

An Access Control System 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— Web Services and Business Processes for Web Services are the new paradigms for the lightweight integration of business from different enterprises.

Security and access control policies for Web Services protocols and distributed systems are well studied and almost standardized, but there is not yet a comprehensive proposal for an access control architecture for business processes. The major difference is that business processes describe complex services that cross organizational boundaries and are provided by entities that sees 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.

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.

I. INTRODUCTION

Middleware has been the enterprise integration buzzword at the end of the past millennium. Nowadays a new paradigm is starting to take 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.

Basic WS are well studied and standardized, for what concerns access control and security. There are also many approaches [1], [2], [3] for controlling access to services in distributed systems, and an advanced standardization process (see for instance the OASIS XACML [4] proposal). With the notable exception of provisional access control [5] and trust negotiation [6], 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 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 “Shakespearian Tour” in Italy: you might decide to go to the city of Shylock, and 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 Pompeous 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 the system has to be close to trust management systems [3];
- 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, 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 the 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 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 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 [7], [8] and at the architectural level [4]. We have proposals for calculi for controlling release of information [9], and procedures for trust negotiations and communication of credentials [6], architecture for distributed access control [4], [2], [1].

What is missing is a way to synthesize *all* these aspects into one access control architecture for business processes of WS, which is the contribution of this 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. Section V explains how we can use logical deduction and logical abduction to build a firm foundation for the interactive process of inferring discloseable credentials from access control policies and from release policies. Next we discuss how everything can be implemented using Business Process themselves. A brief discussion of related works concludes the paper.

II. A PRIMER ON WS AND BUSINESS PROCESSES

A Web Service as defined by the standard² 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.”

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.

²Web Services Conceptual Architecture (WSCA), <http://www-3.ibm.com/software/solutions/webservices/pdf/WSCA.pdf>

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

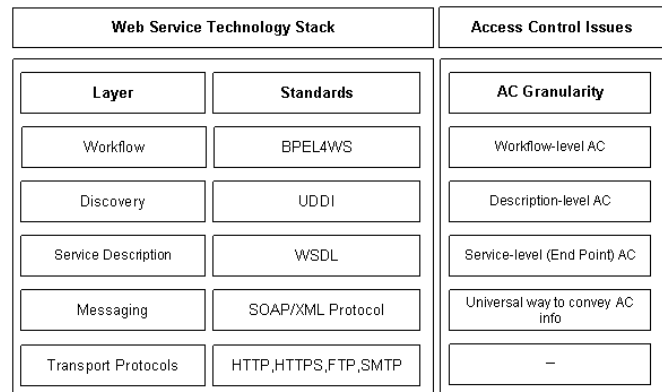


Fig. 1. Web Services Technology Stack & Access Control Issues

The WS architecture, as defined by W3C³, is divided into five layers grouped into three main components - Wire, Description, and Discovery (Fig. 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)⁴.

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 a developer to define an ordered sequence of steps;

³W3C. Web Services Architecture. <http://www.w3.org/TR/ws-arch>.

⁴BPEL4WS specification – <http://www-106.ibm.com/developerworks/webservices/library/ws-bpel/>

