Bruce Christianson, Bruno Crispo, Mark Lomas and Michael Roe, editors. *Security protocols: proceedings of the 5th international workshop*, vol. 1361 of *LNCS*, ÉNS, Paris, France, April 1997. Springer Verlag.

Secure Books: Protecting the Distribution of Knowledge

Ross J Anderson, Václav Matyáš Jr., Fabien A Petitcolas,
(Computer Laboratory, Cambridge CB2 3QG)
Iain E Buchan, Rudolf Hanka
(Medical Informatics Unit, Cambridge CB2 2SR)

University of Cambridge, UK
{rja14, vm206, fapp2, ieb21, rh10}@cam.ac.uk

Abstract. We undertook a project to secure the distribution of medical information using Wax. This is a proprietary hypertext-based system used for information such as treatment protocols, drug formularies, and teaching material. An initial attempt, using digital signatures (in line with a recent European standard) and certificates conforming to X.509 has thrown up a number of interesting problems with current approaches to public key infrastructures. While the X.509 philosophy may be suitable for many electronic commerce applications, signatures on which we may have to rely for many years — such as those on books and contracts — appear to require a different approach.

1 Introduction

Wax is a system for publishing electronic medical books containing information such as treatment protocols, drug formularies and government regulations to which healthcare professionals need frequent access in support of clinical decision-making. Its origins lie in an earlier project [5] and its primary goals were to provide a good enough clinician-computer interface to be used safely without much formal training, and to have the knowledge management structure needed to support local clinical practice.

Protection issues such as assuring the integrity of the information, the authenticity of its source and non-repudiation (in a broad sense) started to arise in late 1996. This resulted in a collaboration between two departments of Cambridge University: the Medical Informatics Unit (the developers of Wax) and the Computer Laboratory, which has a group interested in information security.

The protection priority is to ensure that Wax users can correctly identify the author and publisher of a book on which they rely for clinical decision making, both at the time and if need be in the event of a subsequent dispute. There is

no requirement for strong secrecy properties; some books should be restricted to registered medical professionals only, but this is by virtue of drug adverts that may only be directed at this audience. So there is no more need to encrypt Wax books than to have the 'British Medical Journal' delivered by armed courier.

There is also no requirement to maintain an audit trail of which doctor or other healthcare professional read which chapters of which book. In fact, the maintenance of such a record would probably be considered an intrusion of professional privacy.

Thus Wax provides an interesting case study for the security professional. Unlike many other systems, whose protection is a mix of secrecy, authenticity, integrity and availability properties, Wax's needs are almost exclusively focussed on authenticity and integrity. It is also conceptually simple.

However, when we tried to implement protection mechanisms based on the obvious standards (X.509 [13] for certificates and the recently adopted European medical standard on RSA with exponent 3 for signatures) we ran into unexpected and interesting problems. These raise serious and important questions about the suitability of public key infrastructures currently under construction for certifying the integrity of long-lived objects such as books and contracts, and the authenticity of their signatories.

Section 2 overviews the Wax system. The threat model and security policy are discussed in section 3. The concept of trusted books and granularity of protection are discussed in section 4, while section 5 describes how the prototype system works and section 6 the trust structure. The problems encountered and lessons learned are described in section 7 and 8.

2 Wax

The original Wax project aimed to facilitate the sharing of knowledge between primary and secondary healthcare. The original system presented a user, such as a general practitioner, with a number of protocols for the care of specific diseases that were written and kept up to date by leading specialists in the field. Supporting information such as details of drugs were added, and in the next phase of the project it is planned to add administrative information such as directives from the Department of Health and local Health Authorities. The Wax mechanism also supports locally authored books, which might include policies and procedures developed by a general practice for its own use.

The educational experience of health care professionals is based around hierarchically classified paper systems such as libraries of books. Thus the source information for Wax is arranged into a familiar hierarchical structure: sections – books – shelves – library, and a special book-centred ASCII hypertext system was developed to support this.

