



# SEMANTIC WEB SECURITY AND PRIVACY

<sup>1</sup>N K PRASANNA ANJANEYULU ANNA, <sup>2</sup>SHAIK NAZEER

<sup>1</sup>Institute of Public Enterprise, Hyderabad, India

<sup>2</sup>Bapatla College of Engineering, Bapatla, India

## ABSTRACT

The Semantic Web aims at enabling sophisticated and autonomic machine to machine interactions without human intervention, by providing machines not only with data but also with its meaning (semantics). In the Semantic Web, any two strangers can interact with each other automatically and therefore this assumption does not hold. In this setting, traditional security mechanisms are not suitable anymore. Hence, a semantically enriched process is required in order to regulate an automatic access to sensitive information. Policy-based access control provides sophisticated means in order to support protecting sensitive resources and information disclosure. This paper provides an introduction to policy-based security and privacy protection, by analyzing several existing policy languages. Furthermore, it shows how these languages can be used in a number of Semantic Web scenarios.

**Keywords:** *Semantic Web, Security, Privacy, Machine Internacitons*

## 1. INTRODUCTION

Information provided in the current Web is mainly human oriented. For example, HTML pages are human understandable but a computer is not able to understand the content and extract the right concepts represented there, that is, the meaning of the data. The Semantic Web [1] is a distributed environment in which information is self-describable by means of well-defined semantics, that is, machine understandable, thus providing interoperability (e.g., in e-commerce) and automation (e.g., in search). In such an environment, entities which have not had any previous interaction may now be able to automatically interact with each other. For example, imagine an agent planning a trip for a user. It needs to search for and book a plane and a hotel taking into account the user's schedule. When the user's agent contacts a hotel's website, the latter needs to inform the former that it requires a credit card in order to confirm a reservation. However, the user may probably want to restrict the conditions under which her agent automatically discloses user personal information. Due to such exchange of conditions and personal information, as well as its automation, security and privacy become yet more relevant and traditional approaches are not suitable anymore. On the one hand, unilateral access control is now replaced by bilateral protection. On the other hand, identity-based access control cannot be applied anymore

since users are not known in advance. Instead, entities' properties play a central role. Both these properties and conditions stating the requirements to be fulfilled by the other party, must be described in a machine-understandable language with well-defined semantics allowing other entities to process them. Systems semantically annotated with policies enhance their authorisation process allowing, among others, to regulate information disclosure (privacy policies), to control access to resources (security policies), and to estimate trust based on parties' properties (trust management policies) [2]. Distributed access control has addressed some of these issues though not solved them yet. RBAC (Role-Based Access Control) also does not meet Semantic Web requirements since it is difficult to assign roles to users which are not known in advance. Regarding to user's privacy protection, Platform for Privacy Preferences (P3P) provides a standard vocabulary to describe Web server policies. However, it is not expressive enough and it does not allow for enforcement mechanisms. On the other hand, there is a wide offer of policy languages that have been developed to date [3][4][5][6][7] addressing the general requirements for a Semantic Web policy language: expressiveness, simplicity, enforceability, scalability, and analyzability [8]. These policies can be exchanged between entities on the Semantic Web and therefore they are described using languages with well-founded semantics. The policy languages listed above

differ in expressivity, kind of reasoning required, features and implementations provided, etc. For the sake of simplicity, they are divided according to their protocol for policy exchange between parties, depending on the sensitivity of policies. On the one hand, assuming that all policies are public and accessible, the process of evaluating whether two policies from two different entities are compatible or not consists in gathering the relevant policies (and possibly relevant credentials) from the involved entities and checking whether they match. On the other hand, if policies may be private, it implies that not all policies are known in advance but they may be disclosed at a later stage. Therefore, a negotiation protocol in which security and trust is iteratively established is required [10].

