

Trust But Verify: Authorization for Web Services

Christian Skalka
Department of Computer Science
University of Vermont
skalka@cs.uvm.edu

X. Sean Wang
Department of Computer Science
University of Vermont
xywang@cs.uvm.edu

ABSTRACT

Through web service technology, distributed applications can be built in a flexible manner, bringing tremendous power to applications on the web. However, this flexibility poses significant challenges to security. In particular, an end user (be it human or machine) may access a web service directly, or through a number of intermediaries, while these intermediaries may be formed on the fly for a particular task. Traditional access control for distributed systems is not flexible and efficient enough in such an environment. Indeed, it may be impossible for a web service to anticipate all possible access patterns, hence to define an appropriate access control list beforehand. Novel solutions are needed.

This paper introduces a trust-but-verify framework for web services authorization, and provides an implementation example. In the trust-but-verify framework, each web service maintains authorization policies. In addition, there is a global set of “trust transformation” rules, each of which has an associated transformation condition. These trust transformation rules convert complicated access patterns into simpler ones, and the transformation is done by a requester (the original requester or an intermediary) with the assumption that the requester can be trusted to correctly omit certain details. To verify authorization, the requester is required to document evidence that the associated transformation conditions are satisfied. Such evidence and support information can either be checked before access is granted, or can be verified after the fact in an offline mode, possibly by an independent third party.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—*Security and protection*

General Terms

Security, Design, Performance

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 Secure Web Services, October 29, 2004, Fairfax VA, USA.

Copyright 2004 ACM 1-58113-973-X/04/0010 ...\$5.00.

Keywords

Distributed Authorization, Access Control Logic, Web Services

1. INTRODUCTION

Web services promise a new era of flexibility and power for web applications. However, this power also promises security challenges for web service providers, especially regarding access control, a.k.a. *authorization*. In particular, an end user, either human or machine, may access a web service through a number of intermediary web services, possibly formed on the fly for the task, so that access control is not so simple as validating authorization for the invoker. Additionally, due to the pluralistic and volatile nature of the web, most web services cannot expect to anticipate all access patterns, or have foreknowledge of all possible user identities, so specifying access control policies via an access control list in the usual manner is unrealistic.

Furthermore, not only does the complexity of authorization in web services pose a basic theoretical problem, but also a practical one, since robust security solutions can be too costly for online checks, especially for high-volume web services that are expected to function much like RPC/RMI for applications— that is, expeditiously.

Consider a vendor that accepts credit card payments. The vendor needs to access the card-issuing bank to check the credit card information and obtain payment confirmation; this involves a distributed authorization problem. Strictly speaking, the bank should allow a vendor to charge a card holder only when the card holder authorizes the vendor to do so. Current practice, however, is for the bank to allow vendors access to its automated services without much proof that the card holder actually initiated the transaction. The reason for this is that (1) the actual verification is costly if done for each transaction, and (2) more importantly, it's in any reputable vendor's best interest to be a good “citizen”, hence problems occur only rarely. In case of disputes, the bank can investigate the involved vendors and verify any questionable charges, possibly revising its assumptions about who is to be trusted in case a cheat is discovered. The bank can also request a general audit of vendors. From a vendor perspective, appropriate records should be kept in case of a dispute or an audit.

We call this approach a *trust-but-verify* authorization solution: in the online phase, vendors are trusted to a certain degree, but offline, the option to verify this trust can be pursued. We propose that a trust-but-verify approach is appropriate for web services security in general. Further-

A, B, C, \dots	\in	<i>Atom</i>	<i>atomic principals</i>
P	$::=$	$A \mid P \wedge P \mid P \vee P \mid P \mid P$	<i>principals</i>
p	\in	<i>Prop</i>	<i>primitive propositions</i>
s	$::=$	$p \mid \neg s \mid s \wedge s \mid P \Rightarrow P \mid P \text{ says } s$	<i>formulae</i>

Figure 1: ABLP Syntax

more, we argue that the relation between trust and verification should be meaningful— that is, what is checked offline should *provably* verify what is implicitly trusted online. To this end, we specify a formal framework for characterizing trust-but-verify systems, that requires any implementation to provide a *trust transformation* formalizing the relation between online and offline checking. This also allows a precise characterization of auditing. Returning to the above example, we can define a trust transformation that converts the request “card holder H requests charge on his card A” into “charge card A”; thus, the transformation always takes for granted that the request originates from the card-holder. This trust transformation also serves as the starting point for auditing; verification requires the vendor to deliver proof that the card holder did in fact make the request.

We establish a formal setting for our web service authorization framework in the Calculus of Access Control [2], which we call ABLP logic following previous references in the literature. The use of an access control logic for framework specification and system design is fundamental to our proposal, and we argue that this approach promotes unambiguous specification languages, reliability, and verifiability. Furthermore, as this presentation intends to describe and characterize our architecture, ABLP is at an appealing level of abstraction, allowing understandability and expressiveness in the high-level characterization presented here, but leaving flexibility for low-level implementation details in future work.

The contribution of this paper is to introduce the trust-but-verify framework for access control in web service environments, and formalize the framework conditions within ABLP logic. By characterizing the system within an authorization logic, we also propose rigorous logical foundation for authorization in web service environments.

