NAMES IN CRYPTOGRAPHIC PROTOCOLS

Simone Lupetti, Feike W. Dillema, Tage Stabell-Kulø
Department of computer Science, University of Tromsø
University of Tromsø, N-9037 Tromsø, Norway
Email: {simone, task}@cs.uit.no, feico@pasta.cs.uit.no

Keywords: Names, authentication, cryptographic protocols, security.

Abstract: Messages in cryptographic protocols are made up of a small set of elements; keys, nonces, timestamps, and names, amongst others. These elements must possess specific properties to be useful for their intended purpose. Some of these properties are prescribed as part of the protocol specification, while others are assumed to be inherited from the execution environment. We focus on this latter category by analyzing the security properties of names. We argue that to fulfill their role in cryptographic protocols, names must be unique across correlated sessions i.e. where the massages of one session can be reused in another without detection and that uniqueness must be guaranteed to hold for each participant of these runs. We discuss how uniqueness can be provided and verified by the interested parties. To do so, two different mechanisms are shown possible, namely local and global verification. In both cases we discuss the implications of uniqueness on the execution environment of a cryptographic protocol, pointing out the inescapable issues related to each of the two mechanisms. Finally, we argue that such implications should be given careful consideration as they represent important elements in the evaluation of a cryptographic protocol itself.

1 INTRODUCTION

Cryptographic protocols are surprisingly subtle and so are their interactions with the environments in which they are deployed. Because of this, they should be specified so that the content of all messages is clear and the purpose of the different message elements is stated explicitly, together with the assumptions underlying each of them (Anderson and Needham, 1995). The number of messages and message elements that are relevant for the security of a protocol is normally rather small. Message elements encountered are keys, nonces, timestamps, identifiers and, occasionally, a few others.

Assumptions must be made regarding the properties of these elements, the participants of the protocol and their trustworthiness, their abilities, the properties of the communication infrastructure, and the environment in which the protocol is deployed. For example, cryptographic keys are required to be hard to guess and derive, nonces must be used for the first time, etc.

An assumption that is not maintained represents a potential security threat in that it may be exploited to attack the system. But, even if investigating and specifying the assumptions underlying security protocols is an important activity in the field of computer security, some usually receive more attention than others. Assumptions directly related to the functioning of the protocol are almost always made explicit and discussed at large (Syverson, 1992), while assumptions on the running environment are, more often than not, left unspecified (Canetti et al., 2002).

One consequence of this is that the cost factors considered when evaluating a cryptographic protocol are usually limited to conceptual simplicity, encryption, length, and number of messages (Woo and Lam, 1994, p. 34). External factors such as requiring secure synchronized clocks, synchronous communication or broadcast communication channels are less often considered as integral part of the protocol security (Gong, 1992).

We focus here on the single issue of *names* by analyzing their security requirements. Abadi et al. raised the issue of names in cryptographic protocols by formulating a specific prudent engineering principle:

> *"If the identity of a principal is essential to the meaning of a message, it is prudent to mention the principal's name explicitly in the message"*
>
> (Abadi and Needham, 1996)

They demonstrate the use of this principle with examples showing how some flawed protocols can be fixed by carefully inserting names of principals in the exchanged messages. But, like other message elements, names have properties that must be upheld to justify their use. Using some well-known cryptographic protocols as a starting point, we state that to fulfill their role in cryptographic protocols, names must be unique across sessions of the same protocol that are *correlated* to each other, meaning that they are executed in a time window where replay attacks are possible. How much complexity is added to the system to uphold uniqueness of names is the topic of this paper. We also present two distinct methods to fulfill such a requirement and discuss their peculiarities. In both cases we show how uniqueness of names is a source of external costs and we point out the practical implications of this for the design and evaluation of cryptographic protocols.

The terminology used in the rest of this paper is described in the next section. In Sec. 3, we recall examples of protocol attacks that can be fixed by the proper inclusion of names. These examples are used in Sec. 4 to present two possible ways to provide verifiable unique names. In Sec. 5 we discuss the property of uniqueness and its implications on the name system. We conclude in Sec. 6.