However, specifying policies is as difficult as writing imperative code, getting a policy right is as hard as getting a piece of software correct, and maintaining a large number of them is even harder. Fortunately, ontologies and policy reasoning may help users and administrators on specification, conflict detection and resolution of such policies [11]. This paper first describes how policies are exchanged and how they interact among parties on the Semantic Web, with a brief description of the main Semantic Web policy languages and how ontologies may be used in policy specification, conflict detection and validation. Then, some examples of application scenarios are presented, where policy based security and privacy are used, followed by some important open research issues.

## 2. POLICY BASED INTERACTION AND EVALUATION

Policies allow for security and privacy descriptions in a machine understandable way. More specifically, service or information providers may

use security policies to control access to resources by describing the conditions a requester must fulfil. At the same time, service or information consumers may regulate the information they are willing to disclose by protecting it with privacy policies. Given two sets of policies, an engine may check whether they are compatible, that is, whether they match. The complexity of this process varies depending on the sensitivity of policies and the expressivity of the policies. If all policies are public at both sides, provider and requester, the requester may initially already provide the relevant policies together with the request and the evaluation process can be performed in a one-step evaluation by the provider policy engine and return a final decision. Otherwise, if policies may be private, as it is, for example, typically the case for sensitive business rules, this process may consist of several steps negotiation in which new policies and credentials are disclosed at each step, therefore advancing after each iteration towards a common agreement. In this section we give an overview of both types of languages. The main features of these languages are shown in Table 1. Additionally, we use the running policy “only employees of institution XYZ may retrieve a file” to illustrate an example of each language.

### 2.1 One-Step Policy Evaluation

Assuming that policies are publicly disclosable, there is no reason why a requester should not disclose its relevant applicable policies together with its request. This way, the provider’s policy engine has all the information needed to make an authorisation decision. The KAOS and REI frameworks, specially designed using Semantic Web features and constructs, fall within this category of policy languages, those which do not allow policies themselves to be protected.

Table 1. Comparison of KAOS, REI, PeerTrust and Protune<sup>1</sup>

Policy Languages	Authorization Protocol	Reasoning Paradigm	Conflict Detection	Meta-policies	Loop Detection
KAOS	One-step	DL	Static Detection & Resolution		
REI	One-step	DL+Variables	Dynamic Detection & Resolution	Used for conflict resolution	
Peer Trust	Negotiation	LP + Ontologies			Distributed Tabling
Protune	Negotiation	LP + Ontologies		Used for Driving Decisions	

<sup>1</sup> DL refers to Description Logic while LP stands for Logic Programming





becomes more complex, since not all relevant policies may be available at the time.

Therefore, this process may consist of a several steps negotiation, by disclosing new policies and credentials at each step, and therefore advancing after each iteration towards a common agreement [10]. For example, suppose Alice requests access to a resource at e-shop. Alice is told that she must provide her credit card to be granted access. However, Alice does not want to disclose her credit card just to anyone and she communicates to e-shop that before it gets her credit card, it should provide its Better Business Bureau certification. Once e-shop discloses it, Alice's policy is fulfilled and she provides the credit card, thus fulfilling e-shop's policy and receiving access to the requested resource (see figure 2).




Fig. 2. Policy-driven negotiation between Alice and e-shop

Below, the two most recent languages for policy-driven negotiation are presented. They are also specially designed for the Semantic Web. However, we refer the interested reader to other languages for policy based negotiations [16][17][18], which may be applied to the Semantic Web.

### PeerTrust

PeerTrust [5] builds upon previous work on policy-based access control and release for the Web and implements automated trust negotiation for such a dynamic environment. PeerTrust's language is based on first order Horn rules (definite Horn clauses), i.e., rules of the form " $lit_0 \leftarrow lit_1, \dots, lit_n$ " where each  $lit_i$  is a positive literal  $P_j(t_1, \dots, t_n)$ ,  $P_j$  is a predicate symbol, and the  $t_i$  are the arguments of this predicate. Each  $t_i$  is a term, i.e., a function symbol and its arguments, which are themselves terms. The head of a rule is  $lit_0$ , and its body is the set of  $lit_i$ . The body of a rule can be empty. Definite Horn clauses can be

easily extended to include negation as failure, restricted versions of classical negation, and additional constraint handling capabilities such as those used in constraint logic programming. Although all of these features can be useful in trust negotiation, here are only described other more unusual required language extensions. Additionally, PeerTrust allows the import of RDF based meta-data therefore allowing the use of ontologies within policy descriptions.

