

An Access Control Framework for Business Processes for Web Services*

Hristo Koshutanski Fabio Massacci
Dip. di Informatica e Telecomunicazioni - Univ. di Trento
via Sommarive 14 - 38050 Povo di Trento (ITALY)
{hristo,massacci}@dit.unitn.it

ABSTRACT

Business Processes for Web Services are the new paradigm for the lightweight integration of business from different enterprises.

Whereas the security and access control policies for basic web services and distributed systems are well studied and almost standardized, there is not yet a comprehensive proposal for an access control architecture for business processes. The major issue is that a business process describe complex services that cross organizational boundaries and are provided by entities that see each other as just partners and nothing else.

This calls for a number of differences with traditional aspects of access control architectures such as

- credential vs classical user-based access control,
- interactive and partner-based vs one-server-gathers-all requests of credentials from clients,
- controlled disclosure of information vs all-or-nothing access control decisions,
- abducting missing credentials for fulfilling requests vs deducing entailment of valid requests from credentials in formal models,
- “source-code” authorization processes vs data describing policies for communicating policies or for orchestrating the work of authorization servers.

Looking at the access control field we find good approximation of most components but not their synthesis into one access control architecture for business processes for web services, which is the contribution of this paper.

*This work is partially funded by the IST programme of the EU Commission, FET under the IST-2001-37004 WASP project and by the FIRB programme of MIUR under the RBNE0195K5 ASTRO Project.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM Workshop on XML Security, October 31, 2003, Fairfax VA, USA
Copyright 2003 ACM 1-58113-777-X/03/0010 ...\$5.00.

Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection—*Access controls, Information flow controls*; H.3.5 [Information Storage and Retrieval]: Online Information Services—*Commercial services, Web-based services*; H.4.1 [Information Systems Applications]: Office Automation—*Workflow management*; K.4.4 [Computers and Society]: Electronic Commerce—*Distributed commercial transactions*

General Terms

Management, Design, Security, Languages

Keywords

Web Services, Interactive Access Control, E-Business, Security Management, Distributed Systems Security, Controlled Disclosure.

1. INTRODUCTION

Middleware has been the enterprise integration buzzword at the end of the past millennium. Nowadays a new paradigm is taking hold: Web services (WS for short). Setting hype aside, the major difference between middleware solutions (CORBA, COM+, EJB, etc.) and WS is the idea of lightweight integration of business processes from different enterprises.

The security of basic WS is well studied and standardized [6]. There are also many approaches [35, 37, 4, 16, 13, 5, 33] for controlling access to services and trust management in distributed systems and an advanced standardization process for policies and access control (see for instance the OASIS proposals [12, 26]). However, with the notable exceptions of provisional access control [22] and trust negotiation [36], access control models rest on the idea that the server picks the evidence you sent on who you are (credentials) and what you want (request), checks its evidence on what you deserve (policies) and makes a one-off decision.

Moving up in the WS hierarchy from single services to *orchestration and choreography of WS and business processes* the picture changes. Business processes describe complex services that cross organizational boundaries and are provided by partners.

The paradigmatic example in the WS standards is a travel agent WS that must orchestrate a combination of plane and train tickets, car rental, hotel booking and insurance, each service offered by different partner which may or may not be involved according to the actual unrolling of the workflow.

For example consider the problem of going to a nice “Shakespearean Tour” in Italy: you might decide to go to the city

of Shylock, from there rent a car and travel to Romeo and Juliet's last resort, to jump then on a train and visit the Senate's seat where Pompeus spoke after Caesar's death. However, you might as well decide to travel instead to Germany first and then the train to Verona from there. In the first case you might need to use a car rental company. The second path may require to contact a German train company for the schedule, which is not needed if you land directly in Italy.

Let us now consider the problem of "lightweight" credentials such as the German train discount card or the car rental gold member card. Should the user provide them anyway at the beginning? Obviously not. Should the server orchestrating the process require each partner to publish its policy on discounts? Obviously not.

Such problems are not simply problems of practicality, but have major security implications:

1. Credential vs identity based access control – A WS is something you publish on the Web for everybody to use it, so its design should be fairly close to the principles of trust management systems [5, 33];
2. Orchestrating vs combining – *partners* have different security policies and are just partners and not part of the same enterprise. They may not wish to disclose their policies to the server orchestrating the request. So, we cannot simply combine the policies ('cause we do not have them), we need to orchestrate the request grant/deny/process of many different policies/partners.
3. Interactive vs one-off access control – if partners have different policies they might as well require different credentials to a client. Privacy considerations make gathering all potentially needed credentials from clients difficult. Furthermore, this may simply be impossible. An airline may want to ask confidential information directly to its frequent fliers (e.g., confirmation of religious preferences for food) and not to the Web travel agent orchestrator of the process. This calls for an interactive process in which the client may be asked on the fly for additional credentials and may grant or deny such requests¹.
4. Abducing vs deducing credentials – in most classical formal models of access control we deduce that a request is valid because it is entailed by the combination of the policy and the set of available credentials. Here, a partner must be able to infer the causes of some failed request to ask the missing credentials to the client. The corresponding logical process is no longer deduction but it is abduction. So we must have co-existence of deduction (for deciding access and release of information) and abduction (for explaining failed accesses).
5. Data vs source level communication – the choice of format for messages is always rather complicated, as it calls for the implementation of software that is able

¹Note that the workflow may even take completely different paths based on the results of interaction. For example a rent-a-car operator may require a signed credit card number plus a physical address. The client may deny such requirement and thus another operator may be chosen that only asks for a credit card number.