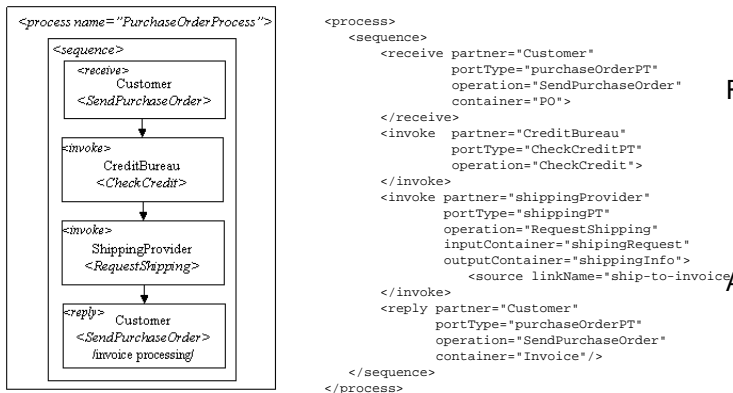


Fig. 2. Example of BPEL4WS Process

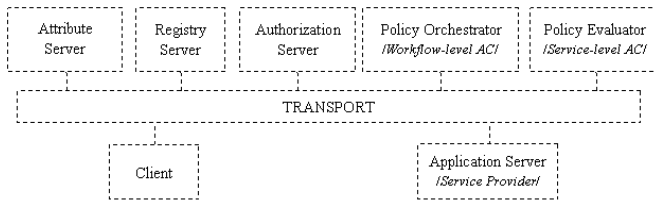


Fig. 3. Cross-section view of the architecture

- <switch> - allows a developer to have branching;
- <while> - allows a developer to define a loop;
- <flow> - allows a 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 continue 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 contents is self explanatory.

III. 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 a cross-section view of the architecture, whereas Figure 4 shows a horizontal view of it. A brief description of the servers shown in the figure is given below.

AttributeServer is responsible for providing group/role membership information as in [1], 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 [2], 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.

To secure the entire architecture we must make some assumptions on the security properties of the lower levels. Obviously we assume authentication, confidentiality, and message integrity at the transport and message levels. So, we assume that we have already in place the proposed standards.

At transport level we assume the adoption of the WS-Security specification⁵ 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-Signature⁶ and W3C XML-Encryption⁷, or the recently release specifications by IBM and Microsoft for WS secure conversations⁸.

Assuming security at lower level, the second key component is the languages and format of communications. We propose here a major innovation: the typical exchange of messages in an 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. It means that instead of sending just messages that have to be interpreted by entities, we truly have

⁵WS-Security – <http://www-106.ibm.com/developerworks/webservices/library/ws-secure>

⁶XML-Signature – <http://www.w3.org/TR/xmlsig-core>

⁷XML-Encryption – <http://www.w3.org/TR/xmlenc-core>

⁸WS-SecureConversation – <http://www.ibm.com/developerworks/library/ws-seccon>

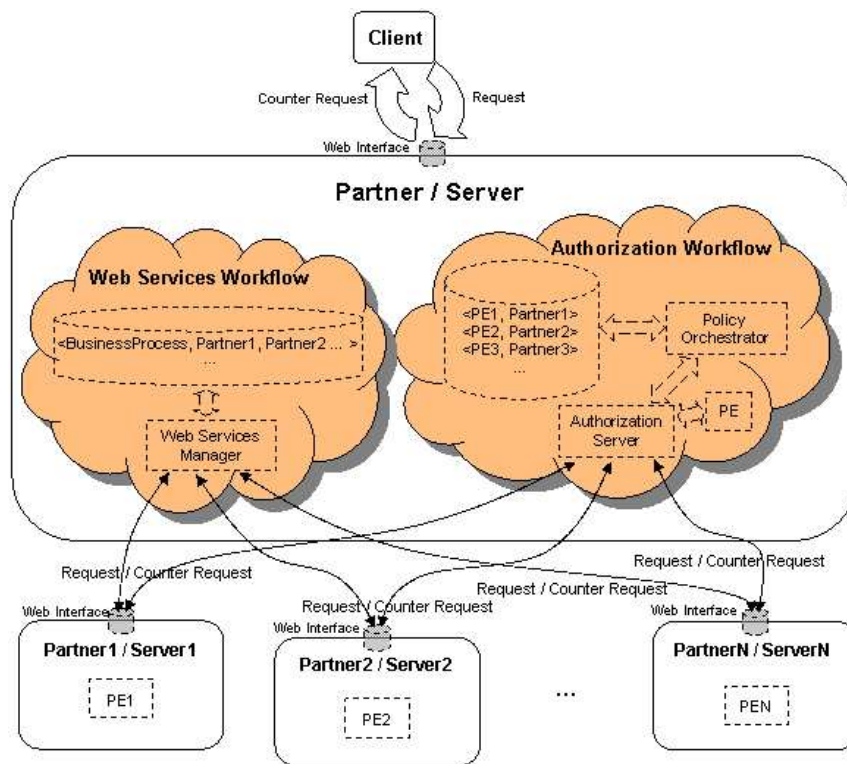


Fig. 4. Horizontal view of the architecture

mobile processes passing from one entity to another indicating themselves what the recipient has to do.