<pre>retrieveFile(fileXYZ) \$Requester ←   employed(Requester) @ institutionXYZ.</pre>	<pre>access('fileXYZ') ←   credential(employee, C),   C.type:employee_id,   C.affiliation:'XYZ'.</pre>
	<pre>access(.) type:decision. access(.) sensitivity:public.</pre>

Fig. 3. Example of PeerTrust (left) and Protune (right) policies

*References to Other Peers* PeerTrust's ability to reason about statements made by other peers is central to trust negotiation. To express delegation of evaluation to another peer, each literal  $lit_i$  is extended with an additional *Authority* argument, that is  $lit_i @ Authority$  where *Authority* specifies the peer who is responsible for evaluating  $lit_i$  or has the authority to evaluate  $lit_i$ . The *Authority* argument can be a nested term containing a sequence of authorities, which are then evaluated starting at the outermost layer. A specific peer may need a way of referring to the peer who asked a particular query. This is accomplished by including a *Requester* argument in literals, so that now literals are of the form  $lit_i @ Issuer \$ Requester$ . The *Requester* argument can also be nested, in which case it expresses a chain of requesters, with the most recent requester in the outermost layer of the nested term. Using the *Issuer* and *Requester* arguments, it is possible to delegate evaluation of literals to other parties and also express interactions and the corresponding negotiation process between parties (see figure 3 for an example). *Signed Rules* Each peer defines a policy for each of its resources, in the form of a set of definite Horn clause rules. These and any other rules that the peer defines on its own are its *local* rules. A peer may also have copies of rules defined by other peers, and it may use these rules to generate proofs, which can be sent to other entities in order to give evidence of the result of a negotiation. A signed rule has an additional argument that says who signed the rule. The cryptographic signature itself is not included in the policy, because signatures are very large and are not needed by this part of the negotiation software. The signature is used to verify that the issuer really





did issue the rule. It is assumed that when a peer receives a signed rule from another peer, the signature is verified before the rule is passed to the DLP evaluation engine. Similarly, when one peer sends a signed rule to another peer, the actual signed rule must be sent, and not just the logic programmatic representation of the signed rule. More complex signed rules often represent delegations of authority.

*Loop detection mechanisms* In declarative policy specification, loops may easily occur and should not be considered as errors. For example, declarative policies may state at the same time that “anyone with write permissions can read a file” and “anyone with read permissions can write a file”. If not handled accordingly, such loops may end up in non-terminating evaluation [25]. In practice, policies, including for instance business rules, are complex and large in number (and typically not under control of a single person) which increases the risk of loops and non-termination during dynamic policy evaluation. A distributed tabling algorithm can handle safely mutual recursive dependencies (loops) in distributed environments. Due to the security context, other aspects like private and public policies and proof generation must be taken into account [25].

### Protune

The P<sub>R</sub>ovisional TrUst NEgotiation framework Protune [7] aims at combining distributed trust management policies with provisional-style business rules and access control related actions. Protune’s rule language extends two previous languages: PAPL [16], which until 2002 was one of the most complete policy languages for trust negotiation, and PeerTrust [5], which supports distributed credentials and a more flexible policy protection mechanism. In addition, the framework features a powerful declarative meta-language for driving some critical negotiation decisions, and integrity constraints for monitoring negotiations and credential disclosure. Protune provides a framework with:

- A trust management language supporting general provisional-style<sup>2</sup> actions (possibly user-defined).
- An extendible declarative meta-language for driving decisions about request

formulation, information disclosure, and distributed credential collection.

