

A Single Path Towards Achieving Information Privacy and Record Authentication

Barry Hudson
Savannah River Site
Bldg 730-B
Aiken, SC 29808
barry.hudson@srs.gov
803/952-7080

Abstract

The Savannah River Site believes that a single user identity infrastructure can be created that is the sole source for information regarding user accounts, encryption keys, non-repudiation, authentication, authority, and access to protected data and resources.

This presentation will outline a 3-year plan that will establish a single, stable user identification mechanism and the consolidated deployment of crypto-based tools such as encryption, digital signature, and electronic records authentication.

Introduction

A number of technologies are competing to be the primary source of users' "digital identity". Until a few years ago, computer accounts were the only digital manifestations an employee possessed. Today, layered applications such as fileserver accounts, e-mail and Web/Firewall identification, and database access privileges require different levels of user identification, so no single recognition of user identity currently exists. Emerging applications such as electronic approval mechanisms, privacy-enhanced mail, encryption, and proof of identity outside of the enterprise (for complex-wide or commercial collaboration) will further complicate the ambiguities unless a single strategy is adopted. This strategy must encompass on-site authentication as well as recognition of and by off-site business partners and collaborators.

Digital identity information will be stored in a "certificate", which identifies the user, organization, roles, and rights. The certificate can be associated to any string of bytes, including an electronic record. The keys for user access to the certificate can also be used to encrypt stored or transmitted information. Thus, a single key infrastructure can provide both user and record authentication as well as privacy.

The Problem: Records Systems Have Unique Requirements

Information systems practices do not map to a records management environment. In most cases, data is protected through the use of an ID/password combination to a system or application. Protection is based on a bulk security approach, at an applications level, rather than at a granular level where data is protected based on its attributes. Because users sometimes handle passwords in a cavalier manner and are forced to use patterns or

familiar terms as mnemonics, the password has been devalued as an authentication mechanism. In business practice, policy or necessity of the moment dictates specific roles and privileges associated with position, allowing haphazard data access. Further, many systems provide access through a common username, such as “guest”. Consequently, the specific individual who provided data or made a change cannot be identified.

Obviously, these approaches are not sufficient for the protection of electronic records in accordance with the requirements and expectations of paper-based systems. Data standards coupled with certificate-based digital signature, implemented through the use of centrally managed keys is the strategy that SRS will use to address the issues of data retrievability, proof of identity, detection of alteration, and privacy.

The Challenges: Making Electronic Systems as Good as Paper Systems

Perhaps the biggest challenge in an electronic records system is whether it will be able to maintain vigilance for protection and retrievability of information far into the future.

Most paper records systems establish:

- Trust that a record is retrievable after 40 or more years
- The file retrieved has not been altered
- The originator of the record can be determined
- Assurance of privacy for information while it is stored.

Accomplishing those same four feats electronically is no small task.

The adoption of electronic records management systems often overlooks these characteristics that make a paper system reliable, trusted, and cost effective. Technology can facilitate electronic records storage and access, but cannot eliminate the need for these fundamental requirements. A system must be built to include these attributes; such capabilities cannot be “added later”. This requires advance planning and the proper use of emerging technology.

Trust that a record is retrievable after 40 or more years

Can the system be managed or migrated so that 40 year-old files are retrievable? First, one must address the longevity of the physical media on which the documents are stored. There are very few 10 year-old disk drives running today, and 40 year-old NASA tapes have proven to be unreadable. Optical and CD-ROM based storage is promising, but one cannot discount the chance of errors and shelf loss. A plausible approach is to store everything magnetically, and migrate it as systems are replaced.

Assuming the computer files themselves will be moved to, or be accessible to successive upgrades of hardware and software, the format may become obsolete, or the hardware to view it may be unavailable. What must be planned so that a Word 6.0 file can be read in 2038? To achieve the technology-independence of paper, systems must be designed, and

electronic formats selected that have integrity outside the records management system. All files should have stand-alone read capability (electronic legibility), and should be stored in a non-proprietary format, with software from many vendors capable of viewing it. TIFF images, HTML, ASCII text, and Acrobat PDF all have advantages and disadvantages.

Native formats such as Word are not viable due to the trend to change the internal representation as new releases of the editing software are released. Today, Word 6.0 will not open a 5 year-old Word 2.0 file without major reformatting and sometimes loss of information. Further discussion is provided in Appendix A and the draft policy is provided in the DOE Implementation Guide for Electronic Records Management (1).

