



SHI3LD: an Access Control Framework for the Mobile Web of Data

Luca Costabello, Serena Villata, Nicolas Delaforge, Fabien Gandon

► To cite this version:

Luca Costabello, Serena Villata, Nicolas Delaforge, Fabien Gandon. SHI3LD: an Access Control Framework for the Mobile Web of Data. Hypertext - 23rd ACM Conference on Hypertext and Social Media - 2012, Jun 2012, Milwaukee, United States. hal-00691277

HAL Id: hal-00691277

<https://inria.hal.science/hal-00691277>

Submitted on 25 Apr 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SHI3LD: an Access Control Framework for the Mobile Web of Data

Luca Costabello, Serena Villata*, Nicolas Delaforge, Fabien Gandon
INRIA Sophia Antipolis
France
firstname.lastname@inria.fr

ABSTRACT

We present Shi3ld, a context-aware access control framework for consuming the Web of Data from mobile devices.

Categories and Subject Descriptors

H.3.5 [Information Storage and Retrieval]: Online Information Services

1. INTRODUCTION

The Web is evolving from an information space for sharing textual documents into a medium for publishing structured data. Recent developments in the Semantic Web field leverage on the RDF uniform data model and on URIs to merge and identify structured data of heterogeneous nature. The Linked Data¹ initiative aims at fostering the publication and interlink of data on the Web, giving birth to the *Web of Data*, an interconnected global dataspace where data providers publish their content publicly [6].

In this paper we describe Shi3ld², an access control framework for querying RDF datastores in mobile environments. The open nature of current Web of Data information and the consumption of web resources from mobile devices may give providers the impression that their content is not safe, thus preventing further publication of datasets, at the expense of the growth of the Web of Data itself. Access control is therefore necessary, and mobile context must be part of the access control evaluation. For a comparison with the related work [1, 5, 7, 8], see [3].

We protect RDF stores by changing the semantics of incoming SPARQL queries, whose scope is restricted to triples included in accessible Named Graphs only [2]. We determine the list of accessible graphs by evaluating pre-defined access policies against the actual mobile context of the requester. Beyond the support for context in control enforcement, our proposal has the advantage of being a pluggable filter for generic SPARQL endpoints, with no need to modify the endpoint itself. We adopt exclusively Semantic Web languages and reuse existing proposals, thus we do not add new policy definition languages, parsers nor validation procedures. We

* Acknowledge support of the DataLift Project ANR-10-CORD-09 founded by the French National Research Agency.

¹<http://linkeddata.org>

²<http://wimmics.inria.fr/projects/shi3ld>

provide protection up to triple level. Our work does not provide yet another context ontology: our model includes base classes and properties only, as we delegate refinements and extensions to domain specialists, in the light of the Web of Data philosophy. For the time being, our framework assumes the trustworthiness of the information sent by the mobile consumer, including data describing context (e.g. location, device features, etc). We do not provide any privacy-preserving mechanism yet, although we are aware that sensible data such as current location must be handled appropriately.

2. THE FRAMEWORK

The access control model is built over the notion of Named Graph [2], thus supporting fine-grained access control policies, including the triple level. We rely on named graphs to avoid depending on documents (one document can serialize several named graphs, one named graph can be split over several documents, and not all graphs come from documents). The model is grounded on two ontologies: **S4AC** deals with core access control concepts and **PRISSMA** focuses on the mobile context. The main component of the **S4AC** model is the Access Policy which defines the constraints that must be satisfied to access a given named graph or a set of named graphs. If the Access Policy is *satisfied* the data consumer is allowed to access the data. Otherwise, access is denied. The constraints specified by the Access Policies concern the data consumer, the device, the environment, or any given combination of these dimensions. We express Access Conditions as SPARQL ASK queries. Each Access Policy is associated to an Access Evaluation Context, an explicit link between the policy and the actual context data used to evaluate the Access Policy. The Shi3ld framework adopts **PRISSMA** which provides classes and properties to model core mobile context concepts, but is not meant to deliver yet another mobile contextual model: instead, well-known Web of Data vocabularies and recent W3C recommendations are reused. We agree on the widely-accepted proposal by Dey [4] and, more specifically, on the work by Fonseca et al.³. The mobile context is seen as an encompassing term, an information space defined as the sum of three different dimensions: the mobile *User* model, the *Device* features and the *Environment* in which the action is performed.

An example of Access Policy associated to a **Read** privilege is shown in Figure 1a. The policy protects the named graph `:alice_data` and allows the access to the named graph only

³<http://bit.ly/XGR-mbui>

```

:policy1 a s4ac:AccessPolicy; ACCESS POLICY
s4ac:appliesTo :alice_data; RESOURCE TO PROTECT
s4ac:hasAccessPrivilege [a s4ac:Read]; ACCESS PRIVILEGE
s4ac:hasAccessConditionSet :acs1.

:acs1 a s4ac:AccessConditionSet;
s4ac:ConjunctiveAccessConditionSet; ACCESS CONDITIONS
s4ac:hasAccessCondition :acl1, :acl2. TO VERIFY

:acl1 a s4ac:AccessCondition;
s4ac:hasQueryAsk
  ""ASK {?context a prisma:Context.
    ?context prisma:user ?u.
    ?u foaf:knows ex:alice#me.}""

:acl2 a s4ac:AccessCondition;
s4ac:hasQueryAsk
  ""ASK {?context a prisma:Context.
    ?context prisma:environment ?env.
    ?env prisma:based_near ?p.
    FILTER (!(?p=ex:ACME_boss#me))}""

