Ontology-Based Matching of Security Attributes for Personal Data Access in e-Health

Ioana Ciuciu[1], Brecht Claerhout[2], Louis Schilders[2], and Robert Meersman[1]

[1] Semantics Technology and Applications Research Laboratory, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium
[2] Custodix
Kortrijksesteenweg 214 b3 Sint-Martens-Latem Belgium 9830
{iciuciu,meersman}@vub.ac.be
{louis.schilders,brecht.claerhout}@custodix.com

Abstract. This paper discusses an interoperability solution (tool) for the internal management of a policy decision engine located at the level of the authorization layer of a service oriented environment. The tool aims to support federated access control in the context of distributed architectures, in which a local authorization policy is not able to recognize all the attributes in the authorization decision requests. The approach is based on an ontology-based interoperation service (OBIS) whose role is to translate security attributes (name-value pairs) from local security vocabularies into the attributes recognized by the central (Master Policy Decision Point) vocabulary based on a security ontology and its domain-specific extensions which provides semantic reasoning services. The approach is validated in an e-Health scenario for the access of patient data for diabetes patient monitoring and disease management.

Keywords: Authentication, Authorization, Security Policies, Ontology, Ontology-based Data Matching, e-Health.

1 Introduction and Motivation

Among the challenges of the Trusted Architecture for Securely Shared Services (TAS^3)[1] project is the interoperability of different access control policies in the context of unified distributed architectures. In this setting, every stakeholder organization describes its authorization policy using its organization-specific vocabulary, and when a policy engine receives an authorization request containing unknown terms, it semantically matches these with the ones locally known by the authorization policy.

In our previous work we have proposed an ontology-based interoperation service (OBIS [1]) which calculates the matching of security concepts extracted from access requests and local authorization policies. This study extends OBIS by proposing a method for mapping the security attributes (name-value pairs specifying the subject, resource and action) corresponding to a local security vocabulary into security

[1] http://tas3.eu/

R. Meersman, T. Dillon, and P. Herrero (Eds.): OTM 2011, Part II, LNCS 7045, pp. 605–616, 2011.

attributes recognized by the central Policy Decision Point (PDP) using a policy ontology and ontology-based matching strategies. The extension of OBIS is called OBIS Domain Mapper. The ontology is grounded in natural language, which enables individuals from different organizations to express their security policies in an intelligible way, thus enforcing the user-centricity.

The proposed method is illustrated on the ontological representation of XACML policies, but the approach applies to other policy languages.

The use case is created with one of the TAS^3 test beds, the Custodix Healthcare demonstrator [2].

The rest of the paper is organized as follows: Section II describes related work; Section III provides background information on the technology being used. The requirements and use case are presented in Section IV. Section V proposes a method for attribute mapping based on ontology-based data matching techniques. Section VI presents our conclusion and suggestions for future work.

2 Related Work

Several approaches exist which aim at resolving semantic access control.

The Semantic Access Control (SAC) Model [3] was specifically designed to enforce ABAC policies in heterogeneous and distributed environments. It maps policies to resources dynamically based on the semantics of policies and resources. The Semantic Access Control Enabler (SACE) [4] was developed to enforce Role-Based Access Control (RBAC) when accessing heterogeneous data from databases.

Verma [5] presents a semantic policy matchmaking for web service policies specified across multiple domains (e.g. security, privacy, trust). KAoS [6] is a semantic policy language and a framework for the specification, management and enforcement of policies within different security domains. A similar approach is presented in [7], concerned with the meaning of contexts to be used directly in an access control policy.

Several approaches propose [8,9,10] semantic reasoning services for policy management based on Semantic Web technologies.

Our approach is slightly different, proposing its own paradigm for semantic reasoning based on an ontology grounded in natural language and on ontology-based data matching strategies.

3 Background

In this section we provide relevant background knowledge related to our approach, namely the knowledge and constraints representation and the policy language used.

