properly validated all policies and their associated rules. The WMS, WFS and WCS services may be secured by policies written in (Geo)XACML when used the corresponding vocabulary.

In Web services, performance is a key issue. The results of experiment show that noticeable computational cost exists with policies enabled comparing to direct access.

The limitations of this solution are related to the proposed vocabulary, which includes only the attributes of Web services requests. The current rules and policies do not include the data properties. This limitation can be addressed by defining a new set of attributes. However, the additional cost of a solution in this direction can directly affect the performance of the PEP, since it needs to get all the data before applying the policies.

Technologies commonly used in information security – such as digital signature, digital certificate and public key infrastructure – solve a type of problem: to identify who is trying to access the system. What this user may actually do in this system is the next challenge. Most e-government systems have few operations and access profiles associated with a rigid data model. But access to heterogeneous spatial data, produced by different institutions, at different times – as a SDI environment – requires a more complex access policies system.

This work presents a solution based on open standards and free software to effectively address this complexity and is being adopted by Geographic Service of the Brazilian Army, one of the geospatial data producers from the Brazilian SDI.

References

[1] G. Anderson and R. Moreno-Sanchez. Building Web-based spatial information solutions around open specifications and open source software. *Transactions in GIS* 7(4):447-466, 2003.

[2] P. Biron, and A. Malhotra. *XML Schema Part 2: Datatypes* Second Edition. 2004. http://www.w3.org/TR/2004/REC-xmlschema-2-20041028. Accessed 13 Nov 2012.

[3] Brazil. *Decreto nº 6.666, de 27 de novembro de 2008*. 2008. http://www.planalto.gov.br/ccivil_03/_Ato2007-2010/2008/Decreto/D6666.htm. Accessed 13 Nov 2012.

[4] G. Câmara, L. Vinhas, K. Ferreira et al. TerraLib: An Open Source GIS Library for Large-scale Environmental and Socio-economic Applications. In G. Hall and M. Leahy, editors, *Open Source Approaches to Spatial Data Handling*, pages 247-270. Berlin, Springer-Verlag, 2008.

[5] INSPIRE. *Network Services Architecture*, Version 3.0. 2008. http://inspire.jrc.ec.europa.eu/reports/ImplementingRules/network/D3_5_INSPIRE_NS_Architecture_v3-0.pdf. Accessed 13 Nov 2012.

[6] A. Matheus and J. Herrmann. *Geospatial eXtensible Access Control Markup Language (GeoXACML)*, Version 1.0. 2008. http://portal.opengeospatial.org/files/?artifact_id= 25218. Accessed 13 Nov 2012.

[7] T. Moses. *eXtensible Access Control Markup Language (XACML)*, Version 2.0. 2005. http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf. Accessed 13 Nov 2012.