- A parameterised negotiation procedure, that gives semantics to the metalanguage and provably satisfies some desirable properties for all possible metapolicies.
- Integrity constraints for negotiation monitoring and disclosure control.
- General, ontology-based techniques for importing and exporting meta-policies and for smoothly integrating language extensions.

The Protune rule language is based on normal logic program rules “ $A \leftarrow L_1, \dots, L_n$ ” where  $A$  is a standard logical atom (called the *head* of the rule) and  $L_1, \dots, L_n$  (the *body* of the rule) are literals, that is,  $L_i$  equals either  $B_i$  or  $\neg B_i$ , for some logical atom  $B_i$ . A *policy* is a set of rules (see figure 3 for an example), such that negation is applied neither to *provisional predicates* (defined below), nor to any predicate occurring in a rule head. This restriction ensures that policies are *monotonic* on credentials and actions, that is, as more credentials are released and more actions executed, the set of permissions does not decrease.

The vocabulary of predicates occurring in the rules is partitioned into the following categories: *Decision Predicates* (currently supporting “allow()” which is queried by the negotiation for access control decisions and “sign()” which is used to issue statements signed by the principal owning the policy, *Abbreviation Predicates* (as described in [16]), *Constraint Predicates* (which comprise the usual equality and disequality predicates) and *State Predicates* (which perform decisions according to the state). State Predicates are further subdivided in *State Query Predicates* (which read the state without modifying it) and *Provisional Predicates* (which may be made true by means of associated actions that may modify the current state like e.g. *credential()*, *declaration()*, *logged(X, logfile name)*). Furthermore, meta-policies consist of rules similar to object-level rules. They allow to inspect terms, check groundness, call an object-level goal  $G$  against the current state (using a predicate *holds(G)*), etc. In addition, a set of reserved attributes associated to predicates, literals and rules (e.g., whether a policy is public or sensitive) is used to drive the negotiator’s decisions. For example, if  $p$  is a predicate, then  $p.sensitivity : private$  means that the extension of the predicate is private and should not be disclosed. An assertion  $p.type : provisional$  declares  $p$  to be a provisional predicate; then  $p$  can be attached to the corresponding action by

<sup>2</sup> Authorizations involving actions and side effects are sometimes called provisional.



asserting  $p.action$  : . If the action is to be executed locally, then we assert  $p.actor$  : self, otherwise assert  $p.actor$  : peer.

### 2.3 Policy specification, conflict detection and resolution

Previous sections described how the Semantic Web may benefit from the protection of resources with policies specifying security and privacy constraints. However, specifying policies may be as difficult as writing imperative code, getting a policy right is as hard as getting a piece of software correct, and maintaining a large number of them is only harder. Fortunately, the Semantic Web can help administrators with policy specification, and detection and resolution of conflicts.

*Policy specification* tools like the KAOS Policy Administration Tool (K-PAT) [3] and the PeerTrust Policy Editor provide an easy to use application to help policy writers. This is important because the policies will be enforced automatically and therefore errors in their specification or implementation will allow outsiders to gain inappropriate access to resources, possibly inflicting huge and costly damages. In general, the use of ontologies on policy specification reduces the burden on administrators, helps them with their maintenance and decreases the number of errors. For example, ontology-based structuring and abstraction help maintain complex software, and so do they with complex sets of policies. In the context of the Semantic Web, ontologies provide a formal specification of concepts and their interrelationships, and play an essential role in complex web service environments, semantics-based search engines and digital libraries. Nejdl et al. [11] suggest using two strategies to compose and override policies, building upon the notions of mandatory and default policies, and formalising the constraints corresponding to these kinds of policies using F-Logic. A prototype implementation as a Protégé plug-in shows that the proposed policy specification mechanism is implementable and effective.