The Solutions offered by Digital Signatures

The final three issues identified in the Challenges section above are the focus of the remainder of this discussion:

1. Is there assurance that the content (or format) was not altered?
2. Can the originator be positively determined?
3. Can the privacy of the data be maintained while it is stored or transmitted?

The 1991 decision by the Comptroller General, which defined the criteria for electronic signatures has been reinforced by the recent release of 21 CFR Part 11: Electronic Records; Electronic Signatures; Final Rule. In summary, they say that an electronic signature is valid if it is unique to the signer and under his sole control, and is capable of being verified. The signature must be linked to the data in such a way that if either the signature or the data is changed, the signature is invalidated upon verification. Further, electronic records can replace paper records when adequate integrity is maintained. Finally, digital signature technology, when properly implemented, can provide the necessary data integrity.

Keys and Certificates: The Foundation for Digital Identity

A digital certificate is a credential issued electronically by a trusted party or your own company. The rigor under which it is created and managed assures that it is unique to you and can be used only by you. When it is created, you must be present, both to prove your identity to the issuer and to receive the electronic keys needed to fully utilize the certificate in the conduct of electronic transactions.

The electronic keys are generated and issued in pairs. The keys are actually very large relatively prime numbers. They are mathematically similar, in that they share some properties. One key is held and known only to the employee. It is called the private key, and is used as positive identity to the electronic signature system. The private key can sign and lock information being sent by the holder, and unlock information encrypted for the holder's eyes only.

The second key is the public key. Having been generated at the same time and associated with the private key, it is capable of limited functions of the owner's private key, but is accessible to anyone in the system. Those limited functions include the ability to recreate portions of the owner's locked and signed information, to the extent it can be verified, but not forged.

The creation of keypairs, one of which is a public key, is called Public Key Infrastructure (PKI). Public keys can be formally exchanged between enterprises to allow for the verification of data across a broad user base. This exchange of public keys is called cross-certification.

The private and public keys are similar to the keys provided with some automobiles. The owner's key can lock and unlock the doors and trunk and start the car. The "valet" key can open the door and start the car, but cannot unlock the trunk. A holder of the public key with an owner's name can identify the owner and drive the car, but cannot unlock or lock items in the secure trunk.

Issuing a digital certificate: Electronic Credentials

This process somewhat resembles the opening of a checking account at a bank. The account owner proves his identity, signs a signature verification card, and obtains an account number. The bank keeps the signature card onfile for later verification of transactions processed under the account. Public knowledge of the account number does not necessarily compromise the money in the account. Funds can normally only be obtained when used in conjunction with the signature. The account owner retains the unique ability to create a signature identical to the one on the signature card.

To obtain a certificate, the employee presents his credentials to the Certificate Officer (CO). This verification of identity through the use of Social Security Card, Passport, or other recognized method is essential in establishing credibility of electronic identity. Personal presence is required for certificate creation and the binding of the keypair to the employee.

The CO generates a certificate and keypair for the employee. The CO digitally signs the certificate and public key, so the key's authenticity can be verified by anybody within the domain. The public key is stored on the key server along with the signed certificate. The certificate provides information regarding the person's identity, common name, roles and privileges in the organization, and other information.

The private key is encrypted using a password known only to the employee, and stored on a floppy or other "token" for use in the digital signature processing. The owner must keep the token under sole control, and never allow access to it by others. The Certificate Officer will manage the certificate and public key server. The CO also manages key exchange with outside entities, and may assert the right to keep a copy of the private key for audit or legal reasons. This so-called key escrow is a highly debated privacy issue, but