2 TERMINOLOGY AND NOTATION

Precise terminology about names and naming systems can be found in (Needham, 1993; Hauzeur, 1986; Benford and Lee, 1993). Our terminology is based on theirs, but has been simplified to keep our presentation intuitive. In particular, we use the term *name* not only for human-readable identifiers, but as a synonym for *identifier*. Both a name and an identifier are taken to be "a string of bits or characters that refers to a resource" (Shoch, 1978). We use the term *namespace* to indicate the set of valid values both for names and identifiers. By *naming system* we mean all the required machinery (set of mechanisms) for creating, using, and revoking names together with enforcing their required properties. Finally, we use the term *principal* (Lampson et al., 1992) to denote the generic source of a request.

Regarding the notation adopted to describe cryptographic constructs we use $\{m\}_{K_{pq}}$ to indicate a message m encrypted by a symmetric key known by the principals p and q. Also, with $\{m\}_{K_p}$ we mean that the message m is encrypted for confidentiality with p's public key while with $\{m\}_{K_p^{-1}}$ we indicate that m is encrypted by p's private key for integrity/authenticity.

3 EXAMPLES

We use the same set of examples as (Abadi and Needham, 1996, Sec. 4) due to their variety. We give only a minimal description of these here and refer the reader to the original article for more details about the attacks and the fixes.

3.1 Example 1

The first example is a key exchange protocol using asymmetric encryption by Denning and Sacco (Denning and Sacco, 1981, p. 535):

Msg 1 $A \rightarrow B : A, B$
Msg 2 $B \rightarrow A : C_A, C_B$
Msg 3 $A \rightarrow B : C_A, C_B, \{\{K_{ab}, T_a\}_{K_a^{-1}}\}_{K_b}$

A and B represent the names of the two principals participating in the protocol: Alice and Bob. The certificates C_A and C_A bind these two principals to their respective public keys K_a and K_b. T_a is a timestamp and K_{ab} is the session key being exchanged.

The problem is that once Bob receives the third message, he can remove the outer encryption layer and re-encrypt the result for an outsider Charlie. In particular, Bob could send:

Msg $3'$ $B \rightarrow C : C_A, C_C, \{\{K_{ab}, T_a\}_{K_a^{-1}}\}_{K_c}$

At this point, and for the duration of the validity of the time stamp T_a, Bob can impersonate Alice to Charlie. The proposed fix is that the third message should contain at least the name of Bob in the inner encrypted section [1]:

Msg $3''$ $B \rightarrow C : C_A, C_B, \{\{B, K_{ab}, T_a\}_{K_a^{-1}}\}_{K_b}$

It is possible now for Charlie to understand that this message is not part of his run of the protocol but was destined for B instead.

3.2 Example 2

A similar flaw is present in the following authentication protocol using symmetric key encryption by Woo and Lam (Woo and Lam, 1996, pp. 42-43):

Msg 1 $A \rightarrow B : A$
Msg 2 $B \rightarrow A : N_b$
Msg 3 $A \rightarrow B : \{N_b\}_{K_{as}}$
Msg 4 $B \rightarrow S : \{A, \{N_b\}_{K_{as}}\}_{K_{bs}}$
Msg 5 $S \rightarrow B : \{N_b\}_{K_{bs}}$

Where K_{as} and K_{bs} are two keys that the server S shares with Alice and Bob, respectively and N_b is a

[1] In the original article both Alice's and Bob's names appear in the fix but it is also pointed out that Alice's name can be deducted from K_a^{-1}. We then omit it here.

nonce generated by Bob. The purpose of this protocol is to allow Bob to check whether Alice is on-line. This protocol is subject to an attack where a third party Charlie impersonates Alice to Bob. Briefly, if Charlie presents itself to Bob both as Charlie and Alice at the same time, Bob is tricked to believe that Alice is present while she is not. The proposed fix again involves the use of names. In this case it requires the last message to include Alice's name:

$$\textit{Msg } 5' \ S \rightarrow B : \{A, N_b\}_{K_{bs}}$$

3.3 Example 3

The last example protocol is an early version (Hickman, 1994) of SSL (Freier et al., 1996), in which Bob can impersonate Alice to a third party Charlie:

$$\textit{Msg 1 } A \rightarrow B : \{K_{ab}\}_{K_b}$$
$$\textit{Msg 2 } B \rightarrow A : \{N_b\}_{K_{ab}}$$
$$\textit{Msg 3 } A \rightarrow B : \{C_A, \{N_b\}_{K_a^{-1}}\}_{K_{ab}}$$