*Conflict detection and resolution.* Semantic Web policy languages also allow for advanced algorithms for conflict detection and its resolution. For example, in section 2.1 it was briefly described how conflicts may arise between policies, either at specification time or runtime. A typical example of a conflict is when several policies apply to a

request and one allows access while another denies it (positive vs. negative authorisation). Description Logic based languages may use subsumption reasoning to detect conflicts by checking if two policies are instances of conflicting types and whether the action classes, that the policies control, are not disjoint. Both KAOS and REI handle such conflicts (like right vs. prohibition or obligation vs. dispensation) within their frameworks and both provide constructs for specifying priorities between policies, hence the most important ones override the less important ones. In addition, REI provides a construct for specifying a general modality priority: positive (rights override prohibitions and obligations override dispensations) or negative (prohibitions override rights and dispensations override obligations). KAOS also provides a conflict resolution technique called ‘policy harmonisation’. If a conflict is detected the policy with lower priority is modified by refining it with the minimum degree necessary to remove the conflict. This process may generate zero, one or several policies as a refinement of the previous one (see [3] for more information). This process is performed statically at policy specification time ensuring that no conflicts arise at runtime.

## 3. APPLYING POLICIES ON THE SEMANTIC WEB

The benefits of using semantic policy languages in distributed environments with automated machine-machine interaction have been described extensible in previous sections. This section aims at providing some examples of its use in the context of the Web, (Semantic) Web Services and the (Semantic) Grid. In all cases, different solutions have been described addressing different scenarios from the point of view of one-step authorization or policy-driven negotiations.

### 3.1 Policies on the Web

The current Web infrastructure does not allow the enforcement of user policies while accessing web resources. Web server authentication is typically based on authentication mechanisms in which users must authenticate themselves (either by means of certificates or typing a user name and password). Semantic Web policies overcome such limitations of the Web. Kagal et al. [4] describe how the REI language can be applied in order to control access to web resources. Web pages are marked up with policies specifying which credentials are required to access such pages. A



policy engine (bound to the web server) decides whether the request matches the credentials requested. In case it does not, the web server could show which credentials are missing. Furthermore, Kolari et al. [20] presents an extension to the Platform for Privacy Preferences (P3P) using the REI language. The authors propose enhancements using REI policies to increase the expressiveness and to allow for existing privacy enforcement mechanisms. PeerTrust can be used to provide advanced policy-driven negotiations on the Web in order to control access to resources [5][21]. A user receives a signed (by a trusted authority) applet after requesting access to a resource. Such an applet includes reasoning capabilities and is loaded in the Web browser. The applet automatically imports the policies specified by the user and starts a negotiation. If the negotiation succeeds, the applet simply retrieve the resource requested or, if necessary, redirects the user to the appropriate repository.

### 3.2 Semantic Web Services

Semantic Web Services aim at the automation of discovery, selection and composition of Web Services. Denker et al. [22] and Kagal et al. [9] suggest extending OWL-S with security policies, written in REI, like e.g. whether a service requires or is capable of providing secure communication channels. An agent may then submit a request to the registry together with its privacy policies. The matchmaker at the registry will filter out non-compatible service descriptions and select only those whose security requirements of the service match the privacy policies of the requester. Differently, Olmedilla et al. [23] propose the use of the PeerTrust language to decide if trust can be established between a requester and a service provider during runtime selection of web services. Modelling elements are added to the Web Service Modeling Ontology (WSMO) in order to include security information in the description of Semantic Web Services. In addition, the authors discuss different registry architectures and their implications for the matchmaking process.

### 3.3 Semantic Grid

Grid environments provide the middleware needed for access distributed computing and data resources. Distinctly administrated domains form virtual organisations and share resources for data retrieval, job execution, monitoring, and data storage. Such an environment provides users with seamless access to all resources they are authorized to access. In current Grid

infrastructures, in order to be granted access at each domain, user's jobs have to secure and provide appropriate digital credentials for authentication and authorisation. However, while authentication along with single sign-on can be provided based on client delegation of X.509 proxy certificates to the job being submitted, the authorisation mechanisms are still mainly identity-based. Due to the large number of potential users and different certification authorities, this leads to scalability problems calling for a complementary solution to the access control mechanisms specified in the current Grid Security Infrastructure (GSI) [24].