Web Service Technology Stack		Access Control Issues
Layer	Standards	AC Granularity
Workflow	BPEL4WS	Workflow-level AC
Discovery	UDDI	Description-level AC
Service Description	WSDL	Service-level (End Point) AC
Messaging	SOAP/XML Protocol	Universal way to convey AC info
Transport Protocols	HTTP,HTTPS,FTP,SMTP	–

Figure 1: Web Services Technology Stack & Access Control Issues

to interpret its meaning. In a Business Process scenario we no longer need messages, but just "mobile" processes. A client will receive a business process so that he can simply execute the source to obtain and send the missing credential. An authorization server can download a business process from a policy orchestrator and obtain the desired authorization.

Looking at the access control field we find a good approximation of most components: we have proposals for combining policies at the logical level [23, 34, 3] and at the architectural level [26]. We have proposals for calculi for controlling release of information [7], procedures for trust negotiations and communication of credentials [36] and architecture for distributed access control [12, 4, 35, 16, 24].

What is missing is a way to synthesize *all* these aspects into one access control framework for business processes of WS, which is the contribution of this paper.

1.1 Plan of the Paper

In the next section we introduce some notion about WS and Business Processes for WS. Then we present our architecture and discuss how the entire message passing scheme can be implemented as "mobile" processes in XML. Section 6 explains how we can use logical deduction and logical abduction to build a firm foundation for the interactive process of inferring disclosable credentials from access control policies and from release policies. A brief discussion of related works concludes the paper.

2. A PRIMER ON WS AND BUSINESS PROCESSES

A Web Service as defined by the standard [21] is "an interface that describes a collection of operations that are network-accessible through standardized XML messaging. A Web service is described using a standard, formal XML notion, called its *service description*. It covers all the details necessary to interact with the service, including message formats (that detail the operations), transport protocols and location."

The idea behind Web services is to encapsulate and make available enterprise resources in a new heterogeneous and distributed way.




Figure 2: Example of BPEL4WS Process

```

<process>
  <sequence>
    <receive partner="Customer"
      portType="purchaseOrderPT"
      operation="SendPurchaseOrder"
      container="PO">
    </receive>
    <invoke partner="CreditBureau"
      portType="CheckCreditPT"
      operation="CheckCredit">
    </invoke>
    <invoke partner="shippingProvider"
      portType="shippingPT"
      operation="RequestShipping"
      inputContainer="shippingRequest"
      outputContainer="shippingInfo">
      <source linkName="ship-to-invoice">
    </invoke>
    <reply partner="Customer"
      portType="purchaseOrderPT"
      operation="SendPurchaseOrder"
      container="Invoice"/>
  </sequence>
</process>

```




Figure 3: Cross-section view of the architecture

The WS architecture, as defined by W3C [32], is divided into five layers grouped into three main components - Wire, Description and Discovery (Figure 1). The *Wire* component comprises the messaging and transport layers with the SOAP protocol and the XML message format. *Discovery* offers users a unified and systematic way to find, discover and inspect service providers over the Internet. There are two standards proposed at this level - Universal Description, Discovery and Integration (UDDI) and Web Service Inspection Language (WSIL).

Moving upward we found the *Service Description* layer and the *Business Process Orchestration* layer. The service description layer is responsible for describing the basic format of offered services (protocols and encodings, where a service resides and how to invoke it). The standard for describing the communication details at this layer is Web Service Description Language (WSDL).

The Business Process Orchestration layer is an extension of the service model defined at the description layer. This layer is responsible for describing the behavior of complex business and workflow processes. Intuitively, business processes are graphs where each node represents a business activity and primitive nodes are in WSDL. The recently released standard at this layer is the Business Process Execution Language for WS (BPEL4WS) [10].

The BPEL4WS primitive activities are the following:

- <invoke> invoking an operation on some Web service;
- <receive> waiting for an operation to be invoked by someone externally;
- <reply> generating the response of an input/output operation;
- <assign> copying data from one place to another.

More complex activities can be constructed by composition:

- <sequence> - allows the developer to define an ordered sequence of steps;
- <switch> - allows the developer to have branching;
- <while> - allows the developer to define a loop;
- <flow> - allows the developer to define that a collection of steps has to be executed in parallel.

An example of compositions of services is shown in Figure 2: a buyer service is ordering goods from a seller service, i.e. the buyer service invokes the order method on the seller

service, whose interface is defined using WSDL. The seller service invokes a credit validation service to ensure that the buyer can pay for the goods and after that continues by shipping the goods to the buyer. The credit validation service can take place at a credit bureau site in a separate security domain. Notice that a number of partners participate in the process that therefore crosses administrative boundaries.

The XML code shown in Figure 2 is a very brief example of the scenario described above in the notations of BPEL4WS primitives. The structure of the processing section is defined by the <sequence> element, which states that the elements contained inside are executed in this order. The node content is self explanatory.

3. ARCHITECTURE

Combining the traditional proposals for distributed access control and the essential components used for Web services we propose here a security architecture for orchestrating authorization of Web Services Processes.

Figure 3 shows view of the architecture. A brief description of the entities is given below.

AttributeServer is responsible for providing group/role membership information as in [35, 37], for instance in the form of membership and non-membership certificates.

RegistryServer is responsible for maintaining relations between services and service providers implementing a particular service. When a Client requests the RegistryServer for a specific service, the latter responds with a list of ApplicationServers implementing the requested service.

AuthorizationServer decouples the authorization logic from the application logic. It is responsible for *locating*, *executing* and *managing* all needed PolicyEvaluators and returning an appropriate result to the ApplicationServer. Also it is responsible for managing all the *interactions* with the Client.

PolicyEvaluator terminology borrowed from Beznosov et al [4], is an entity responsible for achieving endpoint decisions on access control (see Figure 3). All partners involved in a business process are likely to be as different entities, each of them represented by a PolicyEvaluator.

PolicyOrchestrator from the authorization point of view is an entity responsible for the workflow level access and release control. It decides which are the partners that are involved in the requested service (Web service workflow) and on the base of some orchestration security policies to combine the corresponding PolicyEvaluators in a form of a Web process (*Policy Composition*

Process) that is suitable for execution by the `AuthorizationServer`.