The idea is that by providing a unified, hypertext-based library that enables clinicians to get at the information they need quickly and intuitively, and to

update it when appropriate in a controlled manner. The system thus supports a clinical hypertext library appropriate to a healthcare provider such as a general practice or hospital. It also allows users to add their own notes to any topic of a Wax book (these notes are kept constant in a separate file to cope with book updates). A single book or the whole library can be searched for words or phrases, cross references between books can be made, and many books can be open at one time.

An example of the interface is shown in the following figure.

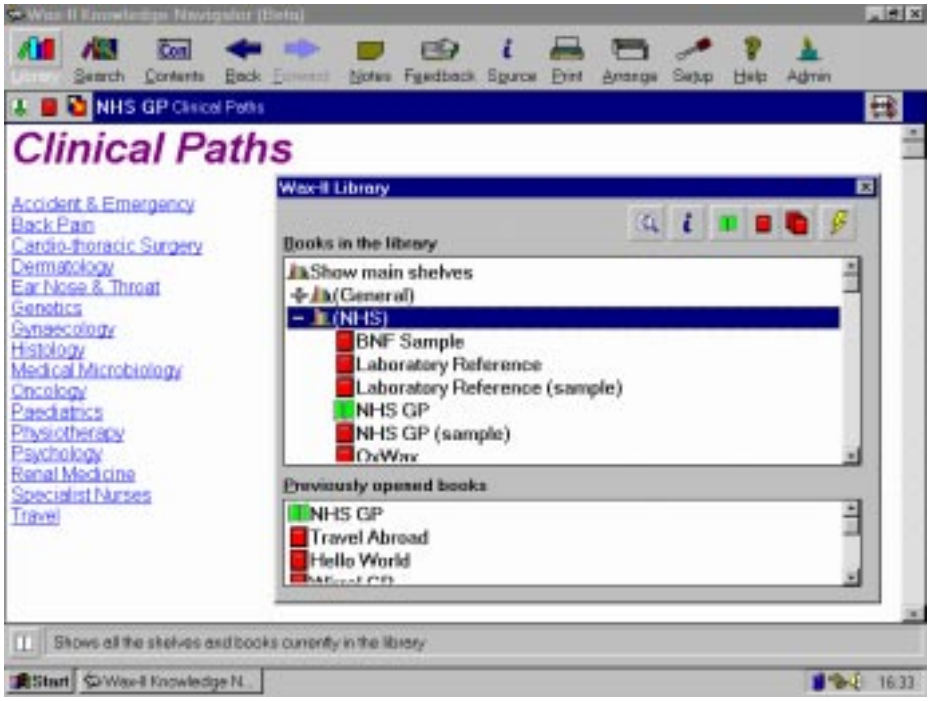

Fig. 1. Wax Interface

The Wax software supports both browsing and authoring and has been optimised for clinical use on Microsoft Windows compatible platforms. Wax books can import and export HTML documents and are designed to be browsed mainly from local storage, as performance is critical to clinical usefulness. The added facilities which are not achievable using HTML and standard web browsers include feedback mechanisms, navigation logs and performance features. Wax software and relevant papers [5, **?**] are at ¡`http://www.medinfo.cam.ac.uk/wax/`¿.

Wax book updates are done using diskettes or e-mail at present. However, from version 2 (due spring 1997) it is proposed to make updates available over the Internet by `ftp` or `http`. This could be particularly helpful in synchronising

book updates between district level libraries and end-user. However, it naturally raised the question of what cryptographic security features might be advisable.

3 Threats and Countermeasures

The overall level of threat against Wax is low — certainly much lower than against systems involving identifiable personal health information. The following primary threats were identified:

1. a book's content could be altered, whether by accident or malice;
2. an incorrect book source (author or publisher) might be claimed;
3. Wax software could be maliciously altered, whether by a general virus or by a more targeted attack;
4. a party involved in a dispute might deny the content of a previously published book, or challenge the date on which the information was published.

The first three are familiar from the general computer security environment [2]; concern about the fourth arises from a case in which a supplier of surgical implants sought to defeat litigation by claiming that it had published warnings at an earlier date than it in fact had done. Thus in addition to the integrity and authenticity of books and software, we want a non-repudiation service that covers both content and publication dates.