K_{ab} is the session key to be used between Alice and Bob, K_b is Bob's public key and N_b is a nonce. The proposed fix is to insert Bob's and Alice's names in the final message [2]:

$$\textit{Msg } 3' \ A \rightarrow B : \{C_A, \{B, N_b\}_{K_a^{-1}}\}_{K_{ab}}$$

3.4 Summary

The three examples show that names assume an essential role in security protocols and that they can be employed as defense against a variety of replay attacks. But while these fixes seem straightforward, some questions arise regarding their implications: What properties, if any, must be provided in order for a name to be used in the security protocol? Should all possible instances of a name enjoy the same properties or do they depend on the particular usage and goal of the name in the protocol? If the security of the protocol relies on some property of the names in use, how can such a property be upheld? Does this impose requirements on the environment in which the protocol is deployed and in particular on the naming system in use? We investigate these questions in the next section with the help of the previous examples.

4 NAMES

The primary role of names is that they let one reference a resource, allowing sharing. In the examples we have presented, however, names are used as a mean to identify a specific protocol run by identifying the protocol participants.

Practical and scalable naming systems designed with sharing in mind usually allow for short-term inconsistencies to achieve high availability and good performance. The Domain Name System (DNS) (Mockapetris, 1987a; Mockapetris, 1987b) is an example of such a naming system; domain name entries are replicated on 2 or more nameservers and can be cached freely on any other server and client for performance. The DNS only offers "eventual consistency" and allows for short-term inconsistencies among the replicas and cached copies. This makes the naming system scalable and responsive in the light of unreliable message delivery, failures, and network partitions. The DNS leaves it to its users to check whether the resource they are referred to is the correct one. Essentially, this means that the DNS is subject to spoofing attacks *by design*. In other words, naming systems designed for resource sharing like the DNS are unsuitable for use in security protocols not being designed for identification purposes.

The question is then what properties names must have to be usable in security protocols. Do we have to remove all potential for naming inconsistencies and sacrifice availability, performance (and thus scalability) for security? The following principle states what must be achieved:

> *When the security of a protocol relies on the identity of a principal, the uniqueness of principals' names across all correlated sessions of the same protocol is required.*

Two or more protocol runs are correlated when executed within a time window where messages of one can be reused in another without detection. This window is determinate by the semantic of the protocol and its messages of a protocol and it is not related to time. For first protocol for example, this window is bounded by (the semantics of) the timestamp T_a in the third message. In Example 2 and 3 instead, since the presented attacks require that the attacker interleaves the messages of its session with the ones of the session he intends to attack, this window is determined by the temporal length of the attacked session.

This principle has been distilled from the previous examples. From them it is straightforward to extrapolate the uniqueness of participant names as the necessary condition for the proposed fixes to hold. For example, with reference to the first protocol, if Bob's and Charlie's names B and C are identical or can be confused with each other (maybe just at bit-string level) the proposed fix does not help anymore to fend off the attack. The same is true also in both the second and the third example.

Uniqueness, on the other hand, is not needed outside the scope of the correlated instances of a specific

[2] As for Example 1 earlier, we omit Alice's name because this can be deduced form her signature.

protocol run since (assuming that messages belonging to different protocols are distinguishable as such) illegitimate messages belonging to another sessions of the same protocol are here detected by other means.

4.1 Guaranteeing uniqueness

In addition to the ability to generate unique names, it must be safe for principals of a security protocol to assume that this uniqueness is maintained even in light of malicious principals. In other words, it must be infeasible for a malicious principal to violate the uniqueness assumption as the basis for an attack. Since uniqueness of names is, in general, required only across different sessions of the same protocol, a principal has two possible ways for verifying this:

Local verification: One or more of the participants in a protocol run verify locally the uniqueness of the names used in that run.