Figure 4 shows an horizontal view of the same architecture with multiple servers.

To secure the entire architecture we must make some assumptions on the security properties of the lower levels. At transport level we assume the adoption of the WS-Security specification [6] that describes enhancements to SOAP messaging to provide message integrity, confidentiality and authentication. For the message level one can use the W3C and IETF specification for XML-Signature [25] and W3C XML-Encryption [14], or the recently release specifications by IBM and Microsoft for WS secure conversations [18, 19].

One of the advantages of using BPEL4WS is that it is possible to implement the entire architecture using BPEL4WS. In this framework, each component is a business process that communicate with others via web services². We plan to use the BPEL generator Collaxa for a sophisticated implementation that includes also the actual verification of credentials³.

At this stage one may wonder why do we need a `PolicyOrchestrator` at all. The `AuthorizationServer` could as well contact all `PolicyEvaluators` on its own. Instead we have decided to decouple the problem in two parts: *deciding* the authorization process and *running* it. The `AuthorizationServer` runs the actual authorization process and thus queries all the `PolicyEvaluators` that are needed. The entity burdened with the task of deciding what authorization process should be run is `PolicyOrchestrator`.

We free the `AuthorizationServer` from bothering about all the details around connections between partners and `PolicyEvaluators`, as well as, `PolicyEvaluator`'s description, location, orchestration, etc. The `PolicyOrchestrator` is responsible for the *Policy Composition Service*: maintaining all relations between resources names (services) and policies, selecting which are the partners involved in the requested process and combining the corresponding `PolicyEvaluators` (as mentioned before) in a policy composition process and link them to the workflow level access and release policies. This is possible because the `AuthorizationServer` can just download and run a business process as we'll discuss in the next section.

4. COMMUNICATION AS “MOBILE” PROCESSES

Assuming security at lower level, the second key component is the format of communications. We propose here a major innovation: the typical exchange of messages in access control system is at “data” level (credentials, policies, requests, objects, etc.) that are interpreted by the recipients. This choice makes the actual implementation of proposed access control infrastructure difficult and often not easily portable. Here, we propose to exchange messages at “source code” level and in particular at the level of business process description. We advocate *mobility of authorization of business processes*. It means that instead of sending just

²This creates a recursive problem of access control: if the the `PolicyEvaluator` publishes its services as a web service, who can access these services? For the time being we'll swipe that problem under the carpet and assume that this is done by a good old access list of authorized `AuthorizationServer` using a suitable authentication mechanism at lower levels.

³At present credentials are just textual expressions and public key operations are not performed.

messages that have to be interpreted by entities, we truly have mobile processes passing from one entity to another indicating themselves what the recipient has to do.

We have decided to use the term *mobile process* because it well expresses the idea of using mobile code together with the functionality of Web processes. The main advantages of using mobile processes in our authorization framework are flexibility and simplicity of entities. The recipient of mobile process is not limited to the functions and computational algorithms that the recipient's logic predefines. Migrations of actors in the system from one server to another is easier with mobile processes and the system as a whole is more flexible. Entities in the framework becomes simpler, having little functionality pre-engineered into them.

Thus, a quick implementation of a server only needs an off-the-shelf interpreter for business processes. Leading this approach at an extreme the `AuthorizationServer` can simply receive a business process from the orchestrator and execute it. The process may still be computationally intensive as an `AuthorizationServer` may have to process thousands or millions of authorization workflows, but it could be logically very simple thus reducing the TCB to the simple execution of certified processes from certified sources.

Another reason is that some `PolicyEvaluators` may decide to disclose their XACML policies to the entity coordinating the proposal and others `PolicyEvaluators` may instead decide to offer an external interface, so that they just specify a container for requests and an output container for its decision. All intermediate choices are possible if we allow the `AuthorizationServer` just to follow an arbitrary business process as certified by the `PolicyOrchestrator`.

5. INTERACTION WITH THE CLIENT

The next important step in advocating mobile processes is to specify a language that is needed for coding them. We have identified it as a *language for communicating interactive requests back to a Client*. This is even in the case when a `Client` is an `AuthorizationServer` waiting for a response either from a `PolicyOrchestrator` or from a `PolicyEvaluator`. This language can be designed with a black box view of the `PolicyEvaluator` but must be easily interpretable from the `Client` side. Thus we propose to use BPEL4WS itself as a language in which requests are coded. The `PolicyEvaluator/PolicyOrchestrator` must represent its request as a WS business process that can then be interpreted and executed by the `Client`. If the `PolicyEvaluator` wants part of the request to be only visible to the `Client` it can use the available XML-crypto features [25, 14] to protect the relevant part.

Loosely speaking we may say that the `Client` starts by executing a simple `<invoke>R</invoke>` and obtain in return either its result or a more complicated process to execute. For example a BPEL4WS interactive request may specify a `<input container>` where to put a digitally signed copy of the travel contract sealed with the public key of the rent-a-car company (a process that can be specified as a `<sequence>` of events).