1.1 Overview of the Paper

In Sect. 2, we provide a brief review of ABLP logic. In Sect. 3, we discuss and formalize the conditions that characterize the trust-but-verify framework in ABLP. In Sect. 4, we give an example of a particular trust-but-verify system, illustrating its adherence to framework conditions, and show how a transaction would be authorized and audited in the system. We conclude with remarks about related and future work in Sect. 5.

2. ABLP LOGIC

In this section we give a brief review of ABLP logic, focusing on those elements of the theory that are most relevant to this paper. For a thorough description and metatheory of the logic the reader is referred to [2].

2.1 Syntax of Principals and Formulae

The syntax of ABLP principals, constituting identities

$s_1 \vee s_2$	\triangleq	$\neg(\neg s_1 \wedge \neg s_2)$	$s_1 \supset s_2$	\triangleq	$\neg s_1 \vee (s_1 \wedge s_2)$
$A \text{ as } R$	\triangleq	$A \mid R$	$A \text{ controls } s$	\triangleq	$(A \text{ says } S) \supset S$

Figure 2: ABLP Formulae Abbreviations

in distributed communications, and formulae, representing statements and beliefs, is defined in Fig. 1. Regarding principles, we will mostly be concerned with atomic principles A and principles $P \mid Q$, pronounced “ P quoting Q ”.

Statements $P \text{ says } s$ generally represent that an assertion s has originated with a principal P . The relation $P \Rightarrow Q$, pronounced “ P speaks for Q ”, denotes that statements uttered by P can also be attributed to Q ; this is clarified in Sect. 2.3, which describes the derivation rules for the system.

2.2 Abbreviations and Notational Conventions

We assume that all principals can perform digital signatures. We let K range over public keys as principals, and always write K_A to denote the public keys of A , and K_A^{-1} the matching private key. The formula $K \text{ says } s$ represents the formula s encrypted under K .

A number of useful abbreviations are defined in Fig. 2. Along with macros for standard logical connectives, these include $A \text{ as } R$, denoting the principal obtained when A assumes the role R , and $A \text{ controls } s$, denoting that A is directly authorized for s . See [2] for a complete explanation of these abbreviations.

2.3 Proof Theory

We write $\vdash s$ to denote that a formulae s is logically derivable, on the basis of the axioms and inference rules of the theory. A selection of the ABLP inference rules, connecting the calculus of principals to the underlying propositional logic, is given in Fig. 3, specifically those which will be relevant to our presentation. We note that the rule names given are of our own devise, for easy reference in the remaining presentation. Also for convenience, we write $s \vdash s'$ iff s' is derivable given assumption s .

Here is an example showing how the logic can be used to model and reason about statements signed by digital signatures associated with particular principals.

EXAMPLE 2.1. *We trust that private keys remain indeed private, so that messages signed with K_J carry the authority of J :*

$$K_J \Rightarrow J$$

Thus, if any statement s is ever signed with J 's private key:

$$K_J \text{ says } s$$

<p style="text-align: center;">TAUT</p> $\frac{s \text{ is a tautology of propositional logic}}{\vdash s}$	<p style="text-align: center;">MODUS PONENS</p> $\frac{\vdash s' \quad \vdash s' \supset s}{\vdash s}$	<p style="text-align: center;">SUBTEXT</p> $\vdash A \text{ says } (s \supset s') \supset (A \text{ says } s \supset A \text{ says } s')$	<p style="text-align: center;">PARROT</p> $\frac{\vdash s}{\vdash A \text{ says } s}$
<p style="text-align: center;">QUOTE</p> $\vdash (B A) \text{ says } s \equiv B \text{ says } A \text{ says } s$	<p style="text-align: center;">SPEAKSFOR</p> $\vdash (A \Rightarrow B) \supset ((A \text{ says } s) \supset (B \text{ says } s))$		

Figure 3: Selected ABLP Inference Rules

By rule SPEAKSFOR:

$$(K_J \Rightarrow J) \supset (K_J \text{ says } s \supset J \text{ says } s)$$

hence by two applications of MODUS PONENS we have:

$$J \text{ says } s$$

That is, any signed message can be taken as a statement of the owner of the signature key.

2.4 Access Control Lists

Access control lists (ACLs) are fundamental to access control systems, providing an explicit association of principals with the privileges for which they're authorized. In the original presentation of ABLP [2], ACLs are conjunctions of statements of the form $P \text{ controls } s$, where s is some privilege. We adapt this approach, letting \mathcal{A} range over ACLs. Furthermore, we designate a subset of *Prop* as the set of *privileges* in the system, letting \mathbf{priv} range over this set. Hence, ACLs are conjunctions of statements $P \text{ controls } \mathbf{priv}$. Our justification for designating atomic propositions as privileges, and our use thereof, is discussed in Sect. 3.1 in more detail.

3. FRAMEWORK DEFINITIONS

Prior to the specifics of system design, the trust-but-verify framework can be precisely characterized as a set of conditions that any implementation must satisfy. In this section, we motivate and describe these conditions, which are stated at a sufficiently high level to allow flexibility in lower-level system design, but are mathematically rigorous.

3.1 Authorization Contexts and Decisions

Web services authorization is based on requests for the service made by invokers. Here we describe our proposed structure for these requests, and for the authorization decision predicated on them.

A *request* is an ABLP assertion s uttered by the invoker of a web service, which the invoker intends to be used by the web service for authorization of its use. In addition to the request, a web service may possess other facts and beliefs, e.g. ACLs and role certifications, that affect the authorization judgement; we assume that these facts and beliefs are expressed as ABLP formulae. The conjunction of these components constitutes an authorization *context*; authorization for a web service is granted upon a particular invocation iff the context of the invocation allows the privilege required for use of the web service to be derived in the ABLP proof theory.