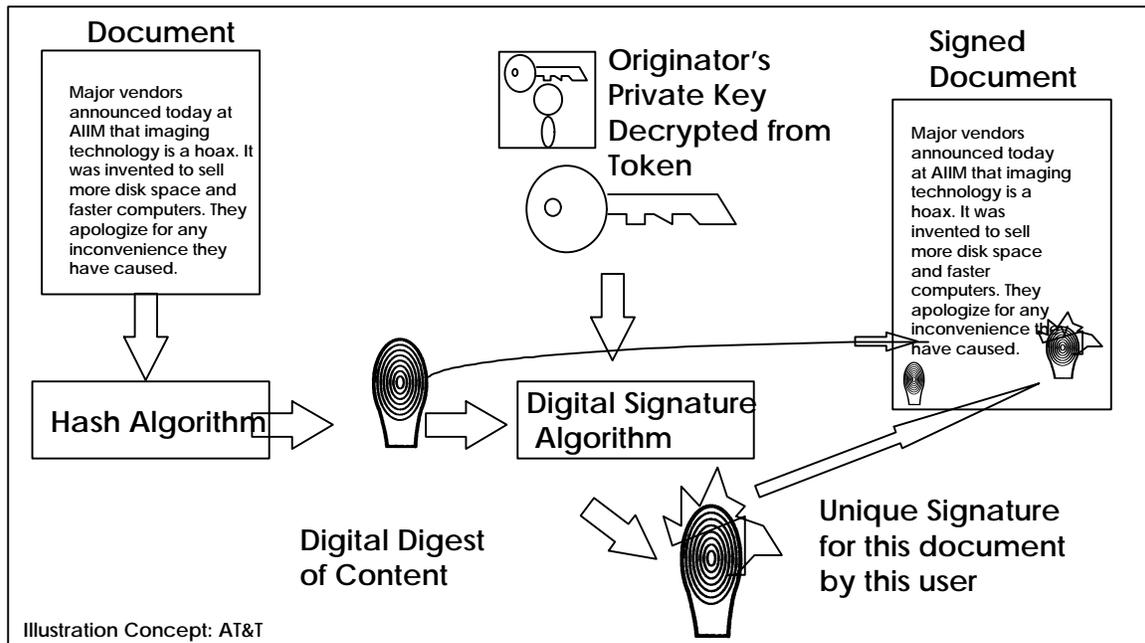
may be necessary to recover information if an employee is terminated or uses the keys maliciously.

Keys are issued with an expiration date, and can be revoked if the token (the private key) is lost or otherwise compromised, or if the employee leaves the company. Keypairs are periodically re-issued as a matter of policy and to prevent successful attacks on the encryption. The typical lifetime of a keypair is two years, though it may vary depending on the security consideration and risk from attempts to break the token encryption.

A few companies provide systems for the generation and management of keys (5,6). Many Federal agencies, especially the FDA (2) and US Postal Service plan to provide certificates and keys for inter-agency use. The Postal Service also plans to sell certificates for commercial and civilian use. Various proposals are being considered for the management of DOE-wide certificates and keys (4).

Signing a Document: Proof of Origination

To prove the originator of a document, the signature must be unique to the signer, linked to the data, and capable of being verified. The digital signature process follows these steps to meet those requirements.



For the document to be signed, a unique, compressed representation of the file is created. This so-called hashing process, creates the digest, a fixed-length string of bytes that can be created only from the document to be signed. It is mathematically unique so that if any bit is different, a different digest would be created. The digest is represented as the thumbprint in the illustration. The use of a digest simplifies the amount of computing needed in later steps, since it is smaller than the document.

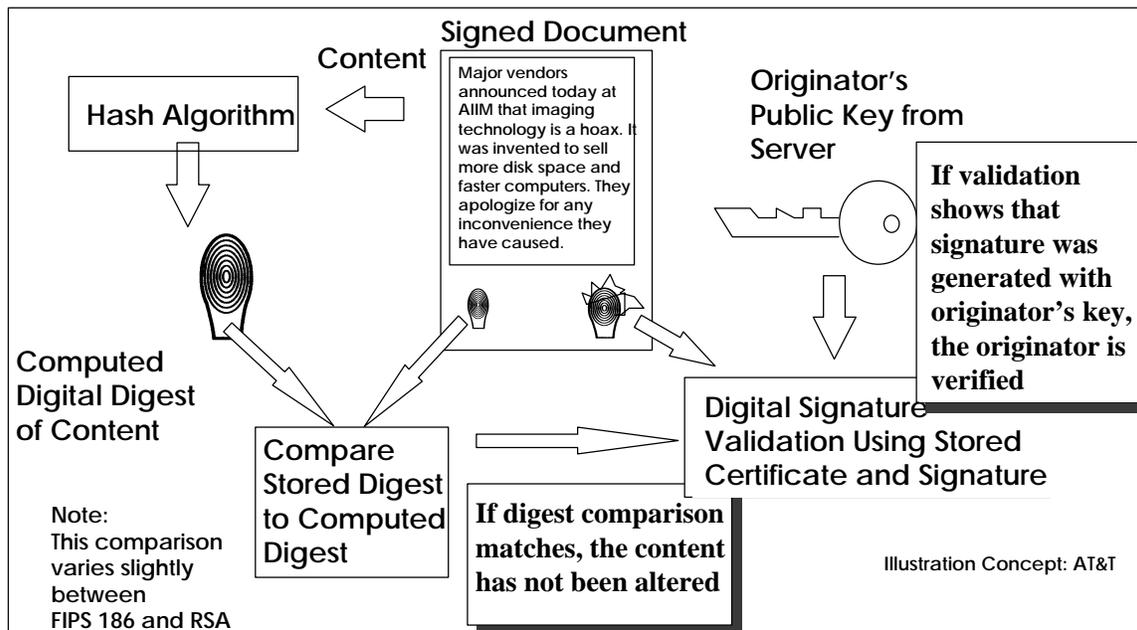
The originator applies his private key to the digest to create the digital signature for this document. This mathematical process assures that the resulting signature could have been created only from this document, and only with this originator's private key. Thus, the signature is linked to the data. The signature is represented by the modified thumbprint in the illustration.