The mobility of authorization processes has a number of advantages. First of all a server simply needs an off-the-shelf interpreter for business processes for a quick implementation. Second we have more flexibility for describing the process leading to an access control decision. Some **PolicyEvaluators** may decide to disclose its XACML policies and therefore send a mobile process, which just describes the evaluation of the policies along some XACML rules. Other **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 so that one can accommodate also provisional access control or the interactive version that we advocate here.

Leading this approach at an extreme the **Authorization-Server** 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 sources⁹.

The role of the **PolicyEvaluator** is to encapsulate the connected with it partner's specific access control model, authorization policy, and requirements with their internal representation, interpretation, and mechanisms for computing an access decision and presenting it as a service using standardized Web service interface (e.g., WSDL).

⁹Recall that we assume that authentication, integrity, and confidentiality are assured at message and transport level.

The entity burdened with constructing the authorization workflow (Figure 4) is the **PolicyOrchestrator**. The **PolicyOrchestrator** functionality can be considered as having two main tasks: first one, called *Policy Composition Service*, is to select which are the partners involved in the requested process and to combine the corresponding **PolicyEvaluators** in a policy composition process, and return it back to the **AuthorizationServer**. After the **AuthorizationServer** having finished the execution of the policy composition process it asks¹⁰ the **PolicyOrchestrator** for applying the workflow level release policies over the results from the execution – the second main task. The process of applying release control policies, called *Release Policy Service*, captures how the final authorization decision should be released to the **Client**.

IV. INTERACTIVE COMMUNICATIONS AS “MOBILE” PROCESSES

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. Flexibility because of 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,

¹⁰This is the case if it is specified in the policy composition process, i.e. depends on the security policies being applied in constructing the policy composition process.

having little functionality pre-engineered into them, as we will see in section VI.

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 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 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 must be 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.

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. [10], 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 [8] 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 implementation¹¹, but rather that the logical language may offer a formal foundation to policy written in BPEL4WS.

V. THE ABDUCTION OF MISSING CREDENTIALS

For the deployment of the architecture, the PolicyEvaluator must be able to determine the set of additional credential 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 [11], [6] 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 [10], [7], [8] 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 [4] gives rules for deriving what is right (evaluating policies) and not rule for understanding what is wrong.

Here, we present an 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;

¹¹This 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.

- 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* [12]. We need to use deduction to infer whether a request can be granted on the basis of the present credentials as in [9], [10], [7], 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 operations.

Due to lack of space, here we just give the basic hint of the formalization.

Definition 1 (Access Control): Let P be a datalog program (or stratified logic program) representing an access control policy, let r be an atom representing a request, let C be a set of atoms representing a set of given credentials, the *request is granted* if and only if $P \cup C \models r$.

Definition 2 (Release Control): Let P be a datalog program (or stratified logic program) representing a release control policy, let d be an atom representing a credential, let C be a set of atoms representing a set of given credentials, the *credential d is disclosable* if and only if $P \cup C \models d$.

Definition 3 (Access Control Explanation): Let P be a datalog program (or stratified logic program) representing an access control policy, let r be an atom representing a request, let C be a set of atoms representing a set of given credentials, let $D_P \supseteq C$ be a set of atoms representing disclosable credentials, an *explanation of missing credentials* $C_M \subseteq D_P$ such that

- 1) $P \cup C \not\models r$
- 2) $P \cup C \cup C_M \models r$
- 3) $P \cup C \cup C_M$ is consistent

The first conditions says that the missing credentials are indeed needed. The second condition says that they are sufficient and the last condition says that they are actually meaningful. In presence of positive Datalog program such as for Bonatti and Samarati's logic [9] and Li's Delegation Logic 1 [7], 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 [10] — which can be easily augmented with credentials — the consistency condition is essential to guarantee that the abducted set of atoms makes sense. Indeed, constraints could make $P \cup C \cup C_M$ inconsistent and therefore it would not make much sense to say that the request r should be granted from a system.