3.1 DOGMA Approach for Ontology Engineering

The common understanding of security policies in this study is based on the Delevoping Ontology Grounded Methodology and Applications (DOGMA, [11]). DOGMA is

a formal ontology engineering framework applying the principles of database design methodology (NIAM/ORM2, [12]) to ontology engineering. DOGMA ontology is grounded in natural language and based on the *double articulation principle* [13], which makes the ontology two layered:

1. The *lexon* base layer, containing a set of simple binary facts, called lexons, which are expressed in semi-natural language;
2. The *commitment* layer that formally defines rules and constraints by which applications may make use of the lexons from the lexon base.

A lexon is defined as a quintuple $\langle \gamma, t_1, r_1, r_2, t_2 \rangle$ representing a fact type. γ is a context identifier that points to a context where two terms, t_1, t_2 are originally defined and disambiguated. r_1, r_2 are two roles that characterize the relationship between t_1 and t_2. For example, *‹ABAC, Subject, performs, performed by, Action›* is a lexon which means "in the context of ABAC, a Subject performs and Action and an Action is performed by a Subject". Table 1 illustrates high level concepts of an ABAC (Attribute Based Access Control) policy represented with lexons.

Table 1. Lexon representation of Subject, Action and Target in the ABAC model

ABAC Policy			
Head term	**Role**	**Co-role**	**Tail term**
SecurityPolicy	controls	controlled by	Action
SecurityPolicy	has	of	Target
SecurityPolicy	written by	writes	Subject
Action	performed by	performs	Subject
Action	performed on	under	Resource

A commitment contains a constraint on a (set of) lexon(s). For instance, we can apply the cardinality constraint on the above lexon, – "only one value is allowed for the action attribute". The commitment language needs to be specified in a language such as OWL[2] or SDRule language [14].

The lexons together with the commitments can be further converted to RDF[3] and OWL in order to make the ontology processable by other applications and by widely adopted semantic reasoners (e.g. Pellet [15]).

3.2 XACML

The eXtensible Control Markup Language (XACML [16]) is an OASIS standard language and architecture for the expression and exchange of access control policies, decision requests and responses.

The XACML policy language is structured in three levels of elements: policyset, policy and rule. A policyset comprises a set of policysets and/or policies, a target, obligations and a policy combining algorithm identifier. A policy comprises a set of rules, a target, obligations and a rule combining algorithm identifier. Finally a rule comprises a

[2] http://www.w3.org/TR/owl-ref/
[3] http://www.w3.org/RDF/

condition, a target and an effect. The target component found in each element type identifies the set of *subjects*, *resources, actions* and *environments* to which it applies.

As illustrated in Table 1, a subject (e.g. physician) requests permission to perform an action (e.g. read) on a resource (e.g. medical diary). A rule is a mapping from a target to a decision, whose value can be either *Permit* or *Deny*. A rule combining algorithm is used to resolve conflicts among all the rules which are applicable and which have different effects.

An *attribute* is the basic unit of an XACML policy. Attributes are characteristics of the subject, resource, action or environment of the access request. An XACML access request therefore consists of a list of attributes-value pairs. The mappings in this study are done between attribute-value pairs in a request (at the level of the central PDP) to attribute-value pairs in the local authorization policy.

4 Requirements and Use Case

4.1 Health Information Network Requirements

The present study is done in the context of the TAS^3 authorization architecture. TAS^3 is a framework for protecting personal data in service oriented environments. It focuses on interoperability and aims to deliver a generic solution useful in a wide range of application domains. At the level of the authorization layer this translates into semantic support for different policy decision engines and policy languages.

TAS^3 primarily puts people into control over their personal data in a service oriented architecture. The e-Health pilot demonstrates how TAS^3 accomplishes this objective in the highly regulated e-health environment, where user centric personal data management translates into:

(1) The possibility for patients to adjust the default e-health domain policies determined by legislation and ethical guidelines, so that their personal data is protected according to their personal preferences on data protection;
   - For example: a patient should have the option not to disclose mental health related information from a replacement physician (out-of-office hours).
   - However, in a highly regulated environment such as healthcare, personal freedom to hide or disclose health information has its boundaries (e.g. where hiding it could result in bad treatment or damage the treating healthcare professionals). These need to be taken into account.