Some of the attackers might be large companies involved in litigation and thus able and willing to use professional techniques. However, these would probably be targeting the non-repudiation mechanisms, and the great majority of attackers — and almost all of those in the first two threat categories above — are likely to be insiders; health-care professionals with limited computer experience. They might intend to change the information within a book, perhaps to cover up malpractice [1].

Although the third threat does not introduce a serious risk in our application, it is the main concern with distribution and installation. The issue facing us is where to draw 'the fine line between prudence and paranoia' when it comes to trusted distribution. The main control envisaged is that the second version of Wax will be distributed in CD form by post. The secondary control is that when Wax is installed and a user is registered, a hash of the installed software is printed on the registration form. The third level control is that we will make available several means of checking the Wax distribution's integrity, including a PGP signature and hashes published in the medical press.

Given these checks, a reasonable level of trust can be placed in the Wax software, and the master public key that it contains. This key can be manually verified at any time, and users are requested to check it against a published value on installation.

At installation, users are also requested to create a public-private keypair with which to their own communications may be signed and verified. Key generation is very similar to the procedure in PGP, except that a hash of the public

key is printed on a registration form, together with the hash of the software mentioned above, and sent to the Wax-Centre by post. The public key is also sent to the Wax-Centre, preferably by email but if this is not possible by diskette or printed in hexadecimal on the registration form. After performing the appropriate due diligence checks (such as verifying medical registration), the Wax-Centre will send the user an identity certificate in the X.509 format.

Integrity checks are also performed whenever a book update is received, and whenever a book is opened (these will be described below, and are to a certain extent customisable by the user). The trust model is that the Wax-Root certifies the publisher, and the publisher certifies the book — taking responsibility for its content to the same extent as in the present world of paper (which is outside the scope of this paper).

The effect of the design is to reduce the problem of the trusted distribution of books to the trusted distribution of the Wax software and master public key. Under the circumstances, we consider that an appropriate level of effort has been expended on trusted distribution; any more effort than this would cross the line into paranoia.

4 Trusted and Untrusted Books

Each publisher certified by a Wax-Root certificate can publish books that the Wax system will consider to be trusted. At the present time, only Wax is a publisher, but other publishers can be added quickly, and in time most healthcare providers will act as publishers for their own local documentation. Users are permitted to have a small number of untrusted books open at any moment; these might be books that they are in the process of writing. However, once an author is finally satisfied with a book, a publisher can be asked to publish it, or the author can acquire the necessary key material for publishing. Once signed by a Wax-certified publisher, the book becomes trusted.

Each publisher is allocated a 'shelf' in the library, and users can clearly see the difference between trusted and untrusted books (different icons and colours). Users will also be warned if a book has been altered.

Users may determine how often books in their library are to be verified — ranging from when first downloaded to whenever opened. The reason for this is that delays must be minimised; the capability of user machines varies widely; and so does the threat environment in which the software is run, ranging from a single handed GP's notebook to a networked server in a large hospital with many temporary staff. So users need to select an appropriate frequency of checking.

Books can be verified in three different ways:

1. Each publisher maintains a catalogue that lists the currently available versions of all books, including not just their names but also their hash values. This catalogue is signed using the publisher's key, which in turn is certified using the Wax-Root key;

2. Each book is also signed by its publisher;
3. Version n of a book contains the hash of version $n-1$ (except for $n=1$).

It is intended to provide a further level of non-repudiation — in view of possible attack by funded organisations involved in litigation — by depositing CDs containing the current Wax library at the UK's statutory copyright deposit libraries or a similar body. (The whole issue of whether electronic media should be subject to the same statutory copyright deposit rules is currently a matter of government consultation in the UK [7].)

A topic that we have still not fully resolved is the granularity of protection. A user with a slow computer (or following a slow hypertext link to an online book chapter) might not want to wait until the whole book has been verified. So there is a case for hashing each chapter of a book and then putting the top level protection on a hash of these hashes. Whether this brings more practical benefits than problems is to be explored empirically.

5 Book Updates

Whatever the internal granularity of protection, external protection (by the catalogue and signature mechanisms) is implemented at the book level. Thus each book has a hash value associated with it, that is protected both by publication in a publisher's catalogue and by being signed.

