Towards an Infrastructure for Authorization

Position Paper

Joan Feigenbaum
AT&T Labs – Research
180 Park Avenue, Room C203
Florham Park, NJ 07932-0971 USA
jf@research.att.com

Abstract

In recent years, there has been a great deal of debate about whether a large-scale "public-key infrastructure" is needed for electronic commerce and, if so, whether the technical difficulty of building and deploying such an infrastructure will impede the growth of electronic commerce. We argue here that much of the controversy is attributable to the fact that the term "public-key infrastructure" has not been clearly and correctly defined. We explain why the informal definition most often associated with the term, *i.e.*, that of a global mapping between users' identities and public keys, is not the right definition for electronic commerce and hence that whether such a mapping can and will be built and deployed with available resources is not an especially pressing question. Finally, we describe an alternative type of infrastructural development that we believe really *would* enable electronic commerce.

1 Introduction

It is our position that the debate over *whether* a large-scale public-key infrastructure is needed for electronic commerce (or "e-commerce") been unsatisfactory, because it got started before developers had a satisfactory answer to the question of *what* a "public-key infrastructure" for e-commerce should consist of. We believe that the statement that some infrastructural development is needed before e-commerce can use public-key cryptography to its full potential is not controversial. Unfortunately, the term "public-key infrastructure" is used very narrowly, and the narrow meaning it has acquired is not particularly well-suited to e-commerce.

It is widely assumed that a public-key infrastructure should provide a way to associate the names of people and businesses with their public keys. That is, the controlling metaphor is that of a phone book. Just as one can look up a name in a phone book, find the line containing the associated phone number, and know immediately how to proceed, one is supposed to be able to call upon the public-key infrastructure to provide a "certificate," extract from it the public key associated with a name, and know immediately how to proceed. For those who believe that such an infrastructure would be useful for e-commerce, the question becomes whether to build a hierarchical, X.509-style structure [1, 2], a PGP-style "web of trust" [9], or something in between. It is our belief that a system of certificates that bind names to keys is not in and of itself very useful for e-commerce but that there is an alternative, more useful meaning that can be given to the term "public-key infrastructure."

In what follows, we focus exclusively on the infrastructure needed to enable widespread use of digital signatures. There are, of course, other ways in which e-commerce applications can use public-key cryptography, but the question of what's needed to deal effectively with digitally signed requests leads naturally to our main point. Because of this focus on signatures, we will use the terms "signing key" and "verification key" interchangeably with "private key" and "public key," respectively.

The crux of our position is that "public-key infrastructure" should not be used primarily to enable *authentication*. Rather, it should be used to enable *authorization*. Signatures on

paper documents in the commercial world often contain both the name and the job title of the signer; a reader who takes action in accordance with a request in such a document does so because he infers from the job title (*not* from the name!) that the signer is authorized to cause the requested action. If a paper signature is not accompanied by a job title or some other explicit token of authority, the reader may be able to infer that the signer has the necessary authority because he knows the signer personally and hence knows what authority the signer has. The point is that, when deciding what action to take, the crucial question for the reader of a paper document is not "who signed this," even in the many situations in which he can answer that question. Rather, the crucial question is "is the signer authorized to do what he wants to do?" The same reasoning applies to e-commerce applications that must process digitally signed requests: A "public-key infrastructure" that enables such applications to decide who signed a request isn't immediately useful; rather, one needs an infrastructure that allows the verifier of a digital signature to decide whether the signer has the authority to do what he wants to do.

This sole purpose of this position paper is to stimulate discussion at a public-key infrastructure session of an e-commerce conference. In particular, it is not our goal here to put forth new results and proposals. All of the technical material alluded to here has been developed (at AT&T Labs and elsewhere) in previous work; please see, *e.g.*, [4, 6, 7, 8] and the references therein for technical development.

2 The Phone-book Metaphor and Why It is Flawed

It is worth examining the origin of the phone-book metaphor and the precise way in which it fails.

In fact, this flawed metaphor is as old as public-key cryptography itself. In their seminal paper on the topic, Diffie and Hellman [5] say that "The enciphering key E can be made public by placing it in a public directory along with the user's name and address." They also claim that all cryptographic problems can be divided into those of "privacy" and those of "authentication," and they propose the following use of public-key cryptography for authentication: "If user A wishes to send a message M to user B, he 'deciphers' it in his secret deciphering key and sends $D_A(M)$. When user B receives it, he can read it and be assured of its authenticity by 'enciphering' it with user A's public enciphering key E_A."