The idea is intuitive and appealing but there is an essential detail that must be taken care of. Notably, the `AuthorizationServer` will receive a number of interactive requests while controlling its workflow and the combination of these requests and the service workflow specification is essential. The simplest solution is to ignore such interaction: all interactive requests are compiled into a `<flow>` and the result

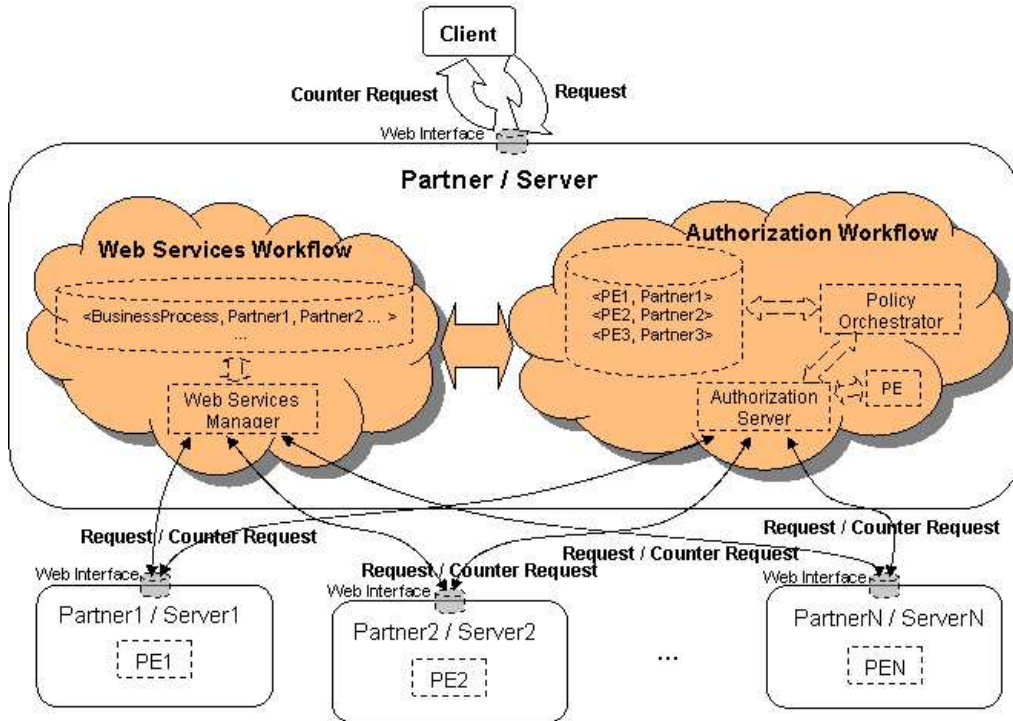


Figure 4: Horizontal view of the architecture

is sent back to the Client. Such solution is hardly satisfactory from the point of view of the Client: we often want to know "why" some additional information is needed. See the example of Figure 2: at some stage somebody may ask for a digitally signed declaration about our address. We may consider this request fair enough from the shipping agent, but not from the credit checking bureau. So, each BPEL4WS interactive request is supplemented with a special tag [root/context]:

- root requests will be compiled with a `<flow>` construct and returned together with the overall result of the computation for contextual requests;
- contextual requests the PolicyOrchestrator will make a copy of the WS process (*not* the authorization process) and replace each step *S* for which an additional request *I* has been called with the request and a context indicating the WS (partner and all) that required the additional credential. The PolicyOrchestrator will then prune the WS process removing all nodes that were not on a path from the root to the newly modified nodes and sends the result to the Client.

The last step is necessary to protect the overall workflow from unnecessary disclosure. Notice that the context workflow is not to be executed by the client. The client must only execute the authorization workflow. So, to avoid complications, we have used a container where the context workflow is stored for the client to see it, if he wants to know why certain credentials are needed.

This combination is sufficiently adequate for most uses, but still it offers the PolicyOrchestrator just the choice of

compiling individual requests rather than combining them. Here we have identified an important point in the PolicyOrchestrator where we need to introduce a new language - a *language for combination of policies and interactive requests at workflow level*. So far, we have not found a proposal that is entirely satisfactory, part because there are not enough case studies of WS Business Processes to guide the selection of policies at workflow level.

The proposal by Bertino et al. [2], is fairly expressive but only focuses on implementing snapshot constraints on a workflow level (i.e. safety properties). So it is not possible to express properties such as "if *Y* is repeatedly true then eventually *X* should happen".

The usage of algebraic constructs based on dynamic logic proposed by Wijesekera and Jajodia [34] seems more promising. Indeed `<invoke>` operation would be mapped into single action, `<sequence>` into sequential compounder, `<switch>` into non deterministic choice (each case represented by a test) and `<flow>` by intersection. This does not mean that we would use dynamic logic for actual implementation⁴, but rather that the logical language may offer a formal foundation to policy written in BPEL4WS.

6. THE ABDUCTION OF MISSING CREDENTIALS

For the deployment of the architecture, the PolicyEvaluator must be able to determine the set of additional creden-

⁴This is less critical than prejudice may suggest. The ML implementation of Peter Patel-Schneider at Bell-Labs can actually crack significant dynamic logic theorems in milliseconds.

tial that are necessary to obtain a service in case of failure. This problem may of course be shifted on the implementors of `PolicyEvaluators`, as the architecture only needs that the outcome of this derivation is mapped into some BPEL4WS process that is then sent to the client.

However, there is no algorithm in either the formal or the practical models of access control and trust negotiations to derive such credentials from the access control policy. The works on trust negotiations [29, 36] focus on communication and infrastructure and assume that requests and counter requests can be somehow calculated from the access policy. The formal models on credential-based access control and policy combination [2, 23, 15, 34] don't treat the problem of inferring missing credentials from failed requests, as they are within the frame of mind of inferring successful requests from present credentials. Also standardization efforts like the XACML proposals [12] gives rules for deriving what is right (evaluating policies) and not rule for understanding what went wrong.

Also a recent proposal by Bonatti and Samarati [7] that has the explicit focus on access and release control is too preliminary and unsatisfactory. In a nutshell, the request is received, the policy rules are filtered for relevance, the relevant rules are partially evaluated and sent to the client. The client will have to figure out which are the credentials (this is not discussed in the paper) and then will evaluate these credentials according its release policy.