As mentioned in Sect. 2, we posit a set of atomic formulae \mathbf{priv} , each of which represent the privilege required to access a particular web service. Authorization for \mathbf{priv} in a

context s is effected by checking validity of $s \vdash \mathbf{priv}$. Thus, our system is inspired by access control mechanisms such as stack inspection [20], which are specialized for program procedure calls, rather than challenge/response systems such as [5], which are adapted to human usage (i.e. web browsing) patterns. Our justification for this is that web service invocation bears a strong similarity to RPC/RMI, as observed in [11, 6], with chains of web service invocations resembling call stacks.

Here is a brief example illustrating the concepts of the framework described thus far:

EXAMPLE 3.1. *Suppose some web service WS requires the privilege \mathbf{priv} to be used, and the web service ACL \mathcal{A} grants this privilege to a principal D , i.e. $\mathcal{A} \equiv \mathcal{A}' \wedge D \text{ controls } \mathbf{priv}$ for some \mathcal{A}' . Suppose also that D invokes WS on its own behalf, making the request $D \text{ says } \mathbf{priv}$. Thus, the authorization context is $D \text{ says } \mathbf{priv} \wedge \mathcal{A}$. Clearly, $D \text{ says } \mathbf{priv} \wedge \mathcal{A} \vdash \mathbf{priv}$, since the context implies both $\vdash D \text{ says } \mathbf{priv}$ and $\vdash D \text{ controls } \mathbf{priv}$, which implies $\vdash \mathbf{priv}$ by MODUS PONENS.*

Naturally, it is desirable for authorization judgements to be decidable. Although ABLP logic is undecidable in general, various presentations have described non-trivial, decidable access control mechanisms [19, 5, 2]. Therefore, it is realistic to make this a formal condition of the trust-but-verify framework:

CONDITION 1. *Let s be an authorization context; then validity of $s \vdash \mathbf{priv}$ is decidable.*

This condition requires any trust-but-verify implementation to provide a decision procedure for validity of authorization judgements, and also implicitly requires the form of authorization contexts to be well-defined. As with all the framework conditions, the formal statement of the condition allows correctness of a decision procedure to be provable, i.e. implementations can (and should) be accompanied with proofs of their adherence to framework conditions.

3.2 Trust Transformations

The distinguishing characteristic of our proposal is the separation of online and offline checking phases, where in the online phase certain elements of authorization are taken for granted, or trusted to hold. This yields a simpler authorization decision, which can be verified more rigorously during the offline phase. However, with security at stake, vague accounts of the relation between these phases does not suffice—rather, we desire a formal relationship, so that offline *verification* of online authorization is meaningful. We embody this notion in the trust transformation, which specifies what elements of online authorization are to be taken for

granted, by specifying how to transform untrusted requests into trusted ones.

A *trust transformation* is a function from ABLP formulae to ABLP formulae. Any trust transformation's domain is formulae in *extrapolated* form, which take into account all components of access control, down to every detail verified during offline checking. The range of any trust transformation is formulae in *trusted* form, which are the “watered down” formulae that exclude restrictions the system takes for granted during online checking. The trust transformation mapping rigorously defines the relation between extrapolated and trusted forms. Since notions of trust can vary depending on the system, we specify the type and necessary preconditions of trust transformations, but the definition of the function itself is left up to a particular system designer. For any extrapolated formula s , we denote its trust transformation as $\llbracket s \rrbracket$. Any trust-but-verify implementation must define extrapolated and trusted authorization forms, and the trust transformation between them, with the requirement that it be a total function on the set of extrapolated statements. Also, we specify that the authorization contexts mentioned in Condition 1 are in trusted form.

EXAMPLE 3.2. *Suppose access control for a web service WS is based on requests made in both signed and unsigned form, so that all requests are of the form:*

$$B \text{ says } s \wedge K_B \text{ says } s$$

Suppose further that in the online component, players are trusted to communicate messages faithfully, and signatures are not checked. All authorization contexts include an ACL \mathcal{A} , which is left unchanged by the trust transformation. Thus, for all B and s , the trust transformation is defined as:

$$\llbracket B \text{ says } s \wedge K_B \text{ says } s \wedge \mathcal{A} \rrbracket = B \text{ says } s \wedge \mathcal{A}$$

Note that this transformation is total for the extrapolated form of requests in this example.

3.3 Auditing

While online checking takes trust into account, the purpose of offline checking is to verify that this trust is warranted. As the trust transformation injects trust into authorization, offline checking inverts the trust transformation, to verify online trust in the offline phase—that is, given a trusted request s , offline checking searches for an extrapolated request of which s is the trust transform. We call this process *auditing*, and require that any trust-but-verify implementation provide a function *audit* that given any trust-transformed context s , retrieves its extrapolated form.

The details of auditing constitute a significant engineering element of any trust-but-verify system. As will be exemplified in Sect. 4, we expect that certain elements of extrapolated statements will be logged for later offline retrieval during auditing, but what elements, and how and where they are logged and retrieve, is at a much lower level of detail than we're concerned with here. But at the abstract level we are concerned with, we can characterize the formal requirements of auditing.