In Figure 5 is shown a logic program showing a university online library access and release rules. The notations for dec-

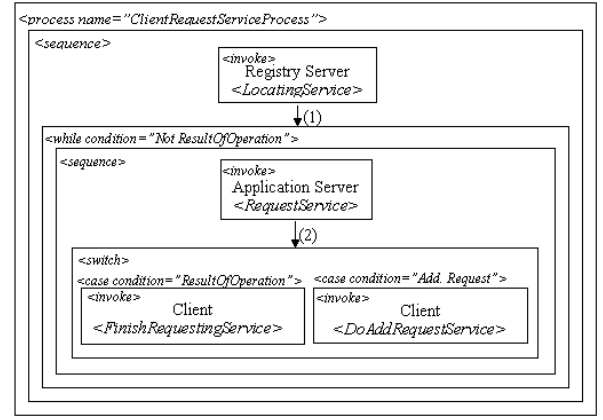


Fig. 6. Client Application Process Diagram

larations, credentials, and services are borrowed from Bonatti and Samarati [9]. 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. Consider 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".

If the PolicyEvaluator is given the declaration $\text{decl}(\text{id}1568)$ and the credential $\text{cred}(\text{card}(\text{user}, \text{john}, \text{id}1568), \text{bib}K)$, together with the request for reading the journal articles on-line. The query $\text{serv}(\text{reading})$ 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:

$$\begin{aligned}
 & \text{decl}(\text{john}, \text{cs}), \text{decl}(\text{id}1568), \\
 & \text{cred}(\text{researcher}(\text{id}1568, \text{cs}), \text{cs}K), \\
 & \text{cred}(\text{card}(\text{user}, \text{john}, \text{id}1568), \text{bib}K), \\
 & \text{cred}(\text{member}(\text{john}, \text{cs}), \text{cs}K), \\
 & \text{cred}(\text{card}(\text{loan}, \text{john}, \text{id}1568), \text{bib}K).
 \end{aligned}$$

The abduction algorithm derive two possible answers for the credentials:

$$\begin{aligned}
 C_{M1} &= \{\text{decl}(\text{john}, \text{cs}), \text{cred}(\text{member}(\text{john}, \text{cs}), \text{cs}K)\} \\
 C_{M2} &= \{\text{cred}(\text{card}(\text{loan}, \text{john}, \text{id}1568), \text{bib}K)\}
 \end{aligned}$$

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.

VI. COMPONENT ALGORITHMS AS BUSINESS PROCESSES

This section shows how we can describe entities in our architecture and how they can communicate each other using BPEL4WS specification.

The Client process is shown in Figure 6. In the figure, after the Client has requested the ApplicationServer for getting a service R , presenting its credentials, there are two cases:

Access Policy:

- serv(query()) ← decl(Id), cred(card(Type, Name, Id), biblioK) (1)
- serv(query(citations)) ← serv(access), cred(member(Name, Dept), K_D), assoc(Dept, K_D) (2)
- serv(booking) ← decl(Name, Dept), cred(card(loan, Name, Id), biblioK) (3)
- serv(reading) ← serv(access), cred(card(loan, Name, Id), biblioK) (4)
- serv(reading) ← cred(academic(Name, UnivId), K_U), assoc(university, K_U) (5)
- serv(reading) ← serv(query(citations)), cred(researcher(Name, Dept), K_D), assoc(Dept, K_D) (6)

Release Policy:

- decl(Name, Dept) ← decl(Id) (7)
- cred(researcher(Name, Dept), K_D) ← decl(Name, Dept), cred(card(Type, Name, Id), bibK) (8)
- cred(member(Name, Dept), K_D) ← decl(Name, Dept) (9)
- cred(card(loan, Name, Id), bibK) ← decl(Id), cred(card(Type, Name, Id), bibK) (10)
- cred(academic(Name, UnivId), K_U) ← decl(UnivId), decl(Name, Dept) (11)