Uszok et al. [25] presents an integration of the KAOS framework into Globus Toolkit 3. Its authors suggest offering a KAOS grid service and providing an interface so grid clients and services may register and check whether a specific action is authorised or not. The KAOS grid service uses the KAOS policy services described in section 2.1 and relies on the Globus local enforcement mechanisms. Alternatively, Constandache et al. [26] describe an integration of policy driven negotiations for the GSI, using semantic policies and enhancing it providing automatic credential fetching and disclosure. Policy-based dynamic negotiations allow more flexible authorisation in complex Grid environments, and relieve both users and administrators from up front negotiations and registrations. Constandache et al. [26] introduces an extension to the GSI and Globus Toolkit 4.0 in which policybased negotiation mechanisms offer the basis for overcoming these limitations. This extension includes property-based authorisation mechanisms, automatic gathering of required certificates, bidirectional and iterative trust negotiation and policy based authorisation, ingredients that provide advanced self-explanatory access control to grid resources.

## 4. OPEN RESEARCH ISSUES

Although there has been extensive research in the last years, there exist still open issues that must be solved [27]. The following provides a non-exhaustive list of issues which have not yet been given enough attention, or that still remain unsolved and crucial challenges in order to have a semantic policy framework adopted in real world applications.

- Adoption of a *broad notion of policy*, encompassing not only access control policies, but also privacy policies, business rules, quality of service, agent



- conversation, mobility policies, etc. All these different kinds of policies should eventually be integrated into a single framework.
- *Strong and lightweight evidence*: Policies make decisions based on properties of the peers interacting with the system. These properties may be strongly certified by cryptographic techniques, or may be reliable to some intermediate degree with lightweight evidence gathering and validation. A flexible policy framework should try to merge these two forms of evidence to meet the efficiency and usability requirements of web applications.
  - These desiderata imply that trust negotiation, reputation models, business rules, and action specification languages have to be integrated into a single framework at least to some extent. It is crucial to find the right tradeoff between generality and efficiency.
  - *Automated policy driven negotiation* is one of the main ingredients that can be used to make heterogeneous peers effectively interoperate.
  - *Lightweight knowledge representation and reasoning* does not only refer to computational complexity; it should also reduce the effort to specialise general frameworks to specific application domains; and the corresponding tools should be easy to learn and use for common users, with no particular training in computers or logic.
  - The last issue cannot be tackled simply by adopting a rule language. Solutions like *controlled natural language syntax for policy rules*, to be translated by a parser into the internal logical format, will definitively ease the adoption of any policy language.
  - *Cooperative policy enforcement*: A secure cooperative system should (almost) never say *no*. Web applications need to help new users in obtaining the services that the application provides, so potential customers should not be discouraged. Whenever prerequisites for accessing a service are not met, web applications should explain what is missing and help the user in obtaining the required permissions. As part of cooperative enforcement, advanced *explanation*

*mechanisms* are necessary to help users in understanding policy decisions and obtaining the

- permission to access a desired service.

## 5. CONCLUSIONS

This paper provides an introduction to policy-based security and privacy management on the Semantic Web. It describes the benefits of using policies and presents four of the most relevant policy languages in the Semantic Web contexts. These four languages are classified according to whether policies are assumed to be public or else may be protected. The former consists of a single evaluation step where a policy engine or a matchmaker decides whether two policies are compatible or not. Examples of this kind of evaluation are the KAOS and REI frameworks. If policies may be protected (by e.g. other policies), the process is not anymore a one-step evaluation. In this case, policies guide a negotiation in which policies are disclosed iteratively increasing the level of security at each step towards a final agreement. Examples of these kind of frameworks are PeerTrust and Protune. Furthermore, Semantic Web techniques can be used to ease and enhance the process of policy specification and validation. Conflicts between policies can be found and even resolved automatically.