Firstly, as mentioned above, *audit* returns the extrapolated form of trusted authorization contexts. Since the trust transformation has been defined formally, we can precisely characterize this condition as follows:

CONDITION 2. *Let s be a trusted context; then if $\text{audit}(s)$ succeeds, $\text{audit}(s) \vdash s'$ such that $\llbracket s' \rrbracket = s$.*

Note that this condition allows a certain degree of flexibility, in that *audit* must return a statement that is *at least* as strong as an extrapolated form of the input, not necessarily an extrapolated form per se.

Furthermore, we say *an* extrapolated form, since it is possible that any given trust transformation is many-to-one. Significantly, it is not even necessary that an extrapolated form of an authorized statement be authorized. However, since auditing seeks to verify trust implicit in an online check, we require that auditing not only return an extrapolated form of input statements, but one that is also authorized for the privilege in question; otherwise, auditing fails. This motivates the third condition of our framework:

CONDITION 3. *Let s be a trusted context and priv be a privilege. If $s \vdash \text{priv}$ holds, then so does $\text{audit}(s) \vdash \text{priv}$.*

It is important to note that this condition does not necessarily require theorem proving on extrapolated forms, but rather provability should follow by adherence to this condition generally. This point is revisited in more detail in Sect. 4.6. Here is an example that illustrates an auditing technique satisfying the specified conditions:

EXAMPLE 3.3. *Given both the online and offline checking scheme, as well as the trust transformation defined in Example 3.2, we define *audit* as follows. First, we assume that while trust transformations discard the signed portion of requests for online checking, the signed portion is actually saved (logged) as part of the implementation. Auditing of any request $B \text{ says } \text{priv}$ will then retrieve the signed portion of the original request discarded by the trust transformation, and verify that it is of the form $K_B \text{ says } \text{priv}$ by decrypting it, yielding the extrapolated context $B \text{ says } \text{priv} \wedge K_B \text{ says } \text{priv} \wedge \mathcal{A}$ as the result of auditing.*

Clearly, this example is vague with regard to how signed portions of requests are saved and retrieved. While these details are naturally addressed in implementations, we revisit this issue with some general suggestions for implementations in Sect. 4.

4. AN IMPLEMENTATION EXAMPLE

To illustrate our ideas, we now describe an example scenario with some implementation details. The example consists of a transaction involving a user, a target web service, and an intermediary web service. While this example is mainly intended to clarify the proposals of the previous section, we also make some substantive suggestions for trust-but-verify implementations in general, particularly advocacy of role keys, carrier authority, and logging.

4.1 Individuals

Many concrete entities take part in web service transactions; machines, users, applications, web services, domains, etc. A complete treatment would consider all possible players, since for example the same application run by the same user on two different machines might inspire different levels of trust, depending on the status of the machines. However, for the purposes of simplicity in this presentation, we will assume the existence of only two sorts of concrete entities

in web service transactions, *users* and *web services*, termed *individuals* taken together.

We also assume given finite, disjoint set of available principal identifiers for user and web services, respectively, and let J range over the former, and WS the latter. Furthermore, we will assume that these names have some known and decidable format, allowing automatic determination of whether a given individual is a user or a web service.

EXAMPLE 4.1. *We posit the following individuals: Joe is a user, and WS_M is an intermediary web service providing service to users, part of a system that includes access to a centralized medical database web service WS_{MDB} (to be continued...)*

4.2 Roles

As in many other systems, *roles* are an important component of our authorization scheme. In particular, due to the inherently volatile and popular nature of the web, web services cannot in general be expected to know the names of every possible invoker *a priori*; thus, authorization for privileges will be granted to known roles, and individuals must prove that they may assume claimed roles.

Furthermore, the same characteristics that inspire the use of roles for authorization, imply that role membership should not be established via explicit role membership lists. Rather, we posit that every known role R is associated with a key K_R , and the ability to sign messages with K_R (that is, proof of possession of the associated private key) is sufficient to establish role membership. Thus, we introduce the axiom:

$$\begin{aligned} &\text{ROLEKEY} \\ &\vdash K_R \text{ controls } (P \Rightarrow R) \end{aligned}$$

So for example, if an individual A wishes to legitimately assume a role R , it could make the statement $K_R \text{ says } (A \Rightarrow R)$; this and the assumed authority of role keys imply $A \Rightarrow R$.

Membership in roles is established by role certificates, the grammar for which is given in Fig. 4. They are conjunctions of role certifications, with **true** the empty certificate.

EXAMPLE 4.1. (Continued) *In our transaction example, we assume that Joe can take on the role of a doctor (denoted D). This means that Joe is capable of signing his requests using K_D . We assume a “trusted medical web services” role, denoted M , and that WS_M can take on role M . We also assume that the web service WS_{MDB} will authorize doctors to use WS_{MDB} . Hence, we have:*

$$D \text{ controls } \mathbf{priv}_{MDB}$$

is in \mathcal{A}_{MDB} . Note that we do not assume role M has access to WS_{MDB} . Role M will be used to “carry” doctor’s requests, as defined below.

4.3 Carrier Authority

In any access control statement $P \text{ controls } \mathbf{priv}$, it is not necessary for P to be atomic, allowing a fine-grained approach to access control. For example, if it is not desirable to grant a role R direct access to **priv**, but only on behalf of a principal D , the ACL can specify $R|D \text{ controls } \mathbf{priv}$, disallowing R direct access to **priv**.