The first problem is that demanding clients to analyse security policies is not acceptable here. We only assume an interpreter of Business Processes on the client side (possibly with some crypto capabilities if some digital signatures are needed) and thus all analysis of logical policies should be performed elsewhere. The second problem is that after a suitable number of queries the entire policy of the server would be disclosed to the client or to the server orchestrating the process. This is hardly acceptable from the perspective of a WS business partner. Furthermore, the relevancy filtering approach only works for flat policies, in which for every request we list all its credentials. The relevancy selection procedure in [7] is not correct already for the simple example that we show in Figure 5.

The other key proposal on trust negotiation by Yu et al. [36], offers a dual view w.r.t Bonatti and Samarati [7]. Loosely speaking, each credential is associated to a policy (a boolean expression) denoting the credentials that a client must have already provided for its safe disclosure, by a step wise process the parties can exchange credentials or policy rules (as in Bonatti and Samarati [7]) until the desired resource is released. The papers provide for safe sequences of disclosure in a rather ad-hoc fashion building upon trees rather than logical formalization. As a consequence they can only treat monotone policies and it is not possible to define notions of consistency of policies and disclosure of policies in presence of constraints (e.g. separation of duty). The major limitation of the paper is that it interlocks the access and the release policy into one. So, as the authors acknowledge [36, page. 21], it is impossible to access resources if some of the needed credentials cannot be disclosed at some point. Furthermore, the need for intermediate credential disclosure calls for a structuring of policy rules that is counter-intuitive from the point of view of access control. For instance, a policy rules may say that for access to the full text of on-line journal article we must have already got the access to brows-

ing the journal table of content, plus additional credentials. Access to table of contents could then specify some simpler set of credentials. For the disclosure process to take place such natural composition is not possible when using Yu et al. framework [36].

We propose a more general and principled approach based on logic that allows for a clean solution of these problems. For sake of simplicity (and popularity), assume that the policy is expressed using Datalog rules or logic programs with the stable model semantics (if we need negation to implement some constraints like separation of duties). What we need is a logical implementation of the following process:

1. the `PolicyEvaluator` receives the credentials and evaluates the request against the policy augmented with the credentials i.e. whether the request is a logical consequence of the policy and the credentials;
2. if the request is granted nothing needs to be done;
3. if the request fails we evaluate the given credential against a release policy of the `PolicyEvaluator` to infer which are the credentials whose need can be disclosed on the basis of the credentials already received;
4. abduce the actually needed credentials by re-evaluating the request against the policy and considering the potentially disclosable credentials determined at the previous step; only the needed credential are communicated to the client.

In a nutshell, what we need for the implementation of `PolicyEvaluator` is to implement two main inference capabilities: *deduction* and *abduction* [31]. We need to use deduction to infer whether a request can be granted on the basis of the present credentials as in [7, 2, 23, 15], we use abduction to explain which minimum set of credentials would be necessary to grant a failed request. Obviously it is not necessary to use logic, what we claim is that the underlying logical constructs that we need for our access decisions are these two conceptually different operation.

DEFINITION 6.1 (ACCESS CONTROL). *Let P_A be a stratified logic program representing an access control policy, r be an atom representing a request, C be a set of atoms representing a set of given credentials, then the request is granted iff $P_A \cup C \models r$.*

DEFINITION 6.2 (RELEASE CONTROL). *Let P_R be a stratified logic program representing a release control policy, d be an atom representing a credential, C be a set of atoms representing a set of given credentials, then the credential d is disclosable iff $P_R \cup C \models d$.*

The notion of release control subsumes the notion of "policy" that is used by Yu et al. [36]. Indeed, a step of the negotiation process by trust builder can now be explained either as a successful entailment (the disclosure of a credential) or the disclosure of a logic rule.

DEFINITION 6.3 (AC FAILURES EXPLANATION). *Let P_A and P_R be a stratified logic programs representing respectively an access control policy and a release control policy, r be an atom representing a request, C be a set of atoms representing a set of given credentials, an admissible explanation(abduction) of missing credentials is any set of credentials C_M such that*

1. $P_A \cup C \not\models r$
2. $P_A \cup C \cup C_M \models r$
3. $P_A \cup C \cup C_M$ is consistent
4. $P_R \cup C \models c$ for all $c \in C_M$

The first condition says that the missing credentials are indeed needed. The second condition says that they are sufficient. The third one says that they are actually meaningful and don't lead to inconsistency (say because of separation of duties). The last condition specifies that credential must be disclosable. In presence of positive Datalog program such as for Bonatti and Samarati's logic [7], Li's Delegation Logic [23], Samarati et al. authorization framework [30], the consistency condition is satisfied by default. In presence of constraints on the execution or negation as failure, as in Bertino et al. Datalog programs for workflow policies [2] — which can be easily augmented with credentials — the consistency condition is essential to guarantee that the abduced set of atoms makes sense. Indeed, constraints could make $P_A \cup C \cup C_M$ inconsistent and therefore it would not make much sense to say that the request r should be granted from a system.

In most definitions of abduction we also have constraints on the *minimality of the solution* wrt some partial ordering. Traditional requirements are minimality wrt set-containment and set-cardinality. In presence of role hierarchies additional ordering can be defined. These requirements make the problem harder than deduction from a computational complexity point of view but are usually desirable from the view points of simplicity of interaction and information-flow control. We discuss this issue in details in another paper [20].

In Figure 5 is shown a sample policy of the university online library access and release rules. The notations for declarations, credentials and services are borrowed from Bonatti and Samarati [7]. Here `decl` means that it is a statement (e.g., identity, address) declared by the client, while `cred` is a statement declared and signed by a key corresponding to some trusted authority. Rule 4 that says "to have access to service `reading` the client should have access to library (presenting Id and some library card) and a loan library card". Rule 10 says "to reveal the need for a loan library credential there should be a declaration of the library's Id and some library credential".

Notice that here is no way to disclose the need for a credential such as `cred(card(user, john, id1568), bibK)`. Such credential must be given. The same is true for the declaration about the university employee id which cannot be disclosed. However, the lack of a rule for disclosure of a credential does not forbid us to use the very same credential in some access rule.