Fig. 5. University Library WS Access and Release Policies

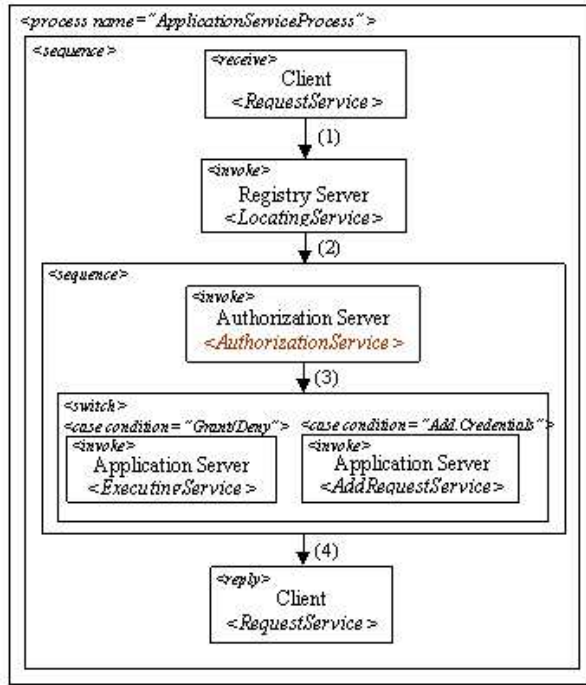


Fig. 7. Application Server Process Diagram

Additional Request - in this case is returned a counter request (a process), indicating what should be done by the Client. After that locally is invoked a service *DoAddRequestService* for executing the required process. Because of the while loop again is requested the service *R* with the result of the process; *ResultOfOperation* - in this case is returned the result of the requested service *R* and the Client's process finishes. The ApplicationServer, after the Client's request for accessing the service *R*, asks the RegistryServer (step 1 in Figure 7) for locating its *AuthorizationService*. After that the *AuthorizationService* is invoked along with Client's credentials and the requested service *R* for taking the authorization decision (step 2 in Figure 7). Then we can switch between explicit

Grant/Deny response returned from the AuthorizationServer in the case of which is executed or not the requested service *R* and the results are returned back to the Client (step 4 in Figure 7), or in the case of additional credentials is executed the *AddRequestService*, which either executes some counter-requirements that have to be presented to the Client or redirects the entire request to the Client (step 4 in Figure 7).

The AuthorizationServer process, shown in Figure 8, is the following: after the *AuthorizationService* has been invoked by the ApplicationServer the *PolicyCompositionService* located in the PolicyOrchestrator is invoked. The result of the service invocation (step 1 in Figure 8) is a policy composition process (e.g., BPEL4WS) indicating what should be done by the AuthorizationServer in order to be taken the final authorization decision. After obtaining the process (step 2 in Figure 8), the AuthorizationServer starts executing it, requesting all needed PolicyEvaluators with respect to that process, i.e. some of them in parallel, others in a sequence etc. Here the policy composition process consists of a sequence indicating that first the AuthorizationServer has to execute all PolicyEvaluators relevant to the requested service *R* orchestrated in a specific way (where the most intuitive structure is a <flow> one indicating execution in parallel, as shown in Figure 8), and after that executing the *ReleasePolicyService* responsible for taking the final access decision. After finishing the policy composition process, the AuthorizationServer returns the final access decision to the ApplicationServer (step 4 in Figure 8).

VII. CONCLUSIONS AND RELATED WORK

As we have already discussed, a number of access control models have been proposed for workflows [10], role based access control on the web [13], entire XML documents [14], [15], tasks [16], and DRM [17], possibly coupled by sophisticated policy combination algorithms. However, they have mostly remained within the classical framework where servers know their clients pretty well: they might not know their names

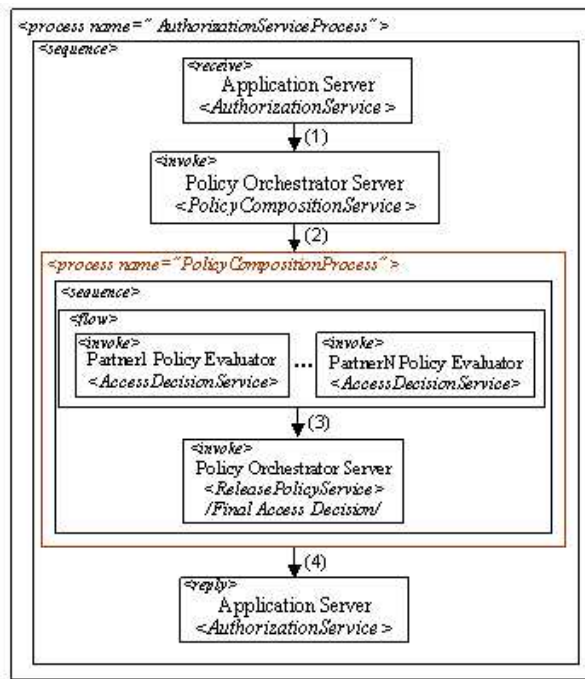


Fig. 8. Authorization Server Process Diagram

but they know everything about what, when, and how can be used by these clients.