However, we’re concerned with access control decisions in the presence of multiple intermediaries— that is, several intervening nodes may transport an authorization request

from source to target. The above scheme would require separate entries for every possible chain of intermediaries; for example, given statements:

$$\begin{aligned} &R_1 \text{ says } R_2 \text{ says } D \text{ says } \mathbf{priv} \\ &R_2 \text{ says } R_1 \text{ says } D \text{ says } \mathbf{priv} \end{aligned}$$

authorization for both would require both of the following statements to be present in the relevant ACL:

$$\begin{aligned} &R_1|R_2|D \text{ controls } \mathbf{priv} \\ &R_2|R_1|D \text{ controls } \mathbf{priv} \end{aligned}$$

Either that, or it would require access statements:

$$\begin{aligned} &R_1|D \text{ controls } \mathbf{priv} \\ &R_2|D \text{ controls } \mathbf{priv} \end{aligned}$$

to be present, along with known relations $R_1 \Rightarrow R_2$ and $R_2 \Rightarrow R_1$. The former solution is clearly cumbersome, unrealistically so given the number of possible intermediaries on the web. The latter is better, but is restrictive, requiring every intermediary to adopt the same role as, or a more powerful role than, its predecessor (although this problem could be alleviated by adapting the “speaks for regarding” relation proposed in [8] as an extension to ABLP). Since the only privilege at issue is **priv**, it is intuitively sufficient for each intermediary to have some sort of authorization for **priv**, as in e.g. stack inspection [20].

However, unlike stack inspection, web service intermediaries should often not be granted direct access to privileges— for example, a web service should not be granted direct access to withdraw cash from an individual’s bank account, but only on behalf of that individual. To maintain this property, and to overcome the drawbacks of the approaches described in the previous paragraph, we introduce the notion of a *carrier authority*. Intuitively, a carrier authority is the authority to carry an authorization request for a particular principal, but not the authority to make the request itself. Formally:

$$R \text{ carries } \mathbf{priv} \text{ for } D \triangleq (R|D \text{ says } \mathbf{priv}) \supset (D \text{ says } \mathbf{priv})$$

In the above example, access requires the carrier authorities $R_1 \text{ carries } \mathbf{priv} \text{ for } D$ and $R_2 \text{ carries } \mathbf{priv} \text{ for } D$, as well as the direct authority $D \text{ controls } \mathbf{priv}$. In general, carrier authority allows a fine-grained and flexible approach to authorization in the context of web services. We call conjunction of carrier authority statements *carrier control lists* (CCLs), which we denote C .

EXAMPLE 4.1. (Continued) *In our transaction example, M is not directly authorized for \mathbf{priv}_{MDB} , but it should be possible for trusted medical web services to carry this privilege for doctors, hence WS_{MDB} defines a carrier control list C_{MDB} that includes:*

$$M \text{ carries } \mathbf{priv}_{MDB} \text{ for } D$$

4.4 Extrapolated Statements

At a high level, we assume that extrapolated contexts possess the following characteristics:

1. All requests are made by individuals in a particular role.
2. All requests are signed by the requesting individual, to establish its authenticity (both on and offline).

$\mathcal{R} ::= K_R \text{ says } (A \Rightarrow R) \wedge \mathcal{R} \mid \mathbf{true}$	<i>role certificates</i>
$\mathbf{req} ::= K_{WS} \text{ says } R \text{ says } \mathbf{req} \mid K_J \text{ says } R \text{ says } \mathbf{priv}$	<i>extrapolated requests</i>
$\hat{s} ::= K_{WS} \text{ says } R \text{ says } \iota \mid K_J \text{ as } R \text{ says } \mathbf{priv} \mid \hat{s} \wedge \mathcal{R}$	<i>indexed statements</i>

Figure 4: Components of Authorization Requests

$$\begin{aligned} \llbracket \mathbf{req} \wedge \mathcal{R} \wedge \mathcal{C} \wedge \mathcal{A} \rrbracket &= \llbracket \mathbf{req} \rrbracket \wedge \mathcal{C} \wedge \mathcal{A} \\ \llbracket K_J \text{ says } R \text{ says } \mathbf{priv} \rrbracket &= R \text{ says } \mathbf{priv} \\ \llbracket K_{WS} \text{ says } R \text{ says } \mathbf{req} \rrbracket &= R \text{ says } \llbracket \mathbf{req} \rrbracket \end{aligned}$$

Figure 5: A Trust Transformation

3. All requests are accompanied by role certificates, to establish the relevant individual's role membership.

Characteristics (1) and (2) together determine the form of extrapolated statements specified in Fig. 4. Characteristic (3) implies that extrapolated authorization contexts will include role certificates, along with requests, ACLs, and CCLs. Hence, extrapolated authorization contexts are statements of the form:

$$\mathbf{req} \wedge \mathcal{R} \wedge \mathcal{C} \wedge \mathcal{A}$$

EXAMPLE 4.1. (Continued) *Joe, a doctor, wishes to use a medical web service WS_M to diagnose a patient, which in turn invokes WS_{MDB} for the patient's medical history. To initialize the appropriate role memberships and authorizations, Joe's request, in extrapolated form, is:*

$$K_J \text{ says } D \text{ says } \mathbf{priv}_{MDB} \wedge K_D \text{ says } J \Rightarrow D$$

This means that Joe (J), speaking as a doctor, wants to access \mathbf{priv}_{MDB} , and Joe (J) establishes that he can take the role of doctor (D).