(2) The possibility for data users to query patients for specific (extraordinary) access requests for data processing (e-consent).
   - For example: a patient could be invited to share existing data into a clinical study.

4.2 Use Case: Federated Data Access

The demonstration environment (Fig. 1) was modeled according to the “distributed health repositories with central access” concept, which forms the basis of many

e-health information sharing initiatives in the EU and the US. In particular, the use case was staged in a Belgian setting.

Central to the system is the Patient Information Location Service (PILS [2]) which is used by professionals (e.g. medical doctors, researchers) to find patient information in distributed repositories. Two types of repositories have been connected in this demonstrator: (1) hospital results servers and (2) summary record repositories, as illustrated in Fig. 1.

Multiple Identity Providers (IdP) exist in the trial, all of which are authoritative with respect to unique user identities and to unique healthcare professional identifiers (similar to the actual situation in Belgium). Finally, there is a privacy management center where patients can set access policies on their personal data.

In the privacy center, patients can specify their personal privacy preferences which are then to be enforced over the health information sharing network for health professionals. The preferences set by the individual patients are translated into XACML policies which are loaded into a central Policy Decision Point (PDP). Service providers participating in the health network are required to forward access requests to that central PDP (if they involve resources covered by the central PDP policies).

Apart from the differences in authorization frameworks, different service providers use different security vocabularies (attribute-value pairs describing the subject, resource and action). In the demonstrator, the OBIS Domain Mapper service instances are responsible for translating local security vocabularies (name value pairs) into the "domain" vocabulary, used in by the central PDP.

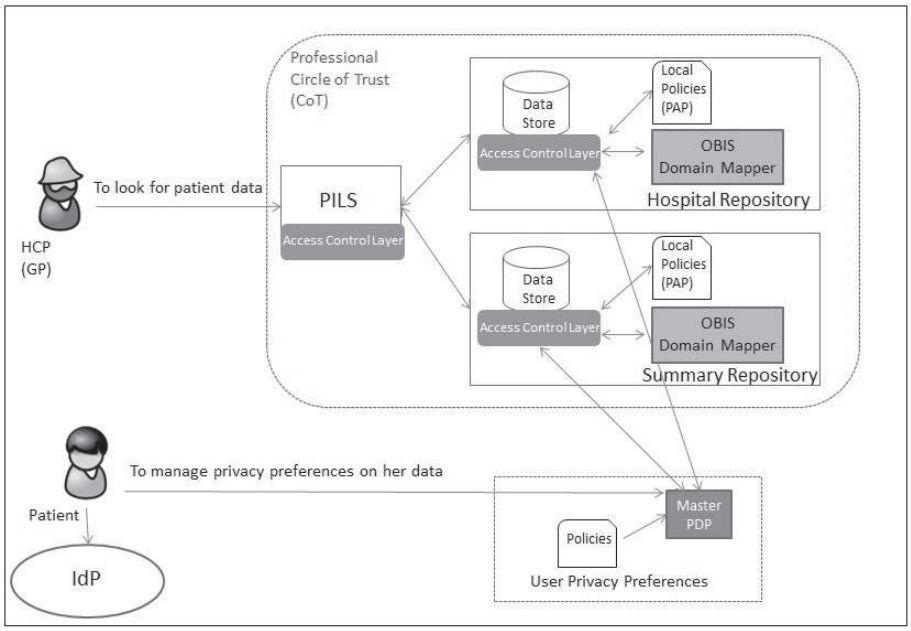

Fig. 1. Health Care use case scenario

Table 2 exemplifies request content of an access request to the central (Master) PDP in the first scenario.

Table 2. Example of content of an access request to the central PDP

Request			Decision
Subject	**Action**	**Resource**	
Subject_type = GP	Action_type = Read	Resource_type = labresult	Permit
Subject_type = GP	Action_type = Create	Resource_type = document	Permit
Any	Any	Any	Deny

5 Ontology-Based Attribute Matching for Access Control

Here we explain how we extend our previous ontology-based interoperation service to support mappings between the attributes in a decision request and the attributes from the access control rules in a local access policy. The method, the tools and the underlying technologies are presented.