If the `PolicyEvaluator` is given the declaration `decl(id1568)` and the credential `cred(card(user, john, id1568), bibK)`, together with the request for reading the journal articles online. The query `serv(reading)` does not follow from the policy and the given declarations and credentials. So, we apply the release policy and infer that the following credentials are disclosable:

$$\begin{aligned} & \text{decl}(\text{john}, \text{cs}), \text{decl}(\text{id1568}), \\ & \text{cred}(\text{researcher}(\text{id1568}, \text{cs}), \text{csK}), \\ & \text{cred}(\text{card}(\text{user}, \text{john}, \text{id1568}), \text{bibK}), \\ & \text{cred}(\text{member}(\text{john}, \text{cs}), \text{csK}), \\ & \text{cred}(\text{card}(\text{loan}, \text{john}, \text{id1568}), \text{bibK}). \end{aligned}$$

The abduction algorithm derives two possible answers for the credentials:

$$\begin{aligned} C_{M1} &= \{\text{decl}(\text{john}, \text{cs}), \text{cred}(\text{member}(\text{john}, \text{cs}), \text{csK})\} \\ C_{M2} &= \{\text{cred}(\text{card}(\text{loan}, \text{john}, \text{id1568}), \text{bibK})\} \end{aligned}$$

Both sets are minimal with respect to the subset inclusion ordering and only C_{M2} is minimal with respect to a set cardinality ordering. In case the first set is chosen the `PolicyEvaluator` will compile a `<flow>` node for sending the requests back to the client.

It could be possible to avoid the presence of the release policy by adding an additional field to a credential that can be requested. Each time any such credential would be required to be true in the logical evaluation of the policy, we trigger an action sending the request to the client.

However, we believe that the separation of access and release policies is useful for practical reason in spite of the additional complication that the two-policies system requires for evaluation. The double query system is immaterial to the human administrator who might simply buy a faster machine. In contrast, the specification of policies is normally done by humans and is costly and error prone process. The integration of release and access policy into one policy would imply that a change in the release policy requires modification to the access policy which might have been unchanged. Furthermore the separation of access and release policies allows for separation of duties among administrators: one administrator can modify the access policy and another the release policy.

The use of abduction of missing credential is sufficient for the description of stateless business processes. Namely for business processes in which only the set of credential exhibited by the client is examined. If `PolicyEvaluators` store past credentials (for instance to avoid that somebody takes up two incompatible roles within a short time frame), it is necessary to revise the access control process to allow for the logical revocation of credentials. This is discussed elsewhere [20].

One of the limitations of the present system is the treatment of disjunctive information, such as "present either a VISA a Mastercard". In some cases it might be that there is a preference between the disjunct (e.g., one brand may charge the merchant a higher fee than the other) and this can be expressed by the p.o. used for the abductive solutions. In absence of ordering information, an arbitrary disjunct will be selected and the next interaction (if the credential is declined by the client) must inform the abduction algorithm that the previous solution is not acceptable. This can be easily done by adding the negation of the declined credentials to the set of presented credentials. However, from the point of view of the client, it could be more informative to receive the disjunction, but it is not clear whether this may turn out to be practical (it complicates the abduction algorithm) or may lead to leakage of information. The current implementation seems to suggest that having two interactions is essentially immaterial to the overall timings.

7. CONCLUSIONS AND RELATED WORK

As we have already discussed, a number of access control models have been proposed for workflows [2], web services [27], role based access control on the web [11, 28], SOAP

Access Policy:

- $$\begin{aligned} \text{serv}(\text{query}()) &\leftarrow \text{decl}(\text{Id}), \text{cred}(\text{card}(\text{Type}, \text{Name}, \text{Id}), \text{biblioK}) & (1) \\ \text{serv}(\text{query}(\text{citations})) &\leftarrow \text{serv}(\text{access}), \text{cred}(\text{member}(\text{Name}, \text{Dept}), K_D), \text{assoc}(\text{Dept}, K_D) & (2) \\ \text{serv}(\text{booking}) &\leftarrow \text{decl}(\text{Name}, \text{Dept}), \text{cred}(\text{card}(\text{loan}, \text{Name}, \text{Id}), \text{biblioK}) & (3) \\ \text{serv}(\text{reading}) &\leftarrow \text{serv}(\text{access}), \text{cred}(\text{card}(\text{loan}, \text{Name}, \text{Id}), \text{biblioK}) & (4) \\ \text{serv}(\text{reading}) &\leftarrow \text{cred}(\text{academic}(\text{Name}, \text{UnivId}), K_U), \text{assoc}(\text{university}, K_U) & (5) \\ \text{serv}(\text{reading}) &\leftarrow \text{serv}(\text{query}(\text{citations})), \text{cred}(\text{researcher}(\text{Name}, \text{Dept}), K_D), \text{assoc}(\text{Dept}, K_D) & (6) \end{aligned}$$

Release Policy:

- $$\begin{aligned} \text{decl}(\text{Name}, \text{Dept}) &\leftarrow \text{decl}(\text{Id}) & (7) \\ \text{cred}(\text{researcher}(\text{Name}, \text{Dept}), K_D) &\leftarrow \text{decl}(\text{Name}, \text{Dept}), \text{cred}(\text{card}(\text{Type}, \text{Name}, \text{Id}), \text{bibK}) & (8) \\ \text{cred}(\text{member}(\text{Name}, \text{Dept}), K_D) &\leftarrow \text{decl}(\text{Name}, \text{Dept}) & (9) \\ \text{cred}(\text{card}(\text{loan}, \text{Name}, \text{Id}), \text{bibK}) &\leftarrow \text{decl}(\text{Id}), \text{cred}(\text{card}(\text{Type}, \text{Name}, \text{Id}), \text{bibK}) & (10) \\ \text{cred}(\text{academic}(\text{Name}, \text{UnivId}), K_U) &\leftarrow \text{decl}(\text{UnivId}), \text{decl}(\text{Name}, \text{Dept}) & (11) \end{aligned}$$
-

Figure 5: University Library WS Access and Release Policies

messages [8], entire XML documents [1, 9], tasks [17] and DRM [27], possibly coupled by sophisticated policy combination algorithms. However, they have mostly remained within the classical framework. Even more liberal models such as those for DRM based on usage [27] has assumed that servers know their clients pretty well: they might not know their names but they know everything about what, when and how can be used by these clients. We have discussed the proposals of Bonatti and Samarati [7] and Yu et al. [36] more in details in Section 6.