This image of a public directory mapping users' names and addresses to their public keys is what I refer to as the "phone-book metaphor." As the theory of public-key cryptography developed and digital signatures were (correctly) identified as a crucial enabler for e-commerce, the power and pervasiveness of the phone-book metaphor led people to assume that e-commerce applications would process signed requests by retrieving from a public directory one or more "certificates" that establish the binding between a user name and a signature-verification key, verifying the digital signature using this key, and, if verification succeeds, deciding whether to perform the requested action using knowledge of *who* signed the request. Later statements of public-key infrastructural requirements for e-commerce were heavily influenced by the phone-book metaphor and by this insufficiently examined assumption about how digital signatures would be used. For example, the 1994 MITRE report for NIST [3] says in its executive summary that "This study addresses the issues related to a Public Key Infrastructure (PKI), which will automatically manage public keys through the use of public key certificates. Each certificate certifies the association between a user's identity and his public key." The technical challenge of creating, distributing, and maintaining very large-scale directories of "public-key certificates" has commanded center stage in e-commerce infrastructural development, and insufficient attention has been paid to the question of whether certificates that associate users' identities with their public keys are actually what e-commerce needs.

What is wrong with the phone-book metaphor? Essentially, the problem is that phone numbers are not analogous to signature-verification keys. In the world of Plain Old Telephone Service (POTS), where traditional, white-pages phone books were perfected, a person or business lived or worked in a fixed building, the building contained one or more

telephones connected by physical wires to a universal phone network, and each telephone had a number. When one located the name of a person or business in a phone book, the phone number listed beside this name had an unambiguous meaning, and there was really only one thing it could be used for, *i.e.*, to make a call to that person or business.

The fact that users of digital signing keys are not analogous to people and businesses with landline POTS numbers has, in one sense, been widely recognized. A great deal of attention has been paid to the observations that signers do not necessarily have fixed locations or long-lived signing keys and, hence, that a database of public-key certificates is more dynamic and distributed than a phone book and hence more difficult to build, maintain, and use. However, insufficient attention has been paid to a more fundamental way in which signature-verification keys are not analogous to phone numbers: They do *not* have an unambiguous meaning, and there is *not* just one thing that an e-commerce application can do with a verified signature. The meaning of a verified signature depends on the application and its policies. A reliable mapping from a verification key to a user's name would not necessarily allow an e-commerce application to decide whether to fulfill the signed request. To be useful, a verification key has to be bound reliably to all information needed to authorize a requested action, and the meanings of "reliably" and "authorize" vary from application to application. Certificates binding names to keys do not necessarily authorize anything and hence do not fulfill the infrastructural needs of e-commerce *no matter how reliably a dynamic, distributed database of certificates can be implemented.*

3 Infrastructure for Authorization

We now give a (partial) list of the infrastructural needs that must be met before digital signatures can be processed by computers on the same scale that paper signatures are now processed by people. Recall that [4, 6, 7, 8] and the references therein present recent and ongoing technical work that addresses some (but not all) of these needs.

Expressive Credentials: Instead of "certificates" that bind users' names to verification keys, e-commerce needs a much richer and more flexible notion of "credentials" that bind verification keys to the full range of information needed to authorize the actions signed by the corresponding signature keys. Examples include the right to approve expenditures up to a certain amount, the right to use a certain system resource, the right to certify that a digital object has a certain property, and the right to delegate authority to another key. Authority can be given to a single key or to a set of keys, because some requested actions require only one appropriate signature and some require more than one. Designing, building, and deploying a sufficiently general credential system that is usable in a wide variety of applications is one of the most important infrastructural needs for e-commerce.

Expressive Policies: A rich and flexible language is needed for policies as well as for credentials. That is, an e-commerce application must be able to encode in its policy a description of the credentials that are needed to authorize the various actions that it may take. Important elements of commercial policies include *delegation* and *trust*. As in the paper world, complex transactions in the electronic world will involve credentials, information, and commitments from many parties, and the developers and users of a particular application will, in general, not be in a position to understand or control everything that's involved. They will need to trust domain experts, and they will need to express precise, conditional delegation of authority to such experts. For example, an application may expect a request to run a program to be signed by a "mobile-code safety" expert, and it may delegate to a professional society the authority to certify such experts. It would delegate to entirely different parties authority over purely financial transactions. E-commerce policies will also need ways to handle conflicts among trusted authorities.

Experience writing and using a wide variety of policies should help us develop the e-commerce equivalent of "standard contracts" that exist in the paper world. These would be a valuable infrastructural component.