In order to demonstrate the benefits and feasibility of Semantic Web policies, several application scenarios are described, namely the Web, (Semantic) Web Services and the (Semantic) Grid. Finally the chapter concludes with a list of open research issues that prevent existing policy languages from being widely adopted. This list is intended to help new researchers in the area to focus on those crucial problems which are still unsolved.

## REFERENCES:

- [1]. Tim Berners-Lee, James Hendler, and Ora Lassila. The Semantic Web. *Scientific American*, may 2001.
- [2]. Grigoris Antoniou, Matteo Baldoni, Piero A. Bonatti, Wolfgang Nejdl, and Daniel Olmedilla. Rule-based policy specification. In Ting Yu and Sushil Jajodia, editors, *Decentralized Data Management Security*. Springer, 2006.





- [3]. Andrzej Uszok, Jeffrey M. Bradshaw, Renia Jeffers, Niranjani Suri, Patrick J. Hayes, Maggie R. Breedy, Larry Bunch, Matt Johnson, Shriniwas Kulkarni, and James Lott. KAoS policy and domain services: Toward a description-logic approach to policy representation, deconfliction, and enforcement. In *POLICY*, page 93, 2003.
- [4]. Lalana Kagal, Timothy W. Finin, and Anupam Joshi. A policy based approach to security for the semantic web. In *The Semantic Web - ISWC 2003, Second International Semantic Web Conference, Sanibel Island, FL, USA, October 20-23, 2003, Proceedings*, Lecture Notes in Computer Science, pages 402–418. Springer, 2003.
- [5]. Rita Gavriloiu, 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 *Lecture Notes in Computer Science*, pages 342–356, Heraklion, Crete, Greece, may 2004. Springer.
- [6]. Moritz Y. Becker and Peter Sewell. Cassandra: Distributed access control policies with tunable expressiveness. In *5th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2004)*, 7-9 June 2004, Yorktown Heights, NY, USA, pages 159–168. IEEE Computer Society, 2004.
- [7]. Piero A. Bonatti and Daniel Olmedilla. Driving and monitoring provisional trust negotiation with metapolicies. In *6th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2005)*, pages 14–23, Stockholm, Sweden, June 2005. IEEE Computer Society.
- [8]. Gianluca Tonti, Jeffrey M. Bradshaw, Renia Jeffers, Rebecca Montanari, Niranjani Suri, and Andrzej Uszok. Semantic web languages for policy representation and reasoning: A comparison of KAoS, Rei, and Ponder. In *International Semantic Web Conference*, pages 419–437, 2003.
- [9]. Lalana Kagal, Massimo Paolucci, Naveen Srinivasan, Grit Denker, Timothy W. Finin, and Katia P. Sycara. Authorization and privacy for semantic web services. *IEEE Intelligent Systems*, 19(4):50–56, 2004.
- [10]. William H. Winsborough, Kent E. Seamons, and Vicki E. Jones. Automated trust negotiation. DARPA Information Survivability Conference and Exposition, IEEE Press, Jan 2000.
- [11]. Wolfgang Nejdl, Daniel Olmedilla, Marianne Winslett, and Charles C. Zhang. Ontology-based policy specification and management. In *2nd European Semantic Web Conference (ESWC)*, volume 3532 of *Lecture Notes in Computer Science*, pages 290–302, Heraklion, Crete, Greece, may 2005. Springer.
- [12]. Jeffrey M. Bradshaw, Andrzej Uszok, Renia Jeffers, Niranjani Suri, Patrick J. Hayes, Mark H. Burstein, Alessandro Acquisti, Brett Benyo, Maggie R. Breedy, Marco M. Carvalho, David J. Diller, Matt Johnson, Shriniwas Kulkarni, James Lott, Maarten Sierhuis, and Ron van Hoof. Representation and reasoning for DAML-based policy and domain services in KAoS and nomads. In *The Second International Joint Conference on Autonomous Agents & Multiagent Systems (AAMAS)*, Melbourne, Victoria, Australia, jul 2003.
- [13]. Mike Dean and Guus Schreiber. OWL web ontology language reference, 2004.
- [14]. Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
- [15]. Lalana Kagal. *A Policy-Based Approach to Governing Autonomous Behaviour in Distributed Environments*. PhD thesis, University of Maryland Baltimore County, 2004.
- [16]. P. Bonatti and P. Samarati. Regulating Service Access and Information Release on the Web. In *Conference on Computer and Communications Security (CCS'00)*, Athens, November 2000.
- [17]. N. Li and J.C. Mitchell. RT: A Role-based Trust-management Framework. In *DARPA*