In order to issue an update of one or more books, a publisher must therefore sign them and also create either a complete new catalogue or a supplementary catalogue, which he also signs. This catalogue, plus the books it refers to, are made available for download.

The user first gets the catalogue and checks its signature all the way back through the certificate chain to the Wax-Root key. Books to be downloaded are then chosen, and Wax instructed to fetch them. Once the books have been downloaded and their integrity checked, an update is made to an index kept locally and it is signed with the user's signing key. A passphrase is solicited to confirm that the user is happy with the new configuration of the library.

6 Trust Structure

As noted above, the primary purpose of the Wax system's security features is to reduce the problem of verifying the authenticity and integrity of downloaded medical and other books to that of verifying the authenticity and integrity of the Wax software and its embedded root key. Only those books whose certificate chain can be followed successfully back to the Wax-Root are considered trusted.

The trust structure thus looks like figure 2 overleaf.

In the layer below the Wax-Root will be found the main publishers, such as the Wax-Centre (the current publisher); in time this may include both official

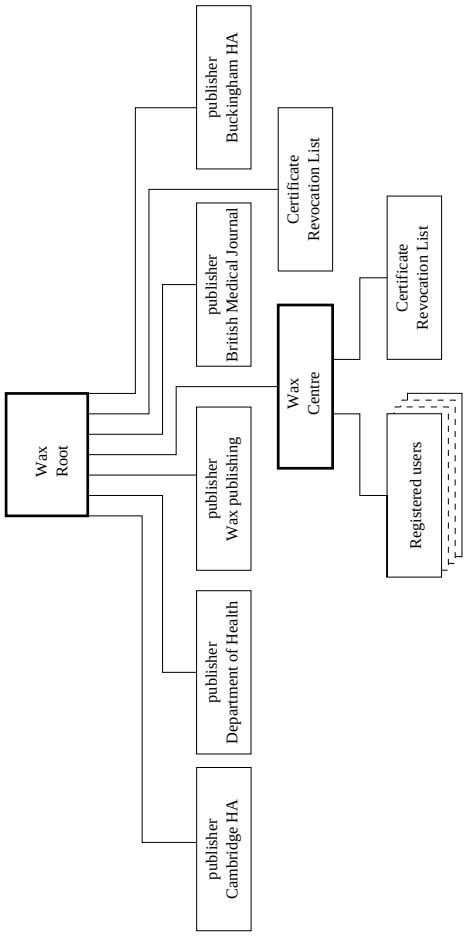

Fig. 2. Foreseen Certification Structure

publishers, such as the Department of Health, commercial publishers, and the publications of professional bodies such as the British Medical Association. As the number of publishers grows, it is envisaged that another one or two levels may be added. Thus, for example, the Department of Health may in time accept responsibility for certifying local health Authorities, who also publish; and the British Medical Association might certify individual general practitioners via the existing structure of Local Medical Committees. The future trust structure in medicine is a matter of current negotiation between interested parties, and nothing we are doing in the Wax project is an attempt to pre-empt or second guess the outcome of these discussions.

But for the time being, to get the system underway, the Wax-Centre will act as the Certification Authority for individual users.

7 The Shortcomings of X.509 Certificates

In the initial design we produced for this system, we assumed that individual users and publishers would have X.509 certificates [13] as these are in some sense the 'standard' format (although there are competing architectures, such as Microsoft's SDSI [10]). The idea was that each user, and each publisher, would have a certificate, and that a Certificate Revocation list would be published daily as a special Wax book. However, once we got into the detailed design phase, it became increasingly apparent that the trust model implicit in X.509 is not particularly appropriate in electronic publishing.

This observation may be the principal value of this paper to the broader research community, and to people involved in building other applications, so it is worth going into in some detail.

X.509 certificates are in many ways similar to credit cards — certificates have an expiry date (like credit cards) and there is a certificate revocation list that performs the same function as the hot card list of a credit card company. Unsurprisingly, the support they offer for general purpose digital signatures is less than ideal.

Many signatures must persist for a long time, such as the signatures on Wax books (assuming for the time being that signatures continue to be used as a protection mechanism). What does it mean, for example, that an author's certificate expires after a set period of time? Does his book suddenly become untrusted? And what does revocation mean in the context of the signature on a book that has already been published?