Local verification means here that the uniqueness check only involves local state and messages belonging to the protocol run itself, not entailing additional communication with third parties. Correlated sessions using the same names are prohibited and serialized as a consequence. Notice that in this solution the integrity of the namespace does not need to be protected since a name is *verified* to be unique every time it is used.

Global verification: A trusted naming system can be relied upon to be resilient to tampering. Such an infrastructure is global and must be designed to guarantee the integrity and the consistency of the namespace in order to provide unique names.

In this case, an authority hands out unforgeable and tamper-proof names for all the principals in the system.

To sum up, the consistency of the namespace must be either enforced globally or verified locally. Local and global verification each have, as we will show, their own applicability range that is delimited both by the protocol in question and by environmental constraints.

4.2 Local verification and verifiers

Not all principals participating in a protocol run are an interested in (verifying) the uniqueness of names. In the previous examples (and arguably in the general case), the principal that is concerned about the uniqueness of names is the potential victim of the attack that is countered by the use of names. Obviously, this role is protocol-dependent. With global verification, all protocol participants are given the same guarantees regarding the properties of the names used and therefore it is irrelevant what principals require uniqueness. A variety of scenarios is possible instead when local verification is used. We discuss some of them using the example protocols we presented earlier.

4.2.1 Example 1

In a protocol run between Alice and Bob, Bob can replay a message belonging to this run in another protocol run that he executes with a third party Charlie. Using these messages, Bob can impersonate Alice to Charlie. To avoid such a replay attack, Bob's name is added to the message being replayed in the attack. This name must be verified to be unique in order to thwart the replay attack.

Charlie cannot derive from the messages he receives from Alice, if and with what other parties Alice is or has been executing correlated runs of the same protocol. Charlie has therefore no way to verify whether his and Bob's names are identical and in use in correlated protocol runs by Alice (hence creating the precondition for a replay attack). Direct verification of the names by Charlie is therefore not possible.

But Charlie could trust Alice to perform the name check on his behalf. Alice, when opening two correlated executions of the protocol with Bob and Charlie, is able to check whether the names B and C (to be) sent in the messages of the two protocol runs are the same or not:

Msg 1 $A \rightarrow B : A, B$
Msg 2 $B \rightarrow A : C_A, C_B$
Msg 3 $A \rightarrow B : C_A, C_B, \{\{B, K_{ab}, T_a\}_{K_a^{-1}}\}_{K_b}$

Msg $1'$ $A \rightarrow C : A, C$
Msg $2'$ $C \rightarrow A : C_A, C_C$
Msg $3'$ $A \rightarrow C : C_A, C_C, \{\{C, K_{ac}, T_a\}_{K_a^{-1}}\}_{K_c}$

If Alice detects such duplicate names used in correlated protocol runs, she can protect Charlie against replay attacks by aborting one or both of the protocol runs. Notice that Charlie already trusts Alice for other security-critical tasks: For picking a good K_{ab} and for not disclosing it to third parties, for example. This means Alice is already part of Charlie's Trusted Computing Base (TCB). It is then reasonable in this case for Charlie to also rely on Alice for protection against this replay attack. On the other hand, trust should not be extended automatically without examination of all the possible implications, at least not in the general case.

To sum up, local verification is possible in this protocol under the condition that the potential victim of the attack, Charlie, trusts the initiator of the protocol, Alice, to verify the uniqueness of the names of his counterparts.

4.2.2 Example 2

Here the designated victim of the attack is Bob. A third party Charlie can induce Bob to believe that Alice is on-line while she is not. Enriching the fifth message with principals' names makes this impossible. Local verification is possible in this protocol also: For the attack to be mounted Charlie initiates two concurrent sessions of the protocol with the victim Bob. In these sessions, Charlie calls himself by two different names; his own and the name of the principal that he is impersonating:

$$\begin{array}{l} \textit{Msg 1} \;\; C \rightarrow B : A \\ \textit{Msg } 1' \; C \rightarrow B : C \\ \quad\quad\quad \vdots \end{array}$$