5.1 Ontology-Based Interoperation Service (OBIS)

OBIS was initially designed as a web service located in the authorization architecture of TAS^3. It provides an interface to perform relation lookups between two terms represented as URIs, corresponding to the Service Requestor (SR) and to the Service Provider (SP) respectively, in order to determine the level of dominance between them.

Given e.g. a name of a resource, OBIS semantically infers the object class of the resource and computes how the authorization propagates in the (role/attribute/action) inheritance hierarchy, while enforcing the constraints in the ontological commitments.

This study proposes an extension of the OBIS service, called OBIS Domain Mapper, which translates the security attributes in the local PDPs into attributes recognized by the central (master) PDP in the TAS^3 authorization architecture. The main difference between the original OBIS service and the one proposed in this paper is that the first one only returns a code indicating the domination relation between two security concepts originating from different policy languages, while with the second approach the mapping between the two security domains (languages) is also provided. This approach is described in the next section.

5.2 Security Attributes Mapping

A method is proposed here for the mapping of security attributes using the Ω-RIDL Mapping Generator tool [17,18]. Ω-RIDL takes in input an XML file representing the access decision request and an ontology file (lexons and commitments) representing the access control policy ontology and returns a Ω-RIDL mapping file which maps the security attributes to concepts in the ontology.

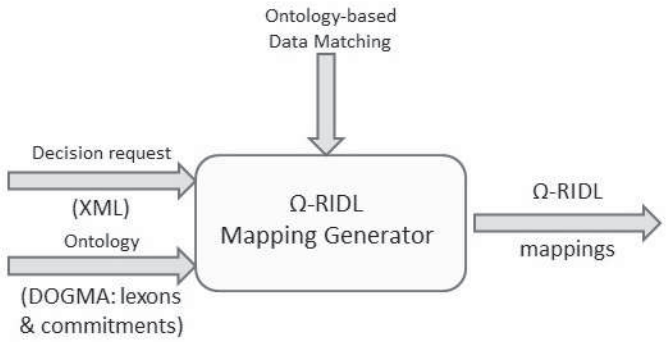

Fig. 2. Ω-RIDL mapping generator architecture

The Ω-RIDL mappings are obtained by applying ontology-based data matching strategies at (1) string level (fuzzy literal similarity, e.g. JaroWinkler); (2) lexical level (synonymous similarity, e.g. based on WordNet[4]) and (3) ontology (lexon graph) level (semantic similarity) in this order (refer to [19] for details on ontology-based data matching strategies).

Ω-RIDL is designed as a web service which is called by the OBIS Domain Mapper service in order to infer the mapping of security attributes between two domains. OBIS Domain Mapper sends a bag of security attributes (name-value pairs representing the subject, resource and action) corresponding to a domain Dom1 in input to Ω-RIDL which performs semantic inference and ontology-based data matching operations and returns another bag of attributes corresponding to another domain, Dom2, as shown in Fig. 3. Previous to calling Ω-RIDL, OBIS performs an explicit translation (mapping) from the local terminology (Dom1, Dom2) to the core ontology (lexon graph), based on the user-defined dictionaries. Then Ω-RIDL performs semantic inference operations on the ontology graph in order to infer the mappings from Dom1 to Dom2. For the moment being we only consider one-to-one mapping of attribute-value pairs. The one-to-many and many-to-one mappings of attribute-value pairs are ongoing work.

Fig. 3. Ω-RIDL invocation

[4] http://wordnet.princeton.edu/

Below we provide an example of a XACML policy rule that returns `Permit` for access requests that have value `physician` for attribute `subject`, value `read` for attribute `action` and value `summary information` for attribute `resource` (see example from Table 2).

```
<Target>
  <Resources>
    <Resource>
      <ResourceMatch
        MatchId="function:string-equal">
        <AttributeValue
          DataType="#string">summary_information
          </AttributeValue>
        <ResourceAttributeDesignator
          AttributeId="…information-class "
          DataType="#string" />
        </ResourceMatch>
      </Resource>
    </Resources>

  <Actions>
    <Action>
      <ActionMatch
        MatchId="function:string-equal">
        <AttributeValue
          DataType=="#string">read
        </AttributeValue>
      <ActionAttributeDesignator
        AttributeId="…action"
        DataType="#string" />
      </ActionMatch>
    </Action>
  </Actions>
</Target>

<Rule RuleId="SenderIsPhysician" Effect="Permit">
  <Description>Physicians can create new documents
  </Description>
  <Condition>
    <Apply>
      FunctionId="function:string-is-in"
      <AttributeValue>
        DataType="#string">physician
      </AttributeValue>
      <SubjectAttributeDesignator
        DataType="#string"
        AttributeId="hcp-type"/>
```

```
    </Apply>
  </Condition>
</Rule>
```