In practice, the digest is created in the background by the document management system. The signer of a document would be asked to provide his token, or floppy with his encrypted private key, to the system. By entering a password (under his sole control), the private key is applied to the hash, creating the unique signature. The digest and signature can be stored separately or attached to the document. Depending on the required granularity of validation, only the signature need be stored.

Verifying a Document: Originator and Content Validation

The process for validating a document satisfies both the need for assurance that the content (or format) was not altered and provides positive determination of the originator.

A digest is created from the provided document original and compared to the provided digest. If the digests do not match, then the document has been altered. Since the hash algorithm creates a digest based only on the content (not the signer), the content can be verified by regeneration of the digest using the same hash algorithm. Comparison of hashes is quicker than comparing the lengthy documents.



To validate the originator of the document, a partial digital signature can be re-created using the public key of the purported originator against the new digest. While the public

key cannot create a valid signature (that would be a forgery), it is designed so that the partial signature can be compared to the actual signature to the extent that it can be verified or rejected.

If the validation shows that the signature was generated with the originators' key, then the originator and content is verified. If the partial signature were generated with a public key of somebody else, it would not contain enough similarities to validate. Further, if the partial signature is compared against a signature created from a different digest, it will not validate.

The comparison of digests assures that the document was not altered, and the comparison of signatures verifies the originator. Successful comparison only of signatures determines both that the content was not altered and that the originator was verified. Failure to verify indicates that either the content or originator are incorrect.

Additionally, since a validated signature could not have been created by anyone else, the originator cannot deny having created it. This property is called non-repudiation. By signing and exchanging small, randomly created pieces of information, users can positively identify each other, thus enabling a trusted electronic relationship for online transactions and information exchange.

It should be well noted, that digital signature cannot *prevent* the alteration of documents, but can only *detect* it. Physical security mechanisms such as non-erasable media and secure communications, when used in conjunction with digital signatures, provide a complete solution for deterring alteration and providing notification if the content is altered.

Encryption: Ensuring privacy of the while it is stored or transmitted

The final piece to the electronic records puzzle is the ability to protect information from those who do not have a need to know, and to prevent unauthorized access. Typical information management systems would attempt solve this through the use of resource or applications-based passwords, but would not be adequate.

Encryption is the reversible manipulation of data in order to prevent any but the intended recipient from reading that data. The same signature keys that allow validation of content and authentication can be used to encrypt documents for distribution to single or multiple recipients. A file can be signed only, encrypted only, or both signed and encrypted.

Encryption algorithms typically operate by using a large number as a "seed" to begin the randomization and encryption. If an employee's public key is used to encrypt a document, only that employee can decrypt that document, by using his private key. User-to-user privacy can be achieved by encryption using the recipient's public key, available from the certificate server.

Encryption to multiple users involves an additional, but clever step. The file is encrypted by software with a random, one-time key, known only to the software. This key is separately encrypted with each recipient's public key. The encrypted key and the encrypted file are sent to each recipient, who decrypts his key, and uses it to unlock the file. Observers who have no key cannot read the file. Files that require signature are typically signed before they are encrypted, so that the human-readable form is the one validated.