Bob registers the two names, A and C, and uses them ad verbatim in the messages of the protocol he sends later. This allows him to verify whether the names used in the different sessions are the same or not. In other words, the designated victim of the replay attack on this protocol is able to verify itself whether the names used in correlated protocol sessions are unique

4.2.3 Example 3

Our third and final example gives us the opportunity to discuss a scenario where local verification by one or more protocol participants is not possible. We report the replay attack on the protocol at full because it was left as an exercise to the reader in the original paper:

$$\begin{array}{l} \textit{Msg 1} \;\; A \rightarrow B : \{K_{ab}\}_{K_b} \\ \textit{Msg } 1' \; B \rightarrow C : \{K_{ab}\}_{K_c} \\ \textit{Msg } 2' \; C \rightarrow B : \{N_b\}_{K_{ab}} \\ \textit{Msg 2} \;\; B \rightarrow A : \{N_b\}_{K_{ab}} \\ \textit{Msg 3} \;\; A \rightarrow B : \{C_A, \{N_b\}_{K_a^{-1}}\}_{K_{ab}} \\ \textit{Msg } 3' \; B \rightarrow C : \{C_A, \{N_b\}_{K_a^{-1}}\}_{K_{ab}} \end{array}$$

Here Bob can impersonate Alice to Charlie. As a consequence, all information sent by Charlie to Alice can be read by Bob. Adding the names of the legitimate participants of each protocol run to the third message makes it impossible for Bob to mount this attack on Charlie. But, as we have seen, names must be verified by potential victim, Charlie in this case. Charlie has no way, however, to verify the uniqueness of names from the messages he receives. Moreover, he cannot rely on his counterpart to perform such a check as was possible in the first example. After all, that counterpart is the principal mounting the attack. Hence, local verification is not possible for this type of protocol and global verification has to be used to guarantee uniqueness.

4.3 Global verification

The most reliable way to enforce the integrity and consistency of a namespace is to use digital identity certificates. These are issued and digitally signed by a central trusted agency, the certification authority (CA), on behalf of all the principals in the system binding each of them to its public key. All of this, along with mechanisms for verifying the validity and revoking certificates constitutes a Public Key Infrastructure (PKI). Identity certificates are a common element in many cryptographic protocols; most of them already require the use of a PKI (see Example 1 and 3). This can provide for a reliable and trusted naming system. When the integrity and the consistency of the namespace is provided using a PKI, one could then replace the names used in the protocol with identity certificates for the named principals. This means:

$$\textit{Msg } 3'' \; B \rightarrow C : C_A, C_B, \{\{C_B, K_{ab}, T_a\}_{K_a^{-1}}\}_{K_b}$$

for the first protocol, and:

$$\textit{Msg } 3'' \; A \rightarrow B : \{C_A, \{C_B, N_b\}_{K_a^{-1}}\}_{K_{ab}}$$

for the third one.

Alternatively, one could use Alice's and Bob's public keys K_a and K_b (or their hashes) in place of their names. This is a common practice in the implementation of cryptographic protocols, but we recommend also to replace references (in subscripts) to the name B with references to the public key K_b already in the protocol description. This clarifies the properties we require of the names. To this end, however, rewriting the protocol to use certificates for principals' names is even better than when simply public keys are used in their place. Using certificates makes explicit in the protocol description the necessity of an authority taking responsibility for the integrity of the namespace.

5 DISCUSSION

We discuss here the two previous strategies: the local one that is based on punctual verification of names uniqueness and the global one that relies on a global names space guaranteed to be consistent.

As opposite as for naming systems that are specifically designed for sharing of resources (like the DNS), a namespace whose main requirements are to be consistent and non tamperable must sacrifice performance, scalability and availability because of its dependencies to a PKI (Slagell and Bonilla, 2004; Stabell-Kulø and Lupetti, 2005). Moreover, by relying on a PKI, the system inherits all its open problems like privacy concerns (in setting where they are relevant) and revocation-related issues (Brand, 2000; Millen and Wrigh, 1999).

Still, there are several valid reasons to use a PKI to provide a globally consistent namespace. In cases where public key cryptography (and hence a PKI) is in use or already required (Example 1 and 3), such a solution comes at zero additional cost in that the costs for deploying the PKI have been already sustained. Also, as already shown, local verification is not always possible, forcing in these cases to resort on a global and consistent namespace.