When a physician tries to view a `vaccination fiche` in one of the repositories through the health information sharing infrastructure, the following happens: Inside the contacted repository an access control request for a `read` action on a `vaccination fiche` (resource) is triggered (to eventually determine if the `vaccination fiche` can be shown to the physician). This request is to be evaluated by a local access control decision engine according to locally formulated policies (which do e.g. also deal with access rules for "locally originated" requests). However, for this type of access through the health information network, also the central PDP needs to be queried (access control decisions by different engines are eventually to be combined).

The security vocabulary used in the local repository (which is typically implementation specific) is not aligned with the more generic security vocabulary used in the wider health domain (used in the policies handled by the central PDP). The local access control request can thus not be evaluated as such by the central PDP.

A generic approach to translation of access requests from one domain vocabulary to another is provided by the OBIS Domain Mapper. This component translates attributes (e.g. XACML name value pairs) from one security domain to another (here from a "local" repository into the vocabulary used as reference in the distributed environment).

More specifically, in the described example, the OBIS Domain Mapper will look into the ontology hierarchy and constraints via the Ω-RIDL mappings and will infer that a `vaccination fiche` document (as known in the repository) classifies as `summary information` according to the central policy vocabulary.

5.3 Access Control Policy Ontology

Fig. 4 illustrates the access policy ontology used in the e-Health scenario. The focus in the figure is on the "target" concept, showing its constituents hierarchically (see the circles). The semantic relations are of the type 'part-of' and 'is-a' (which grow or shrink a set), indicating the (security-specific) domination relation between the concepts. OBIS computes the domination relation between two concepts in the ontology using AND/OR graphs. The figure includes core concepts from the TAS^3 security ontology (e.g. 'subject', 'action', 'resource') linked to application-specific concepts derived from the e-Health scenario (e.g. 'patient-id', 'hcp-type').

Every component in the above described scenario commits to this ontology. Every stakeholder organization (or department), must provide a mapping file between its own terminology and the core ontology. This task is the responsibility of the security officer of every participant organization. The mapping files will serve to translate the local concepts to central ones as a preliminary step before performing the ontology-based data matching (inference) with Ω-RIDL.

Table 3 shows mappings between the local policies of the Hospital data repository and the central ontology. The first mapping concerns a subject mapping from the local

hospital repository A, where the subject attributes are expressed as `name = employee_type` and `value = nurse`, to the central vocabulary used by the Master PDP, where `name(employee_type)` maps to `hcp-type` and `value(nurse)` maps to `nurse`.

The second row shows a resource attribute mapping from the local vocabulary of hospital repository A, where the resource `name = file_type` and the resource `value = vaccination type`, to the central vocabulary where `name` maps to `document-type` and `value` maps to `summary information`.

Table 3. Mappings between the local (organization-specific) terminology and the core ontology

Repository	Term	Concept in the ontology
Hosp. RepositoryA	employee_type = nurse	hcp-type = nurse
Hosp. RepositoryA	file_type = vaccination fiche	document-type = summary information
HIV Center	Doc-type = medication fiche	Sensitivity-indicator = HIV
	Location = Brussels	
	Patient-Name = Herve	