There are various applications for encryption that allow individual privacy of stored files. Encryption is used with the possible consequence that if the key is lost or forgotten, the data cannot be retrieved. Records managers should avoid storage of encrypted information, especially if there is no escrow of keys.

Conclusion

The establishment of a single, stable user identification mechanism will be required prior to deployment of crypto-based tools such as encryption and digital signature. Longevity of user/owner identity is important due to a requirement that files having cryptographic attributes will need to be verified as much as 40 years or more in the future against the same identity mechanisms used to encrypt or sign them. Creation and maintenance of X.509 certificates would be a giant step towards accomplishing all of the goals outlined herein. By its nature, certificates will also support e-mail directories, encryption, database access control, and some single-sign-on packages (SSO). An industry-accepted implementation would likely involve the use of physical devices, called "tokens to couple the use of passwords with personal presence at the moment of the transactions.

Industry standards and other external sources will drive decisions in this area. Vendor trends, NARA and The Association of Image and Information Management (AIIM) are particularly good sources for opinion in these issues (7). Further, each DOE site policy and practice will be subject to DOE-HQ policies in order to assure a consistent, minimum set of expectations and practices. Draft DOE policy for electronic records is provided in the DOE Implementation Guide for Electronic Records Management (1). Draft policy for PKI is provided in Chapter 9, DOE Telecommunications and Security Manual (3).

Complex-wide coordination will be required in order to assure eventual participation in complex-wide collaboration and privacy-enhanced mail. Each site's Public Key Infrastructure (PKI) will be cross-certified with others in the complex, building towards a complex-wide trust for secure e-mail, digital signature, and authentication.

References

1. DOE G 241.X-1, Electronic Records Management Guide for Use with 36 CFR Chapter XII - Part 1234. (draft March 11, 1998). For more information, contact Howard Landon, HR-41, Room 8F-084, FORRESTAL, howard.landon@hq.doe.gov. (202) 586-6344.
<http://www.explorer.doe.gov:1776/pdfs/draftordtext/241/g241x-1.pdf>
2. Food and Drug Administration, 21 CFR Part 11: Electronic Records; Electronic Signatures; Final Rule, April 1997
<http://www.access.gpo.gov/nara/cfr/index.html>
3. Chapter 9 (Draft), DOE Telecommunications and Security Manual, available from Clem.Boyleston@hq.doe.gov
4. DOE Working Group for Digital Signature. <http://www.hr.doe.gov/ucsp/ds.html>
5. Website for Entrust Technologies, Ltd. www.entrust.com
6. Website and FAQ for RSA Technologies. <http://www.rsa.com/rsalabs/newfaq/>
7. Website for AIIM, International. www.aiim.org

Author Bio

Barry Hudson is a Staff Technical Lead for Information Technology at the Savannah River Site in Aiken, SC. His current efforts concentrate on identification and pilots of business enablers, including groupware, crypto-based authentication, and electronic document management. He is active in a number of corporate, government, and national groups promoting the use of electronic documents, multimedia and CD-ROM in the business setting. He has made numerous presentations on document management and multimedia, including Seybold World, AIIM, CD-ROM Expo, and DOE's Inforum and AIMC conferences

Appendix A: Some Considerations for Electronic Records Systems

Many so-called electronic records systems track the location, retention, and data about paper records. Since such systems do not store the content, many considerations for electronic storage of the record itself have not yet been faced or resolved. In a totally electronic system, these issues must be decided:

1. Operated according to accepted practices

Is the system operated according to a documented set of procedures, limiting the personnel who have access to make changes to the system or data? Can the system store a wide variety of files, in formats widely accessible across the industry? If index data and content are both stored, are they linked sufficiently to each other? Is the index information considered to be “records”, or just the content?

2. Meets customer and legal requirements

What is the definition of an electronic record? Can the records acceptance process be conducted electronically? Must an electronic record have a paper counterpart? Is there a proper procedure for total removal of data after retention expires? Is there a security risk associated with having information searchable? Will the system be expandable to accommodate the volumes for the next decade? If the retrieved document has identical content, but different formatting, is it still a viable record?

3. Maintains vigilance for protection and retrievability of information

Can the system be migrated so that 40 year-old files are retrievable? Is there assurance that the content (or format) was not altered? Can the originator be positively determined?