If we look at the proposals for distributed access control architectures [35, 37, 4, 16] the common thread is decoupling access control logic from application logic and possibly distribute the access control component. However we are still within the same administrative boundaries.

For instance Woo and Lam [35] propose that the **ApplicationServer** offloads its authorization policy to an **AuthorizationServer**. After evaluating the policy the **AuthorizationServer** hands out authorization certificate to the **Client**, which the **Client** has to present along with its request.

An architecture close to ours has been proposed by Beznosov et al. [4]. Authorizations are managed by an **Authorization Service** and its **Access Decision Object (ADO)**. The **ADO** obtains references to all **PolicyEvaluators** related to the **Client's** request, asks a decision combinator for combining decisions according to a combination policy and returns the decision back to the **Client**. Also Akenti Policy Engine [16], OASIS system [13] and Adage system [37] share the idea of an **AuthorizationServer** communicating with application and various **IdentityServers** to obtain credentials for the **Client**.

In comparison with the OASIS framework [12]: the **ApplicationServer** acts as **PEP**; the **PolicyEvaluator** acts partly as **PAP** in the case of making available policies to the **AuthorizationServer** and acts partly as **PDP** in taking authorization decisions and providing them to the **AuthorizationServer**; the **AuthorizationServer** acts partly as "context handler" in receiving requests from the **ApplicationServer** and sending access decisions back to it and in collecting attributes from an **AttributeServer**. It acts partly as **PDP** in the case of taking an authorization decision from policies returned by **PolicyEvaluators** (acting as **PAPs**) and applying some rules on them;

the **PolicyOrchestrator** acts partly as "context handler" in requesting (giving a source to) the **AuthorizationServer** that interprets the source and returns the result back to it. It also acts partly as **PDP** in taking authorization decisions on the base of applying some policies available on the workflow level – acting in this case as **PAP**.

In most proposals, the possibility that servers may get back to the calling **Clients** with some counter requests is not considered. This even in the case where the **Client** is actually an **AuthorizationServer** querying different **PolicyEvaluator** servers.

In this paper we have proposed a solution to address the challenges of WS processes: a possible architecture for the authorization of business processes for Web services. We have identified an interactive access control model as a way for protecting security interests wrt disclosure of information and access control of both servers and clients. Logical abduction is the solid semantical foundation upon which interaction can be build.

In the model a **Client** interacts (contracts) with the servent in order to finalize the necessary set of credentials needed to satisfy all partners' requirements related to the process. We propose to use "mobile" processes as messages exchanged in the architecture and specify how entities in the architecture can be implemented using WS processes themselves.

Future work is in the direction of studying the complexity of the combined deduction and abduction process, for the particular restricted policy that are typically used in formulating workflow and WS security policies.

8. REFERENCES

- [1] BERTINO, E., CASTANO, S., AND FERRARI, E. On specifying security policies for Web documents with an XML-based language. In *Proceedings of the Sixth ACM Symposium on Access control models and technologies* (2001), ACM Press, pp. 57–65.
- [2] BERTINO, E., FERRARI, E., AND ATLURI, V. The specification and enforcement of authorization constraints in workflow management systems. *ACM Transactions on Information and System Security (TISSEC)* 2, 1 (1999), 65–104.