On the other hand, local verification does not rely on distributed state and does not add dependencies between distributed parties. However, even when local verification is possible it has its own share of limitations and implementation concerns.

Local verification requires principals to keep track of all names used in their correlated protocol sessions. All concurrent executions of the same protocol are also correlated if session identifiers are not used. Threading is a popular way to implement concurrently running instances of a protocol, and the list of names in use would be a resource shared between all threads requiring synchronized access. This list does not only grow linearly with the number of concurrent protocol sessions, but it may create a bottleneck because it is shared between all threads. Notice also, that cache of known "good" names can not be used as would be possible when global verification is used because the uniqueness of names must be verified at every new instantiation of the protocol.

Last but certainly not least, local verification is dependent on what principal in the protocol represents the attacker and what principal represents the verifier. This makes local verification dependent on the protocol itself. An *a priori* analysis of the protocol must be performed then to assure the feasibility and the correctness of the verification process. Such an analysis can only be performed for known replay attacks. Unlike a global verification, local verification can not therefore be used as as a *prevention* mechanism but only as a remedy for known exploits. In contrast, global verification provides global integrity for names, such that these can be employed also as a precautionary measure for all protocols countering attacks that are both known and not known to the implementor of the protocol.

5.1 Local vs. global uniqueness

We have argued earlier that the peculiar, but not unusual, use of names in cryptographic protocols makes system-wide uniqueness not necessary in all cases. Uniqueness of names in security protocols is required only during well-defined time intervals, i.e. for the duration of a protocol session, and in a bounded space, i.e. across sessions of the same protocol that are correlated.

In practice, however, the scope of the uniqueness is often extended both in space and time for reasons sometimes not dependent from the protocol itself. Global uniqueness, while not strictly required by the cryptographic protocol, may be convenient because the same global naming system is already in use to identify all principals in the system. This is the argument that leads, for example, to the adoption of public keys (and as consequence of a PKI) in places of names.

Uniqueness of names in time becomes an issue also when the name system implements revocation and re-allocation of names. For example, in IPSEC, IP addresses are used to name communication endpoints (Kent and Atkinson, 1998). When IP addresses are dynamically assigned (using the DHCP protocol, for example), protocol affiliation of messages should be extended also to subsequent sessions. Otherwise, a message intended for Bob, hence containing his name B, could be used in a replay attack against Charlie once Bob's name (IP address) has been reassigned to Charlie. Uniqueness of names may have to be maintained therefore also over time, for sessions that are not correlated.

It is always possible, however, to identify different protocol runs using session identifiers. This removes the reliance on (the uniqueness of) the names of protocol participants but, while naming sessions explicitly is effective, the integrity of session identifiers must be provided as well, possibly implying a substantial modification of the original protocol. When names have direct security implications, recycling of names should be carefully scrutinized (if not completely avoided).

The alternative to name recycling is to use lingering names across subsequent sessions. Such names are usually referred to as *persistent* or *inescapable* (Ellison et al., 1999). To implement inescapable names, however, may be a non-trivial exercise for a systems designer. For example, the namespace used must be "big enough" to never generate colliding names for different principals. A design flaw or implementation bug allowing a wrap-around of the name space may have serious and direct security implications. Determining how big is "big enough" is complicated by the fact that this is not merely determined by the worst-case usage rate of names in the system, but also by the worst-case abuse rate. In other words, unless the rate of name consumption is bounded somehow, a determined attacker is able to exhaust a name space independently of how big it is (Douceur, 2002).

5.2 Other uses of names

In some circumstances, names are used in the protocol as precautionary measure or for achieving secondary security goals. This is the case, for example, of the Otway-Rees mutual authentication protocol (Otway and Rees, 1987):

Msg 1 $A \rightarrow B : M, A, B, \{N_a, M, A, B\}_{K_{as}}$
Msg 2 $B \rightarrow S : M, A, B, \{N_a, M, A, B\}_{K_{as}},$
$\{N_b, M, A, B\}_{K_{bs}}$
Msg 3 $S \rightarrow B : M, \{N_a, K_{ab}\}_{K_{as}}, \{N_b, K_{ab}\}_{K_{bs}}$
Msg 4 $B \rightarrow A : M, \{N_a, K_{ab}\}_{K_{as}}$

With the help of a trusted server S, Alice and Bob authenticate each other by exchanging a session key K_{ab}. M is a session identifier generated by Alice while N_a and N_b are nonces generated by Alice and Bob, respectively.