'Planned obsolescence' may make sense in the world of software publishing, while in the world of electronic commerce, it is perfectly reasonable that the keys used to sign credit card transactions should be replaced every two years — just as physical credit cards are. This limits the period of time in which a live key remains open to compromise; it also limits the size of the certificate revocation list.

However, if a medical publisher goes to the trouble of having a textbook on a medical condition or procedure prepared by an eminent clinician, or where a drug company publishes the results of a clinical trial on a new drug, then we may wish that the integrity and usability of that information be protected for many decades. (The copyright in the work, for example, will typically persist until 70 years after the author's death.)

We are therefore faced with the problem of how we can assure the authenticity and integrity of Wax books for a long period of time, while at the same time not exposing signing keys to compromise by requiring that an individual maintain a single signing key, and migrate it from one system to another, for the whole of his life. (This is not just a medical publishing problem: other applications where durable and secure signatures are needed include not just archived medical records, but all kinds of legal contracts.)

We believe that the long term solution to the problem of long term trust will involve a thorough re-examination of many of the assumptions that have grown up around public key cryptography. We believe that catalogue mechanisms will provide a better way of protecting many long lived objects such as books and contracts; where an object needs to be authenticated and published more quickly than the catalogue update cycle will allow, and strong non-repudiation is a requirement, then a one-time signature may be ideal. After all, so long as a signing key still exists, so does the possibility that a compromise might occur and lead to problems with the validity of the signature.

So the typical two-to-three year lifetime of an X.509 certificate has the curious property that it is too short for long-lived trust, and too long for strong non-repudiation!

Other problems with X.509 in the medical publishing context include firstly that X.509 certificates vouch for identities rather than roles, while we need support for roles: notices signed by 'The Chief Medical Officer' should be bound with the office rather than with the individual who is presently its holder; and secondly that X.509 does not support dual control. There are of course workarounds: dual authorship can always be noted by having a book signed separately by its two authors, and a two-out-of-three arrangement for (say) the Wax-Root key can be achieved using threshold signatures. However, it would be much preferable to have explicit protocol level support for such features.

8 Catalogue Based Trust

The effect of these discoveries on the Wax project has been to persuade us to make the publisher's catalogue — rather than the signature of the publisher or the author — the primary check on integrity and authenticity as new books are downloaded; and the user's internal, signed, directory to be the main check as existing books are opened. The publishers' signatures on the books themselves are merely a secondary check.

So the primary function of the author's signature is now to vouch for the authenticity of his own list of installed books and thus ensure that important files have not been tampered with by other users. This is not entirely straightforward; another user could have replaced the list with one of his own, whether by accident or malice, and as this would be signed by the other user, it is not sufficient merely to check the certificate chain back to Wax-Root. This reduces essentially to the trusted path problem; we propose to display the certificate owner's name in the title of the list of books as well as in the title of every open book. Given that the Wax software is adequately trusted, this gives us an adequate level of confidence.

The secondary function of the user signing key is to validate books sent to a publisher. This is relatively unimportant (except with very prolific authors) as the author could always read a hash over the phone instead.

As an interim measure, we have decided to set a short lifetime (1 year) for the publishers' signing keys.

In the long term, we believe that trust in published matter will not rely on the certification of individual identity, but on cataloguing, notarisation and timestamping. In this model, each page may contain its own hash in a header, but will certainly contain the hashes of all trusted pages to which it points as an extension of the URL (these can be checked using a Java applet or other convenient mechanism). A publisher's root page will contain the top-level hash for his catalogue; the root page hash will be checkable by some other means (such as notarisation in a paper journal).

In this model, the main use of digital signature mechanisms may be to authenticate pages that are created on demand (such as current exchange rates), in which case mechanisms such as RSA may be used, and to provide flexible links to pages that are updated frequently (such as secondary care access information) or at short notice (such as an official warning about a drug side-effect). In the latter case, we suspect that on thorough analysis a prudent designer may well use one-time signatures, or at the very least conventional signature keys that are destroyed rapidly but whose corresponding verification keys remain valid in perpetuity.