```

(a)

```

:bobCtx{
:ctx1 a prisma:Context;
prisma:user :usr1;
prisma:device :dev1;
prisma:environment :env1.
} THE CONSUMER'S CONTEXT

:usr1 a prisma:User;
foaf:name "Bob";
foaf:knows ex:alice#me.
} THE USER DIMENSION

:dev1 a prisma:Device;
soft:deviceSoftware :dev1sw.
:dev1sw a soft:DeviceSoftware;
soft:operatingSystem :dev1os.
:dev1os a soft:OperatingSystem;
common:name "Android".
} THE DEVICE DIMENSION

:env1 a prisma:Environment;
prisma:motion "no";
prisma:nearbyEntity :ACME_boss#me;
prisma:currentPOI :ACMEoffice.
:ACMEoffice a prisma:POI;
prisma:poiCategory example:Office;
prisma:poiLabel example:ACMECorp.
} THE ENVIRONMENT DIMENSION
}

```

(b)

Figure 1: The Access Policy protecting :alice_data (a) and Bob's sample mobile context in TriG notation (b).

if the consumer (i) knows Alice, and (ii) is not located near Alice's boss. Figure 1b visualizes a sample mobile context featuring all the dimensions described above. The user, Bob, knows Alice and is currently at work, near his and Alice's boss. Bob is using an Android device and is not moving.

Our Access Control Manager is designed as a pluggable component for SPARQL endpoints. As mobile consumer query the SPARQL endpoint to access content, context data is sent with the query and cached as a named graph using SPARQL 1.1 update language statements. Each time a context element is added we use an **INSERT DATA**, while we rely on a **DELETE/INSERT** when the contextual information is already stored and has to be updated. Summarizing, the mobile client sends two SPARQL queries: the first is the client query to the datastore, the second provides contextual information (e.g. Figure 1b). The client query is filtered by the Access Control Manager instead of being directly executed on the SPARQL endpoint. The Access Control Manager selects the set of policies affecting the client query, i.e. those with a matching Access Privilege. The Access Conditions (SPARQL ASK queries) included in the selected policies are executed. For each verified policy, the associated

```

PREFIX bibo: <http://purl.org/ontology/bibo/>
SELECT *
WHERE {?review a bibo:Article}

```

(a)

```

PREFIX bibo: <http://purl.org/ontology/bibo/>
SELECT *
FROM :peter_reviews NAMED GRAPH
ACCESSIBLE BY THE CONSUMER
WHERE {?review a bibo:Article}

```

(b)

Figure 2: Bob's SPARQL query (a) and the secured one (b).

named graph is added to the set of accessible named graphs. The client query is sent to the SPARQL endpoint with the addition of the **FROM** clause(s). Query execution is therefore performed only on the accessible named graphs, given the consumer contextual information. The result of the query is returned to the consumer.

An example of client query is shown in Figure 2a, where Bob wants to access all the datastore (including Alice data) from the context described in Figure 1b. The Access Conditions included in the policies are evaluated against the actual context data of the mobile consumer. In our example, the identification of the named graph(s) accessible by Bob returns only the graph **:peter_data**. Alice data is forbidden because Access Conditions evaluation leads to a **false** answer with Bob's context (Bob is near Alice's boss). The Manager adds the **FROM** clause to constrain the execution of the client query only on the allowed named graph. The "secured" client query is shown in Figure 2b. For the implementation details of Shi3ld and its evaluation, see [3].

3. REFERENCES

- [1] F. Abel, J. L. De Coi, N. Henze, A. W. Koesling, D. Krause, and D. Olmedilla. Enabling Advanced and Context-Dependent Access Control in RDF Stores. In *Procs of ISWC-2007, LNCS 4825*, pages 1–14, 2007.
- [2] J. J. Carroll, C. Bizer, P. J. Hayes, and P. Stickler. Named graphs. *J. Web Sem.*, 3(4):247–267, 2005.
- [3] L. Costabello, S. Villata, N. Delaforge, and F. Gandon. Ubiquitous access control for sparql endpoints: Lessons learned and future challenges. In *Procs of WWW Companion*, 2012.
- [4] A. K. Dey. Understanding and using context. *Personal Ubiquitous Computing*, 5:4–7, 2001.
- [5] G. Flouris, I. Fundulaki, M. Michou, and G. Antoniou. Controlling Access to RDF Graphs. In *Procs of FIS-2010, LNCS 6369*, pages 107–117, 2010.
- [6] T. Heath and C. Bizer. *Linked Data: Evolving the Web into a Global Data Space*. Morgan & Claypool, 2011.
- [7] O. Sacco and A. Passant. A Privacy Preference Ontology (PPO) for Linked Data. In *Procs of LDOW-2011*, 2011.
- [8] A. Toninelli, R. Montanari, L. Kagal, and O. Lassila. A semantic context-aware access control framework for secure collaborations in pervasive computing environments. In *Procs of ISWC-2006, LNCS 4273*, pages 473–486, 2006.