- [3] BETTINI, C., WANG, X. S., AND JAJODIA, S. An architecture for supporting interoperability among temporal databases. In *Temporal Databases: Research and Practice* (1998), vol. 1399 of *Lecture Notes in Computer Science*, Springer Verlag, pp. 36–55.
- [4] BEZNOV, K., DENG, Y., BLAKLEY, B., BURT, C., AND BARKLEY, J. A resource access decision service for CORBA-based distributed systems. In *Proceedings of 15th IEEE Annual Computer Security Applications Conference. (ACSAC '99)* (1999), IEEE Press, pp. 310–319.
- [5] BLAZE, M., FEIGENBAUM, J., IOANNIDIS, J., AND KEROMYTIS, A. D. The role of trust management in distributed systems security. In *Secure Internet programming: security issues for mobile and distributed objects*. Springer-Verlag, 1999, pp. 185–210.
- [6] BOB ATKINSON, ET AL. *Web Services Security (WS-Security)*. IBM, Microsoft, VeriSign, April 2002. <http://www-106.ibm.com/developerworks/webservices/library/ws-secure/>.
- [7] BONATTI, P., AND SAMARATI, P. A unified framework for regulating access and information release on the web. *Journal of Computer Security* 10, 3 (2002), 241–272.
- [8] DAMIANI, E., DI VIMERCATI, S. D. C., PARABOSCHI, S., AND SAMARATI, P. Fine grained access control for SOAP E-services. In *Proceedings of the tenth international conference on World Wide Web* (2001), ACM Press, pp. 504–513.
- [9] DAMIANI, E., DI VIMERCATI, S. D. C., PARABOSCHI, S., AND SAMARATI, P. A fine-grained access control system for XML documents. *ACM Transactions on Information and System Security (TISSEC)* 5, 2 (2002), 169–202.
- [10] FRANCISCO CURBERA, ET AL. *Business Process Execution Language for Web Services (BPEL4WS)*. BEA, IBM, Microsoft, 7 2002. <http://www-106.ibm.com/developerworks/webservices/library/ws-bpel/>.
- [11] GIURI, L. Role-based access control on the web. *ACM Transactions on Information and System Security (TISSEC)* 4, 1 (2001), 37–71.
- [12] GODIK, S., AND MOSES, T. *eXtensible Access Control Markup Language (XACML)*. OASIS, February 2003. www.oasis-open.org/committees/xacml/.
- [13] HINE, J. A., YAO, W., BACON, J., AND MOODY, K. An architecture for distributed OASIS services. In *IFIP/ACM International Conference on Distributed systems platforms* (2000), Springer-Verlag New York, Inc., pp. 104–120.
- [14] IMAMURA, T., DILLAWAY, B., AND SIMON, E. *XML-Encryption Syntax and Processing*. W3C, December 2002. <http://www.w3.org/TR/xmlenc-core/>.
- [15] JAJODIA, S., SAMARATI, P., SUBRAHMANIAN, V. S., AND BERTINO, E. A unified framework for enforcing multiple access control policies. In *Proceedings of the 1997 ACM SIGMOD international conference on Management of data* (1997), ACM Press, pp. 474–485.
- [16] JOHNSTON, W., MUDUMBAI, S., AND THOMPSON, M. Authorization and attribute certificates for widely distributed access control. In *Proceedings of Seventh IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE '98)* (1998), IEEE Press, pp. 340–345.
- [17] JOSHI, J. B. D., AREF, W. G., GHAFOR, A., AND SPAFFORD, E. H. Security models for web-based applications. *Communications of the ACM* 44, 2 (2001), 38–44.
- [18] KALER, C., AND NADALIN, A. *Web Services Secure Conversation (WS-SecureConversation)*. IBM and Microsoft, 12 2002. <http://www.ibm.com/developerworks/library/ws-secon/>.
- [19] KALER, C., AND NADALIN, A. *Web Services Trust Language (WS-Trust)*. IBM and Microsoft, 12 2002. <http://www.ibm.com/developerworks/library/ws-trust/>.
- [20] KOSHUTANSKI, H., AND MASSACCI, F. A logical model for security of web services. Tech. Rep. IIT TR-10/2003, First International Workshop on Formal Aspects of Security and Trust (FAST), Istituto di Informatica e Telematica, Pisa, Italy, September 2003. Editors: Theo Dimitrakos and Fabio Martinelli.
- [21] KREGER, H. *Web Services Conceptual Architecture (WSCA 1.0)*, May 2001. <http://www-3.ibm.com/software/solutions/webservices/pdf/WSCA.pdf>.
- [22] KUDO, M., AND HADA, S. XML document security based on provisional authorization. In *Proceedings of the 7th ACM conference on Computer and Communications Security* (2000), ACM Press, pp. 87–96.
- [23] LI, N., GROSOFF, B. N., AND FEIGENBAUM, J. Delegation logic: A logic-based approach to distributed authorization. *ACM Transactions on Information and System Security (TISSEC)* 6, 1 (2003), 128–171.
- [24] LI, N., MITCHELL, J. C., AND WINSBOROUGH, W. H. Design of a role-based trust-management framework. In *Proceedings of IEEE Symposium on Security and Privacy, 2002. S&P* (2002).
- [25] MARK BARTEL, ET AL. *XML-Signature Syntax and Processing*. W3C, IETF, February 2002. <http://www.w3.org/TR/xmlsig-core/>.
- [26] OASIS SECURITY SERVICES TC. *Security Assertion Markup Language (SAML)*. OASIS, November 2002. www.oasis-open.org/committees/security/.
- [27] PARK, J., AND SANDHU, R. Towards usage control models: beyond traditional access control. In *Seventh ACM Symposium on Access Control Models and Technologies* (2002), ACM Press, pp. 57–64.
- [28] PARK, J. S., AND SANDHU, R. RBAC on the Web by smart certificates. In *Proceedings of the fourth ACM workshop on Role-based access control* (1999), ACM Press, pp. 1–9.
- [29] ROSCHEISEN, M., AND WINOGRAD, T. A communication agreement framework for access/action control. In *Proceedings of the Symposium on Security and Privacy* (1996), IEEE Press, pp. 154–163.
- [30] SAMARATI, P., REITER, M. K., AND JAJODIA, S. An authorization model for a public key management service. *ACM Transactions on Information and*

System Security (TISSEC) 4, 4 (2001), 453–482.

- [31] SHANAHAN, M. Prediction is deduction but explanation is abduction. In *Proceedings of IJCAI '89* (1989), Morgan Kaufmann, pp. 1055–1060.
- [32] W3C. Web Services Architecture.
<http://www.w3.org/TR/ws-arch>.
- [33] WEEKS, S. Understanding trust management systems. In *IEEE SS&P-2001* (2001).
- [34] WIJESSEKERA, D., AND JAJODIA, S. Policy algebras for access control the predicate case. In *Proceedings of the 9th ACM conference on Computer and Communications Security* (2002), ACM Press, pp. 171–180.
- [35] WOO, T. Y. C., AND LAM, S. Designing a distributed authorization service. In *Proceedings of Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies. INFOCOM* (1998), vol. 2, IEEE Press, pp. 419–429.
- [36] YU, T., WINSLETT, M., AND SEAMONS, K. E. Supporting structured credentials and sensitive policies through interoperable strategies for automated trust negotiation. *ACM Transactions on Information and System Security (TISSEC)* 6, 1 (2003), 1–42.
- [37] ZURKO, M., SIMON, R., AND SANFILIPPO, T. A user-centered, modular authorization service built on an RBAC foundation. In *Proceedings of the IEEE Symposium on Security and Privacy* (1999), IEEE Press, pp. 57–71.