- Information Survivability Conference and Exposition (DISCEX)*, Washington, D.C., April 2003.
- [18]. Jim Trevor and Dan Suciu. Dynamically distributed query evaluation. In *Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, Santa Barbara, CA, USA, May 2001.
- [19]. 25. Miguel Alves, Carlos Viegas Damásio, Daniel Olmedilla, and Wolfgang Nejdl. A distributed tabling algorithm for rule based policy systems. In *7th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2006)*, London, Ontario, Canada, June 2006. IEEE Computer Society.
- [20]. Pranam Kolari, Li Ding, Shashidhara Ganjugunte, Anupam Joshi, Timothy W. Finin, and Lalana Kagal. Enhancing web privacy protection through declarative policies. In *6th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2005)*, pages 57–66, Stockholm, Sweden, jun 2005. IEEE Computer Society.
- [21]. Steffen Staab, Bharat K. Bhargava, Leszek Lilien, Arnon Rosenthal, Marianne Winslett, Morris Sloman, Tharam S. Dillon, Elizabeth Chang, Farookh Khadeer Hussain, Wolfgang Nejdl, Daniel Olmedilla, and Vipul Kashyap. The pudding of trust. *IEEE Intelligent Systems*, 19(5):74–88, 2004.
- [22]. Grit Denker, Lalana Kagal, Timothy W. Finin, Massimo Paolucci, and Katia P. Sycara. Security for daml web services: Annotation and matchmaking. In *The Semantic Web - ISWC 2003, Second International Semantic Web Conference, Sanibel Island, FL, USA, October 20-23, 2003, Proceedings*, Lecture Notes in Computer Science, pages 335–350. Springer, 2003.
- [23]. Daniel Olmedilla, Rubén Lara, Axel Polleres, and Holger Lausen. Trust negotiation for semantic web services. In *1st International Workshop on Semantic Web Services and Web Process Composition (SWSWPC)*, volume 3387 of *Lecture Notes in Computer Science*, pages 81–95, San Diego, CA, USA, jul 2004. Springer.
- [24]. Grid Security Infrastructure. <http://www.globus.org/security/overview.html>.
- [25]. Andrzej Uszok, Jeffrey M. Bradshaw, and Renia Jeffers. Kaos: A policy and domain services framework for grid computing and semantic web services. In *Trust Management, Second International Conference, iTrust 2004, Oxford, UK, March 29 - April 1, 2004, Proceedings*, Lecture Notes in Computer Science, pages 16–26. Springer, 2004.
- [26]. Ionut Constandache, Daniel Olmedilla, and Wolfgang Nejdl. Policy based dynamic negotiation for grid services authorization. In *Semantic Web Policy Workshop in conjunction with 4th International Semantic Web Conference*, Galway, Ireland, November 2005.
- [27]. P.A. Bonatti, C. Duma, N. Fuchs, W. Nejdl, D. Olmedilla, J. Peer, and N. Shahmehri. Semantic Web policies - A discussion of requirements and research issues. In *3rd European Semantic Web Conference (ESWC)*, Lecture Notes in Computer Science, Budva, Montenegro, jun 2006.