4.5 Trust Transformation

The extrapolated context described above uses encryption to determine authenticity of statements, but at significant cost. Messaging can become quite complex, even in our simplified model, motivating a more efficient online checking technique, wherein many elements of extrapolated checking are "trusted away". Thus, we make the following simplifications for more efficient online checking:

1. Individuals are trusted to make valid claims about role membership.
2. Authenticity of requests is assumed for any intermediary; for example, if the request $R_1 \text{ says } R_2 \text{ says } s$ is received, then we trust that R_1 truly said " $R_2 \text{ says } s$ " and R_2 truly said " s ".

In practice, these assumptions are clearly too simplistic, in that they allow any web service invoker to assume any role membership. Some justification for that claim should be provided, e.g. a role key signature on the "top-level" of the request, so that instead of $R_1 \text{ says } R_2 \text{ says } s$, the signed

request $K_{R_1} \text{ says } R_2 \text{ says } s$ would be communicated. However, for the purposes of this example we will set this issue aside.

We formalize these trust assumptions via the trust transformation defined in Fig. 5. Note that the transformation eliminates role certifications and private key signatures for individuals, leaving just trusted role statements. We assert that authorization for trusted contexts is decidable (and efficient), satisfying Condition 1:

LEMMA 1. *Let s be an extrapolated context; then for all \mathbf{priv} , validity of $\llbracket s \rrbracket \vdash \mathbf{priv}$ is decidable.*

The result follows thanks to the similarity of the current authorization scheme with stack inspection; a decision procedure is easily defined as a modification of that given in [18].

EXAMPLE 4.1. (Continued) *By the trust transformation rules in Fig. 5, we have:*

$$\begin{aligned} \llbracket K_J \text{ says } D \text{ says } \mathbf{priv}_{MDB} \wedge K_D \text{ says } J \Rightarrow D \rrbracket &= \\ \llbracket K_J \text{ says } D \text{ says } \mathbf{priv}_{MDB} \rrbracket &= \\ D \text{ says } \mathbf{priv}_{MDB} & \end{aligned}$$

That is, Joe will simply send $D \text{ says } \mathbf{priv}_{MDB}$ to WS_{MDB} . Upon receiving Joe's request, WS_M composes the extrapolated statement $s_1 \wedge s_2$, where:

$$\begin{aligned} s_1 &\triangleq K_{WS_M} \text{ says } M \text{ says } D \text{ says } \mathbf{priv}_{MDB} \\ s_2 &\triangleq K_M \text{ says } WS_M \Rightarrow M \end{aligned}$$

Using the same trust transformation, we have:

$$\llbracket s \rrbracket = \llbracket s_1 \rrbracket = M \text{ says } D \text{ says } \mathbf{priv}_{MDB}$$

Hence, WS_M will simply send:

$$M \text{ says } D \text{ says } \mathbf{priv}_{MDB}$$

to WS_{MDB} .

4.6 Logging and Auditing

For auditing, it is necessary to reconstruct the signed statements and role certificates of extrapolated forms. In general, we imagine that auditing will be driven by the logging of signed statements and certificates in the online stage; although they're not used in online checking, they're saved and retrieved for offline verification. The details constitute a significant engineering problem and are beyond the scope of this paper.

Firstly, assuming that signed statements and signatures are logged, *who* exactly does the logging is another question. For example, everyone could be required to log their own statements, or the request statements they receive, or the source or target machine could be required to log all relevant statements, or perhaps a distinguished machine could

$$\begin{aligned}
\text{aud}(R \text{ says } \mathbf{priv}, \iota) &= \text{let } \hat{s} \wedge \mathcal{R} = \text{lookup}(\iota) \text{ in} \\
&\text{assert}(\mathcal{R} = K_R \text{ says } J \Rightarrow R) \text{ for} \\
&\text{assert}(\hat{s} = K_J \text{ says } R \text{ says } \mathbf{priv}) \text{ for} \\
&\hat{s} \wedge \mathcal{R} \\
\text{aud}(R \text{ says } s, \iota) &= \text{let } \hat{s} \wedge \mathcal{R} = \text{lookup}(\iota) \text{ in} \\
&\text{assert}(\mathcal{R} = K_R \text{ says } J \Rightarrow R) \text{ for} \\
&\text{assert}(\hat{s} = K_J \text{ says } R \text{ says } \iota') \text{ for} \\
&\text{let } s' \wedge \mathcal{R}' = \text{aud}(s, \iota') \text{ in} \\
&(K_J \text{ says } R \text{ says } s') \wedge (\mathcal{R} \wedge \mathcal{R}')
\end{aligned}$$

Figure 6: An Auditing Technique

be established as a log server, etc. We propose a model wherein these details are abstract; we posit *log locations* ι , and a lookup function that, given a location ι , returns the statements and certificates at that location. The concrete form and definition of locations ι and lookup, as well as logging conventions, determine the particulars for a given implementation.

Furthermore, noting that requests are single statements involving possibly multiple intermediaries, to obtain maximal flexibility in logging we would like the ability to “break up” statements into those parts made by each intermediary, in case each is independently responsible for logging. Thus, rather than statements such as $K_{WS} \text{ says } R \text{ says } s$, we introduce *indexed* statements as defined in Fig. 4, allowing auditing to “follow the trail” to reconstruct extrapolated statements from multiple log locations.