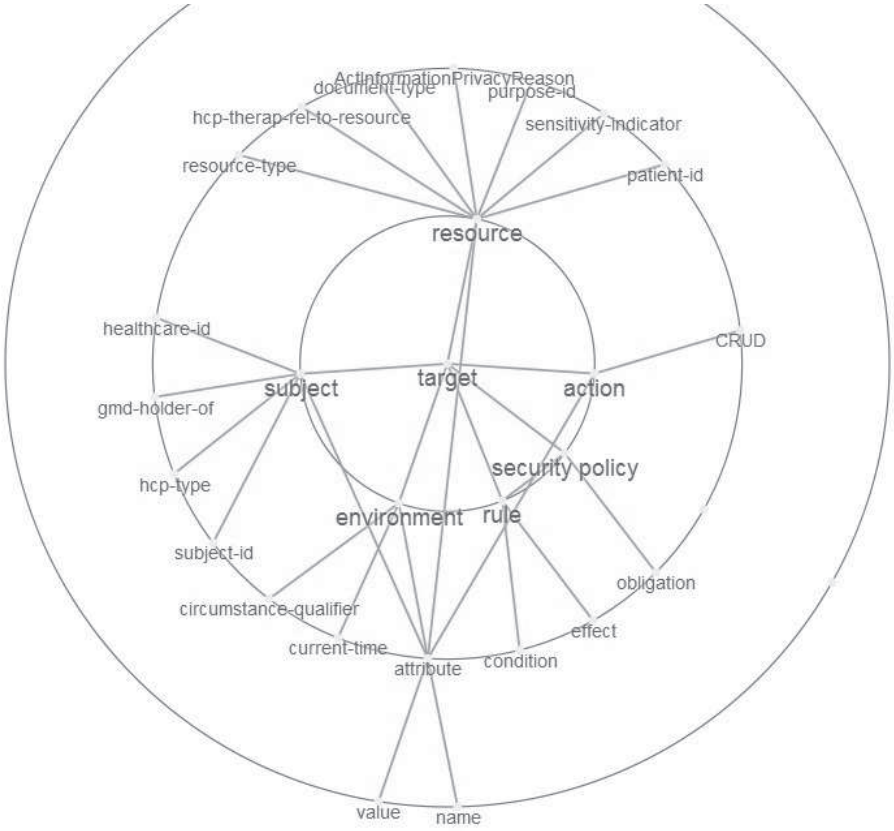

Fig. 4. The "Target" concept in the access control policy ontology

6 Conclusion and Future Work

The paper presents an extension, OBIS Domain Mapper, of a previously proposed ontology-based interoperation service which enables attribute mappings between local and central security vocabularies for the internal management of a policy decision engine in the context of service oriented architectures. The approach is based on the DOGMA ontology which enables semantic reasoning and on the Ω-RIDL mapping generator which performs the mapping based on ontology-based data matching strategies.

The OBIS service is (1) user-centric, enabling end-users to manage and protect their personal data through the creation of control access policies, without needing to know specific details about the security domains of remote service providers; (2) based on a security policy ontology grounded in natural language; (3) automated, through an integrated architecture which ensures OBIS is called by credential validation services and policy decision points; (4) autonomous, being designed as a web service to operate in an open, distributed and dynamic environment; and (5) secure, enabling query-only requests via SSL/TLS links.

Future work will involve the implementation of more sophisticated Ω-RIDL mappings by introducing additional constraints and evolving the ontology with more sophisticated authorization policies, including concepts such as obligations, delegation of authority, and separation of duty. The evaluation of the results with the ODMF (Ontology-based Data Matching Framework) evaluation benchmark is also planned as future work.

Additional functions, which extend the xacml rule engine by reasoning functions, are ongoing work.

Acknowledgments. This paper is supported by the EC FP7 TAS[3] (Trusted Architecture for Securely Shared Services) project. The authors would like to thank all TAS[3] project partners for their contribution to the research.