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 one of the earliest work on distributed access control by Woo and Lam [1] 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. [2]. 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.

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 built.

In the model a Client interacts (contracts) with the servant 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 specified how entities in the architecture

can be implemented using WS processes themselves.

REFERENCES

- [1] T. Y. C. Woo and S. Lam, "Designing a distributed authorization service," in *Proceedings of Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies. INFOCOM*, vol. 2. IEEE Press, 1998, pp. 419–429.
- [2] K. Beznosov, Y. Deng, B. Blakley, C. Burt, and J. Barkley, "A resource access decision service for CORBA-based distributed systems," in *Proceedings of 15th IEEE Annual Computer Security Applications Conference. ACSAC '99*. IEEE Press, 1999, pp. 310–319.
- [3] M. Blaze, J. Feigenbaum, J. Ioannidis, and A. D. Keromytis, "The role of trust management in distributed systems security," *Secure Internet programming: security issues for mobile and distributed objects*, pp. 185–210, 1999.
- [4] S. Godik and T. Moses, *eXtensible Access Control Markup Language (XACML)*, OASIS, February 2003, www.oasis-open.org/committees/xacml/.
- [5] M. Kudo and S. Hada, "XML document security based on provisional authorization," in *Proceedings of the 7th ACM conference on Computer and Communications Security*. ACM Press, 2000, pp. 87–96.
- [6] T. Yu, M. Winslett, and K. E. Seamons, "Supporting structured credentials and sensitive policies through interoperable strategies for automated trust negotiation," *ACM Transactions on Information and System Security (TISSEC)*, vol. 6, no. 1, pp. 1–42, 2003.
- [7] N. Li, B. N. Grosz, and J. Feigenbaum, "Delegation logic: A logic-based approach to distributed authorization," *ACM Transactions on Information and System Security (TISSEC)*, vol. 6, no. 1, pp. 128–171, 2003.
- [8] D. Wijesekera and S. Jajodia, "Policy algebras for access control the predicate case," in *Proceedings of the 9th ACM conference on Computer and Communications Security*. ACM Press, 2002, pp. 171–180.
- [9] P. Bonatti and P. Samarati, "A unified framework for regulating access and information release on the web," *Journal of Computer Security*, vol. 10, no. 3, pp. 241–272, 2002.
- [10] E. Bertino, E. Ferrari, and V. Atluri, "The specification and enforcement of authorization constraints in workflow management systems," *ACM Transactions on Information and System Security (TISSEC)*, vol. 2, no. 1, pp. 65–104, 1999.
- [11] M. Roscheisen and T. Winograd, "A communication agreement framework for access/action control," in *Proceedings of the Symposium on Security and Privacy*. IEEE Press, 1996, pp. 154–163.
- [12] M. Shanahan, "Prediction is deduction but explanation is abduction," in *Proceedings of IJCAI '89*. Morgan Kaufmann, 1989, pp. 1055–1060.
- [13] L. Giuri, "Role-based access control on the web," *ACM Transactions on Information and System Security (TISSEC)*, vol. 4, no. 1, pp. 37–71, 2001.
- [14] E. Bertino, S. Castano, and E. Ferrari, "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*. ACM Press, 2001, pp. 57–65.
- [15] E. Damiani, S. D. C. di Vimercati, S. Paraboschi, and P. Samarati, "A fine-grained access control system for XML documents," *ACM Transactions on Information and System Security (TISSEC)*, vol. 5, no. 2, pp. 169–202, 2002.
- [16] J. B. D. Joshi, W. G. Aref, A. Ghafoor, and E. H. Spafford, "Security models for web-based applications," *Communications of the ACM*, vol. 44, no. 2, pp. 38–44, 2001.
- [17] J. Park and R. Sandhu, "Towards usage control models: beyond traditional access control," in *Seventh ACM Symposium on Access Control Models and Technologies*. ACM Press, 2002, pp. 57–64.