Compliance Checking: An e-commerce application presented with a signed request for action and a set of credentials must

have a well-understood procedure for deciding whether the credentials prove that the request complies with local policy. Ideally, a compliance checker would not just say "no" whenever the credentials fail to authorize the request but would, when appropriate, inform the requester about which additional credentials are needed for authorization or about an alternative action that is authorized by the credentials provided. Formalizing the notion of "credentials' proving that a request complies with a policy" is challenging and subtle. E-commerce development would benefit greatly from widespread deployment of an application-independent, general-purpose notion of "proof of compliance" that is explained, formalized, proven correct, and implemented in a standard package, thus freeing developers of individual applications from the need to reinvent the wheel. Applications that use a standard compliance checker could be assured that the answer returned for any given input (*i.e.*, a request, a policy, and a set of credentials) depends only on the input and not on any implicit policy decisions (or bugs) in the design or implementation of the compliance checker. As policies and credentials become more diverse and complex, the issue of assuring correctness will become especially important, and modularity of function with a clean separation between the role of the application and the role of the compliance checker will make further development more manageable.

Credential Issuers: We have already pointed out that an infrastructure that enables applications to figure out who signed what does not meet the needs of e-commerce, because the name of a signer cannot necessarily be mapped to the authority or privileges of the signer. A closely related point is that the issuer of a credential, *i.e.*, the party that binds a verification key to some appropriately encoded authorization information, should have legitimate power to issue it. This means, among other things, that the issuer should be technically competent to use whatever hardware and software is needed to create credentials, that it should possess whatever domain expertise is needed to understand the authority it is conferring on a key and how that authority will be used, and that it should have the resources and the motivation to take the commercial risks associated with authorization of the resulting actions. Thus, another explanation of the inadequacy of traditional "certificates" is that the ability to vouch for the association of a person and a name (*i.e.*, the ability that would be needed to create a credential that binds that name to the person's verification key) is in general unrelated to the ability to vouch for the person's authority to take relevant commercial actions. For a suitably diverse and widespread credential base to develop, we will need diverse and creative credential issuers. These issuers could be stand-alone businesses or agencies, or they could be divisions of larger organizations; the crucial point is that they be in a legitimate position to confer authority on keys.

Credential Distribution and Retrieval: In both the traditional "PKI" vision and the alternative vision we're advocating here, there is a need for more work on credential distribution and retrieval. Credential issuers will have to market their services effectively, and large-scale applications may have to manage big, dynamic databases of credentials.

Social and Legal Evolution: Although this is outside the expertise of the author (and hence beyond the scope of this position paper), it is necessary to recognize the many "infrastructural" needs for e-commerce that are wholely nontechnical. In particular, liability and risk management will be as important in the electronic world as they are in the paper world. Businesses and consumers will have to learn what to trust in the electronic world before mechanisms for trust management can be used effectively. We hope that effective laws really do *evolve* along with the relevant technology and that "legislative solutions" that don't make technological sense are not imposed by politicians unwilling to wait for evolution.

4 Acknowledgments

It is a pleasure to thank all of the people with whom I've had extensive discussions about electronic commerce, most notably Matt Blaze, Carl Ellison, Barb Fox, Jack Lacy, Brian LaMacchia, Butler Lampson, Ninghui Li, Marianne Mueller, Ron Rivest, and Martin Strauss.

References

[1] "Information Technology – Open Sys-

tems Interconnection – The Directory Authentication Framework," Recommendation X.509, ISO/IEC 9594-8.

[2] International Telegraph and Telephone Consultative Committee (CCITT). "The Directory Authentication Framework, Recommendation X.509," 1993 update.

[3] S. Berkovits *et al.*, "Public-Key Infrastructure Study: Final Report," National Institute of Standards and Technology, April 1994.

[4] M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis, "The KeyNote Trust Management System," `http://www.cis.upenn.edu/~angelos/keynote.html`

[5] W. Diffie and M. Hellman, "New Directions in Cryptography," *IEEE Transactions on Information Theory*, IT-22 (1976), pp. 644–654.

[6] C. Ellison, "SPKI Certificate Documentation," `http://www.pobox.com/~cme/html/spki.html`

[7] J. Feigenbaum, "Overview of the AT&T Labs Trust Management Project: Position Paper," in *Proceedings of the 1998 Cambridge University Security Protocols International Workshop*, Springer, Berlin, to appear. Available in preprint form as AT&T Technical Report 98.10.1 `http://www.research.att.com/library/trs/TRs/98/98.10/98.10.1.body.ps`

[8] R. Rivest and B. Lampson, "Cryptography and Information Security Group Research Project: A Simple Distributed Security Infrastructure" `http://theory.lcs.mit.edu/~cis/sdsi.html`

[9] P. Zimmermann, *PGP User's Guide*, MIT Press, Cambridge, 1994.