This architecture is intrinsically congenial to web-based publishing and it may therefore succeed in the market place. In addition, it follows the Rossnagel principle that electronic trust structures should reflect those in existing practice [11] — a principle jointly agreed by the BMA and the Department of Health. If it does succeed — for whichever of these reasons — then the pragmatic approach being implemented in Wax will be as compatible with it as one can expect to be with standards that are not yet written. Our approach does however make such prudent use as can be made of the existing X.509 architecture that many governments and other organisations are struggling to put into the field; and if X.509 does come to underpin electronic publishing, then we are compatible with that too.

9 Conclusion

We have developed a mechanism to assure the authenticity and integrity of electronic books. Although our particular application was medical, many of the lessons learnt apply to publishing in general, to digital contracts, and indeed to any application in which we need to assure the trustworthiness of digital objects over long periods of time.

The main lesson learned was that the trust structure embodied in X.509 and related standards is not suited for such applications. Indeed, it may turn out that digital signatures are not the appropriate tool for the job, but rather secure cataloguing and notarisation services based on trees of hash values.

We managed to achieve a pragmatic compromise that enables us to move ahead with an initial solution that is broadly compatible with both the public key and cataloguing approaches. However, as time passes and more functionality is added, we expect that ultimately one or other of these design options will have to be closed off. We feel that for both market and other reasons, the protection of published matter will come to rely on mainly cataloguing, with digital signature techniques used principally to provide flexible links from a catalogue-based trust structure to items whose content varies more quickly than the catalogue update cycle.

A catalogue-based trust structure may assuage the fears of law enforcement agencies over crypto proliferation. In many applications (such as conventional book publishing) there is no need for secrets at all (except as part of local mechanisms for logon, trusted path and the like); even where digital signatures are used to link to urgent notices, there is every incentive to use one-time signatures, or at the very least signatures whose verification keys have very much longer lifetimes than their signing keys. In any case, the apparent requirement for a central identity authentication service evaporates.

In conclusion, the long term protection of the authenticity and integrity of digital objects is far from being an adequately solved problem.

Acknowledgements

Václav Matyáš Jr. would like to thank the Royal Society for supporting his research through a Postdoctoral Fellowship.

References

1. "Nurse sacked for altering records after baby's death", K Alderson, *The Times*, November 95, p 6
2. "Why Cryptosystems Fail", RJ Anderson, in *Communications of the ACM* vol 37 no 11 (November 1994) pp 32–40
3. "Security in Clinical Information Systems", RJ Anderson, published by the British Medical Association, January 1996

4. "The Eternity Service", RJ Anderson, *Pragocrypt 96*, proceedings published by CTU Publishing House, Prague, ISBN 80-01-01502-5, pp 242–252
5. "Decision Support for Primary Care Using the Path.Finder System", Iain Buchan, Heather Heathfield, Tom Kennedy, Peter Bundred, in *British Journal of Healthcare Computing v 13 n 6*, pp 20-22, July 1996
6. "Exchanging Clinical Knowledge via Internet", IE Buchan, R Hanka, in *MEDNET 96*, proceedings to be published as a CD-ROM
7. "Government plans to save e-media for sake of nation", *Computer Weekly*, February 20th 1997 p 16
8. Good Medical Practice, General Medical Council, UK
9. "GP Practice computer security survey", RA Pitchford, S Kay, *Journal of Informatics in Primary Care*, September 95, pp 6–12
10. "SDSI – A Simple Distributed Security Infrastructure", RL Rivest, B Lampson, at ¡`http://theory.lcs.mit.edu/~rivest/publications.html`¿, presented at USENIX 96 and CRYPTO 96, April 30, 1996
11. "Institutionell-organisatorische Gestaltung informationstechnischer Sicherungsinfrostrukturen", A Roßnagel, in *Datenschutz und Datensicherung (5/95) pp 259–269*
12. "Secure Hash Standard", National Institute of Standards and Technology, *NIST FIPS PUB 180*, U.S. Department of Commerce, May 1993
13. "Information technology – Open Systems Interconnection – The directory: Authentication framework", *ITU-T Recommendation X.509*, November 1993