The first instances of A and B in Msg. 1 are not protected by encryption having therefore no role in the security of the protocol. This use of names allows Bob to easily lookup the keys he needs for decrypting the remainder of the message. Bob is able this way to quickly discard messages and protocol runs that are not targeted at him or that are initiated by principals he is not interested in. Names are used as optimization and as a seemingly inexpensive first line of defense against simple Denial of Service (DOS) attacks.

For this purpose, the names still must be used to discriminate between principals and, as such, must be unique. Guaranteeing the uniqueness of these names, however, may be too costly for achieving such a secondary.

The potential lack of uniqueness of names has consequences commensurate to the importance of their role. When names are used for optimizations, duplicate names may have less severe consequences than when they are directly needed for the correctness of the protocol. It may not be cost effective to guarantee the uniqueness of names in such cases.

5.3 Design alternatives

In the protocols we have discussed, names are used to fend off replay attacks. All of these attacks require a principal to run two or more correlated executions of the same protocol with different counterparts. A message belonging to one session can be maliciously used this way in another(and *vice-versa*). In these cases, the attack can be thwarted by serializing all executions of the protocol at the victim. Note that, according our definition of concurrency, to just execute a protocol after another may not be enough if the windows of vulnerability extends also after a run is finished (as for the Example 1). in the general case. to serialize means therefore to execute only one protocol during each of this periods.

A possible solution to fend off replay attacks is then to restrict the execution environment of the protocol and so forbid multiple correlated executions of the protocol by the same principals. The original versions of our example protocols (without the additional use of names) would be secure in such a setting. Of course, this kind of restriction on the execution environment may negatively affect the system's performance, e.g. its responsiveness and maximum throughput.

Some would argue that relying on the execution environment of the protocol for protection to replay attacks is inferior from adding protection to the cryptographic protocol itself by adding names. However, names to be used in a cryptographic protocol must be secured either by the local implementation or by an external infrastructure. This places both solutions on equal footing, or some might even argue that a solution that does not require a change in the protocol is even preferable. Solutions against replay attacks do not merely rely on the design of the cryptographic protocol, but always rely on external support. The key point is then what are the different externalities and the economic and technical cost associated to each solution.

6 CONCLUSIONS

Names used in authentication protocols must enjoy uniqueness across correlated sessions of the same protocol to be effective. This property is assumed to be provided in the execution environment of the protocol and must be verifiable by its participants. For this, a system can either use a global naming system that provides unforgeable names to every principal in the system or resort to local verification by the concerned parties only.

In the case of a trusted naming system, this infrastructure becomes part of the trusted computing base of each principal of the protocol. The most common solution for such a service is to rely on a PKI and identity certificates. This kind of infrastructures are complex, costly to deploy and manage, and not risk-free (Ellison and Schneier, 2000).

On the other hand, local verification of uniqueness of principal's names is not always possible. Even when it is, local verification is a design choice that depends on the protocol in consideration and does not allow a general "one-for-all" implementation. Reimplementing the solution every time increases the probability of design errors and implementation bugs. In addition, as opposed to the PKI solution, it is not possible to use local verification as a preventive measure but only as a remedy to known attacks.

We argue that each solution has different implications, both technical and economic, associated with it. This warrants close consideration of what strategy is the best for a specific system and protocol. In either case, these costs may overwhelm the ones traditionally associated with cryptographic protocols such as the costs of communication and encryption. An *a* *priori* estimation of these costs should be made at design time because these may dominate total costs in the end and be a crucial factor in the (technical and economic) success of the system.

Finally, we argued that a system designer should not disregard alternative solutions for protecting from replay attacks. After all, we showed that protection against replay attacks by adding names to the protocol is not a self-contained solution. Alternative protection schemes that rely completely on measures taken in the environment may then very well be the simplest and the most cost-effective.