Thus, given these definitions, we impose the following discipline:

1. If a user J wishes to make a request for a privilege \mathbf{priv} in role R when invoking web service WS , the statement $R \text{ says } \mathbf{priv}$ will be sent to WS , and the statement:

$$(K_J \text{ says } R \text{ says } \mathbf{priv}) \wedge (K_R \text{ says } J \Rightarrow R)$$

will be logged at a fresh location ι ; this location will be determined by or communicated to WS .

2. If a web service WS wishes to propagate a request s to a web service WS' in role R , and the log location associated with s is ι , then WS sends the statement $R \text{ says } s$ to WS' , and the indexed statement:

$$(K_{WS} \text{ says } R \text{ says } \iota) \wedge (K_R \text{ says } WS \Rightarrow R)$$

will be logged at a fresh location ι' ; this location will be determined by or communicated to WS' .

If a web service WS receives a request s , with associated log location ι , online checking will be done with respect s , which is in trusted form. Auditing will use s together with location ι to reconstruct the extrapolated form of s ; in particular, given such s and ι , in a security context containing ACL \mathcal{A} and CCL \mathcal{C} , we define:

$$\text{audit}(s \wedge \mathcal{C} \wedge \mathcal{A}) = \text{aud}(s, \iota) \wedge \mathcal{C} \wedge \mathcal{A}$$

where aud is defined as in Fig. 6. In the definition of aud , the function $\text{assert}(P)$ blocks execution iff the predicate P is not valid. Therefore, aud reconstructs an authorized extrapolated statement from an authorized trusted contexts, and logged indexed statements, satisfying Condition 2; the result follows by induction on the trusted request:

LEMMA 2. *Suppose s is a trusted context, and $\text{audit}(s)$ returns s' ; then $\llbracket s' \rrbracket = s$.*

Furthermore, aud fails if it cannot reconstruct an extrapolated statement that is not authorized for the relevant privilege. Note that this property is implicit in the definition of aud ; it is not necessary to actually perform automated ABLP theorem proving on the extrapolated statement:

LEMMA 3. *Suppose s is a trusted context, $s \vdash \mathbf{priv}$ is valid, and $\text{audit}(s)$ returns s' ; then $s' \vdash \mathbf{priv}$ is valid.*

The key to this result is to note that $\text{audit}(s) \vdash s$, since aud only reconstructs valid role certifications and signed statements, via the assertions embedded in its definition.

EXAMPLE 4.1. (continued) *We assume that the transaction involving Joe, WS_M and WS_{MDB} adhere to the trust transformation and logging discipline described above. Let $s = M \text{ says } D \text{ says } \mathbf{priv}_{MDB}$. Now, it is clearly the case that:*

$$s \wedge \mathcal{C}_{MDB} \wedge \mathcal{A}_{MDB} \vdash \mathbf{priv}$$

so online checking succeeds. Furthermore, if WS_{MDB} audits this statement, it will obtain:

$$\begin{aligned}
&\text{audit}(s \wedge \mathcal{C}_{MDB} \wedge \mathcal{A}_{MDB}) \\
&= \\
&(K_{WS_M} \text{ says } M \text{ says } K_J \text{ says } D \text{ says } \mathbf{priv}_{MDB}) \\
&\wedge (K_D \text{ says } J \Rightarrow D) \wedge (K_M \text{ says } WS_M \Rightarrow M) \\
&\wedge \mathcal{C}_{MDB} \wedge \mathcal{A}_{MDB}
\end{aligned}$$

and we note that $\llbracket \text{audit}(s) \rrbracket = s \wedge \mathcal{C}_{MDB} \wedge \mathcal{A}_{MDB}$, and $\text{audit}(s) \vdash \mathbf{priv}$, as indicated by Lemma 2 and Lemma 3.

5. DISCUSSION

In this section we conclude with observations about related work, remarks on directions for future work, and a brief summary of the main points of the paper.

5.1 Related Work

Our system is a novel web services application of an idea that has been explored in other settings, namely previous authorization frameworks based on ABLP logic. Most closely related in this regard is proof carrying authorization (PCA) [4, 3], a framework for specifying and enforcing webpage access policies (though the logic used there is not ABLP, but an application specific variant). However, that system comprises a general framework for webpage access

control, so expressible policies are potentially more complicated than those we propose for web services. Furthermore, there is no distinction between online and offline checking in that framework as currently conceived, though our trust-but-verify approach could be adapted to it.

Other authorization systems founded in ABLP logic include that used in the Taos operating system [21], essentially a direct implementation of a subset of ABLP logic. Also, Wallach et al. have formalized the “security-passing style” of the Java stack inspection mechanism in a subset of ABLP [18, 20], which has served as a foundation for the SAFKASI programming language-based security architecture [19].

The SDSI/SPKI architecture [8, 16] is another authorization system for distributed communication. Their security model is similar to ABLP, but is based on a system of local names and emphasizes delegation. The semantics of SPKI/SDSI has been shown to be embeddable within ABLP plus an additional “speaks for regarding” primitive [12, 1], so the applicability of trust-but-verify to a SDSI/SPKI setting is suggested by our formulation. Delegation logic [14] and RT [15] are more recently proposed, logically well-founded authorization frameworks, based on datalog. No formulations of either SDSI/SPKI, delegation logic, or RT currently comprise trust-but-verify mechanisms.