References

1. Ciuciu, I., Zhao, G., Chadwick, D.W., Reul, Q., Meersman, R., Vasquez, C., Hibbert, M., Winfield, S., Kirkham, T.: Ontology-based Interoperation for Securely Shared Services. In: Proc. IEEE Int. Conf. on New Technologies, Mobility and Security (NTMS 2011), Paris, France (2011)
2. Claerhout, B., Carlton, D., Kunst, C., Polman, L., Pruis, D., Schilders, L., Winfield, S.: Pilots Specifications and Use Case Scenarios, TAS[3], Deliverable D9.1, Trusted Architecture for Securely Shared Services (2010), `http://tas3.eu/`
3. Yague, M., Gallardo, M., Mana, A.: Semantic access control model: a formal specification. In: Proc. 10th European Symposium on Research in Computer Security, pp. 23–24 (2005)
4. Mitra, P., Liu, P.: Semantic access control for information interoperation. In: Proc. 11th ACM Symposium on Access Control Models and Technologies, pp. 237–246 (2006)
5. Verma, K., Akkiraju, R., Goodwin, R.: Semantic matching of web service policies. In: Proc. 2nd Int. Workshop on Semantic and Dynamic Web Processes, pp. 79–90 (2005)

6. Uszok, A., Bradshaw, J.M., Lott, J., Breedy, M.R., Bunch, L., Feltovich, P.J., Johnson, M., Jung, H.: New developments in ontology-based policy management: Increasing the practicality and comprehensiveness of KAoS. In: Proc. IEEE Workshop on Policies for Distributed Systems and Networks, pp. 145–152 (2008)
7. Dersingh, A., Liscano, R., Jost, A., Finnson, J., Senthilnathan, R.: Utilizing semantic knowledge for access control in pervasive and ubiquitous systems. Mobile Netw. Appl. 15, 267–282 (2010)
8. Damiani, E., De Capitani di Vimercati, S., Fugazza, C., Samarati, P.: Extending Policy Languages to the Semantic Web. In: Koch, N., Fraternali, P., Wirsing, M. (eds.) ICWE 2004. LNCS, vol. 3140, pp. 330–343. Springer, Heidelberg (2004)
9. Smith, M., Schain, A., Clark, K., Griffey, A., Kolovski, V.: Mother, May I? OWL-based Policy Management at NASA. In: OWLED (2007)
10. Ferrini, R., Bertino, E.: Supporting RBAC with XACML+OWL. In: SACMAT, pp. 145–154 (2009)
11. Spyns, P., Tang, Y., Meersman, R.: An Ontology Engineering Methodology for DOGMA. J. of App. Ontology 3(1-2), 13–39 (2008)
12. Halpin, T.: Information Modeling and Relational Databases: From Conceptual Analysis to Logical Design. Morgan Kaufmann, San Francisco (2001)
13. Spyns, P., Meersman, R., Jarrar, M.: Data Modeling Versus Ontology Engineering. SIGMOD Record: Special Issue on Semantic Web and Data Management 31(4) (2002)
14. Tang, Y., Meersman, R.: SDRule Markup Language: Towards Modeling and Interchanging Ontological Commitments for Semantic Decision Making. In: Handbook of Research on Emerging Rule-Based Languages and Technologies: Open Solutions and Approaches. IGI Publishing, USA (2009) ISBN: 1-60566-402-2
15. Sirin, E., Parsia, B., Grau, B.C., Kalyanpur, A., Katz, Y.: Pellet: A practical OWL-DL reasoned. J. of Web Semantics (2007)
16. OASIS "eXtensible Access Control Markup Language" (XACML) Version 2.0 OASIS Standard (2005)
17. Trog, D., Tang, Y., Meersman, R.: Towards Ontological Commitments with Ω-RIDL Markup Language. In: Ontologies, Databases and Applications of Semantics, Villamoura, Portugal (2007)
18. Verheyden, P., De Bo, J., Meersman, R.: Semantically Unlocking Database Content Through Ontology-Based Mediation. In: Bussler, C.J., Tannen, V., Fundulaki, I. (eds.) SWDB 2004. LNCS, vol. 3372, pp. 109–126. Springer, Heidelberg (2005)
19. Tang, Y., De Baer, P., Zhao, G., Meersman, R., Pudkey, K.: Towards a Pattern-Driven Topical Ontology Modeling Methodology in Elderly Care Homes. In: Meersman, R., Herrero, P., Dillon, T. (eds.) OTM 2009 Workshops. LNCS, vol. 5872, pp. 514–523. Springer, Heidelberg (2009)