7 ACKNOWLEDGMENTS

There are those that will be thanked.

REFERENCES

Abadi, M. and Needham, R. (1996). Prudent engineering practice for cryptographic protocols. *IEEE Transactions on Software Engineering*, 22(1):6–15.

Anderson, R. J. and Needham, R. M. (1995). Robustness principles for public key protocols. In *CRYPTO '95: Proceedings of the 15th Annual International Cryptology Conference on Advances in Cryptology*, pages 236–247, London, UK. Springer-Verlag.

Benford, S. and Lee, O.-K. (1993). Collaborative naming in distributed systems. *Distributes Systems Engineering*, 1:67–74.

Brand, S. (2000). *Rethinking Public Key Infrastructures and Digital Certificates - Building in Privacy*. The MIT Press.

Canetti, R., Meadows, C., and Syverson, P. (2002). Environmental requirements for authentication. In *Proceedings of the International Symposium on Software Security*, volume 2609 of *Lecture Notes in Computer Science*, pages 339–355. Springer-Verlag.

Denning, D. E. and Sacco, G. M. (1981). Timestamps in key distribution protocols. *Commun. ACM*, 24(8):533–536.

Douceur, J. (2002). The sybil attack. In *Proceedings of the IPTPS02 Workshop*.

Ellison, C., Frantz, B., Lampson, B., Rivest, R., B.Thomas, and Ylonen, T. (1999). SPKI certificate theory. RFC 2693, The Internet Society.

Ellison, C. and Schneier, B. (2000). Ten risks of PKI: What you're not being told about public key infrastructure. *Computer Security Journal*, 16(1):1–7.

Freier, A. O., Karlton, P., and Kocher, P. C. (1996). The SSL Protocol - Version 3.0. RFC, Netscape Communications Corp.

Gong, L. (1992). A security risk of depending on synchronized clocks. *Operating Systems Review*, 26(1):49–53.

Hauzeur, B. M. (1986). A model for naming, addressing and routing. *ACM Trans. Inf. Syst.*, 4(4):293–311.

Hickman, K. (1994). The SSL Protocol. RFC, Netscape Communications Corp.

Kent, S. and Atkinson, R. (1998). Security architecture for the internet protocol. RFC 2401, The Internet Society.

Lampson, B., Abadi, M., Burrows, M., and Wobber, E. (1992). Authentication in distributed systems: theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310.

Millen, J. and Wrigh, R. (1999). Certificate revocation the responsible way. In *Proceedings of Computer Security, Dependability, and Assurance: From Needs to Solutions (CSDA'98)*, pages 196–203. IEEE Computer Society.

Mockapetris, P. (1987a). Domain names - concept and facilities. RFC 1034, The Internet Society.

Mockapetris, P. (1987b). Domain names - implementation and specification. RFC 1035, The Internet Society.

Needham, R. M. (1993). Names. *Distributed systems (2nd Ed.)*, pages 315–327.

Otway, D. and Rees, O. (1987). Efficient and timely mutual authentication. *ACM Operating Systems Review*, 21(1):8–10.

Shoch, J. (1978). Inter-networking naming, addressing, and routing. In *Proceedings of IEEE Computer Conference, COMPCON (Washington, D.C.)*, pages 72–79. IEEE Association.

Slagell, A. J. and Bonilla, R. (2004). PKI scalability issues. Technical Report 040918, National Center for Advanced Secure System Research.

Stabell-Kulø, T. and Lupetti, S. (2005). Public-key cryptography and availability. In *Proceedings of SAFECOMP 2005*, pages 222–232.

Syverson, P. (1992). Knowledge, belief, and semantics in the analysis of cryptographic protocols. *Journal of Computer Security*, 1(3):317–334.

Woo, T. Y. C. and Lam, S. S. (1994). A lesson on authentication protocol design. *Operating Systems Review*, 28(3):24–37.

Woo, T. Y. C. and Lam, S. S. (1996). Authentication for distributed systems. In Stallings, W., editor, *Practical Cryptography for Data Internetworks*, volume 25. IEEE Computer Society.