In [17], a web services authorization system is defined, allowing specification of security policies in temporal logic, which are translated into reference monitors embedded in applications software. However, their approach is focused on complex policies for usage patterns similar to [3], and they make no online/offline checking phase distinction.

Related work on web services authentication includes an XML-based logic for web services authentication [6, 10], that is embedded within the applied Pi-calculus [9], allowing verification of web service security protocols via both human and machine proof techniques. It is likely that their approach will be relevant to future considerations of web services request authentication for our model.

5.2 Future Work

We envision a number of possibilities for future work. The most obvious direction is the development of trust-but-verify systems at a lower level, including an extension of SOAP messaging syntax to comprise ABLP formulae, and the detailed formulation of a system architecture for implementations. On the theoretical front, we also plan to embed our authorization framework within a richer formalism, such as that proposed in [10], allowing the semantics of authorization to be expressed and verified in a realistic threat model.

In addition to stand alone web services, authorization for composite web services [7, 13] also promises to be an interesting topic for future work; access control in that setting presents a number of unique issues for investigation.

5.3 Conclusion

In this paper, we have introduced a trust-but-verify framework for web services. We have used ABLP logic to establish a formal setting for framework design, and specified the conditions that any trust-but-verify implementation must satisfy. The central ideas we have presented are the separation of online and offline authorization phases, the notion of a trust transformation that establishes a meaningful relation between these phases, and a characterization of auditing for offline verification of online checking. We have presented

broad strokes of an example system, and in so doing have made some suggestions for implementations in general, including the use of role key certifications and carrier authority for flexible authorization schemes in an open web environment.

6. REFERENCES

- [1] M. Abadi. On SDSI's linked local name spaces. *Journal of Computer Security*, 6(1-2):3–21, 1998.
- [2] M. Abadi, M. Burrows, B. Lampson, and G. Plotkin. A calculus for access control in distributed systems. *ACM Trans. Prog. Lang. Syst.*, 15(4):706–734, 1993.
- [3] A. W. Appel and E. W. Felten. Proof-carrying authentication. In G. Tsodik, editor, *Proceedings of the 6th Conference on Computer and Communications Security*, Singapore, Nov. 1999. ACM Press.
- [4] L. Bauer. *Access Control for the Web via Proof-carrying Authorization*. PhD thesis, Princeton University, 2003.
- [5] L. Bauer, A. W. Appel, and E. W. Felten. Mechanisms for secure modular programming in java. Technical Report TR-603-99, Princeton University, Computer Science Department, July 1999.
- [6] K. Bhargavan, C. Fournet, and A. D. Gordon. A semantics for web services authentication. In *Proceedings of the 31st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 198–209. ACM Press, 2004.
- [7] V. Christophides, R. Hull, G. Karvounarakis, A. Kumar, G. Tong, and M. Xiong. Beyond discrete e-services: Composing session-oriented services in telecommunications. In *Proceedings of the Workshop on Technologies for E-Services (TES)*, Rome, Italy, 2001.
- [8] C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylonen. SPKI certificate theory. RFC 2693, Sept. 1999.
- [9] C. Fournet and M. Abadi. Hiding names: Private authentication in the pi calculus. In *Proceedings of the International Symposium on Software Security*, number 2609 in LNCS, pages 317–338. Springer-Verlag, November 2003.
- [10] A. Gordon, K. Bhargavan, C. Fournet, and R. Pucella. Tulufale: A security tool for web services. In Springer, editor, *Formal Methods for Components and Objects*, LNCS, 2003.
- [11] A. D. Gordon and R. Pucella. Validating a web service security abstraction by typing. In *Proceedings of the 2002 ACM workshop on XML security*, pages 18–29. ACM Press, 2002.
- [12] J. Howell and D. Kotz. A formal semantics for SPKI. Technical Report 2000-363, Dartmouth College, 2000.
- [13] R. Hull and J. Su. Tools for design of composite web services. In *ACM SIGMOD*, pages 958–961, 2004.
- [14] N. Li, B. N. Grosz, and J. Feigenbaum. Delegation Logic: A logic-based approach to distributed authorization. *ACM Transaction on Information and System Security (TISSEC)*, Feb. 2003. To appear.
- [15] N. Li, J. C. Mitchell, and W. H. Winsborough. Design of a role-based trust management framework. In *Proceedings of the 2002 IEEE Symposium on Security*

- and Privacy*, pages 114–130. IEEE Computer Society Press, May 2002.
- [16] R. Rivest and B. Lampson. SDSI — a simple distributed security infrastructure, 1996.
 - [17] E. G. Sirer and K. Wang. An access control language for web services. In *Proceedings of the seventh ACM symposium on Access control models and technologies*, pages 23–30. ACM Press, 2002.
 - [18] D. S. Wallach. *A New Approach to Mobile Code Security*. PhD thesis, Princeton University, 1999.
 - [19] D. S. Wallach, A. W. Appel, and E. W. Felten. SAFKASI: a security mechanism for language-based systems. *ACM Trans. Softw. Eng. Methodol.*, 9(4):341–378, 2000.
 - [20] D. S. Wallach and E. W. Felten. Understanding java stack inspection. In *Proceedings of 1998 IEEE Symposium on Security and Privacy*, Oakland, CA, May 1998.
 - [21] E. Wobber, M. Abadi, M. Burrows, and B. Lampson. Authentication in the Taos operating system. Technical Report 117, DEC Systems Research Center, 130 Lytton Avenue, Palo Alto, Ca 94301, December 1993.