

Network Working Group
Request for Comments: 1114

S. Kent
BBNCC
J. Linn
DEC
IAB Privacy Task Force
August 1989

Privacy Enhancement for Internet Electronic Mail:
Part II -- Certificate-Based Key Management

STATUS OF THIS MEMO

This RFC suggests a draft standard elective protocol for the Internet community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

ACKNOWLEDGMENT

This RFC is the outgrowth of a series of IAB Privacy Task Force meetings and of internal working papers distributed for those meetings. We would like to thank the members of the Privacy Task Force for their comments and contributions at the meetings which led to the preparation of this RFC: David Balenson, Curt Barker, Matt Bishop, Morrie Gasser, Russ Housley, Dan Nessel, Mike Padlipsky, Rob Shirey, and Steve Wilbur.

Table of Contents

1. Executive Summary	2
2. Overview of Approach	3
3. Architecture	4
3.1 Scope and Restrictions	4
3.2 Relation to X.509 Architecture	7
3.3 Entities' Roles and Responsibilities	7
3.3.1 Users and User Agents	8
3.3.2 Organizational Notaries	9
3.3.3 Certification Authorities	11
3.3.3.1 Interoperation Across Certification Hierarchy Boundaries	14
3.3.3.2 Certificate Revocation	15
3.4 Certificate Definition and Usage	17
3.4.1 Contents and Use	17
3.4.1.1 Version Number	18
3.4.1.2 Serial Number	18
3.4.1.3 Subject Name	18
3.4.1.4 Issuer Name	19
3.4.1.5 Validity Period	19
3.4.1.6 Subject Public Component	20

3.4.1.7 Certificate Signature	20
3.4.2 Validation Conventions	20
3.4.3 Relation with X.509 Certificate Specification	22
NOTES	24

1. Executive Summary

This is one of a series of RFCs defining privacy enhancement mechanisms for electronic mail transferred using Internet mail protocols. RFC-1113 (the successor to RFC 1040) prescribes protocol extensions and processing procedures for RFC-822 mail messages, given that suitable cryptographic keys are held by originators and recipients as a necessary precondition. RFC-1115 specifies algorithms for use in processing privacy-enhanced messages, as called for in RFC-1113. This RFC defines a supporting key management architecture and infrastructure, based on public-key certificate techniques, to provide keying information to message originators and recipients. A subsequent RFC, the fourth in this series, will provide detailed specifications, paper and electronic application forms, etc. for the key management infrastructure described herein.

The key management architecture described in this RFC is compatible with the authentication framework described in X.509. The major contributions of this RFC lie not in the specification of computer communication protocols or algorithms but rather in procedures and conventions for the key management infrastructure. This RFC incorporates numerous conventions to facilitate near term implementation. Some of these conventions may be superseded in time as the motivations for them no longer apply, e.g., when X.500 or similar directory servers become well established.

The RSA cryptographic algorithm, covered in the U.S. by patents administered through RSA Data Security, Inc. (hereafter abbreviated RSADSI) has been selected for use in this key management system. This algorithm has been selected because it provides all the necessary algorithmic facilities, is "time tested" and is relatively efficient to implement in either software or hardware. It is also the primary algorithm identified (at this time) for use in international standards where an asymmetric encryption algorithm is required. Protocol facilities (e.g., algorithm identifiers) exist to permit use of other asymmetric algorithms if, in the future, it becomes appropriate to employ a different algorithm for key management. However, the infrastructure described herein is specific to use of the RSA algorithm in many respects and thus might be different if the underlying algorithm were to change.

Current plans call for RSADSI to act in concert with subscriber organizations as a "certifying authority" in a fashion described

later in this RFC. RSADSI will offer a service in which it will sign a certificate which has been generated by a user and vouched for either by an organization or by a Notary Public. This service will carry a \$25 biennial fee which includes an associated license to use the RSA algorithm in conjunction with privacy protection of electronic mail. Users who do not come under the purview of the RSA patent, e.g., users affiliated with the U.S. government or users outside of the U.S., may make use of different certifying authorities and will not require a license from RSADSI. Procedures for interacting with these other certification authorities, maintenance and distribution of revoked certificate lists from such authorities, etc. are outside the scope of this RFC. However, techniques for validating certificates issued by other authorities are contained within the RFC to ensure interoperability across the resulting jurisdictional boundaries.

2. Overview of Approach

This RFC defines a key management architecture based on the use of public-key certificates, in support of the message encipherment and authentication procedures defined in RFC-1113. In the proposed architecture, a "certification authority" representing an organization applies a digital signature to a collection of data consisting of a user's public component, various information that serves to identify the user, and the identity of the organization whose signature is affixed. (Throughout this RFC we have adopted the terms "private component" and "public component" to refer to the quantities which are, respectively, kept secret and made publically available in asymmetric cryptosystems. This convention is adopted to avoid possible confusion arising from use of the term "secret key" to refer to either the former quantity or to a key in a symmetric cryptosystem.) This establishes a binding between these user credentials, the user's public component and the organization which vouches for this binding. The resulting signed, data item is called a certificate. The organization identified as the certifying authority for the certificate is the "issuer" of that certificate.

In signing the certificate, the certification authority vouches for the user's identification, especially as it relates to the user's affiliation with the organization. The digital signature is affixed on behalf of that organization and is in a form which can be recognized by all members of the privacy-enhanced electronic mail community. Once generated, certificates can be stored in directory servers, transmitted via unsecure message exchanges, or distributed via any other means that make certificates easily accessible to message originators, without regard for the security of the transmission medium.

Prior to sending an encrypted message, an originator must acquire a certificate for each recipient and must validate these certificates. Briefly, validation is performed by checking the digital signature in the certificate, using the public component of the issuer whose private component was used to sign the certificate. The issuer's public component is made available via some out of band means (described later) or is itself distributed in a certificate to which this validation procedure is applied recursively.

Once a certificate for a recipient is validated, the public component contained in the certificate is extracted and used to encrypt the data encryption key (DEK) that is used to encrypt the message itself.

The resulting encrypted DEK is incorporated into the X-Key-Info field of the message header. Upon receipt of an encrypted message, a recipient employs his secret component to decrypt this field, extracting the DEK, and then uses this DEK to decrypt the message.

In order to provide message integrity and data origin authentication, the originator generates a message integrity code (MIC), signs (encrypts) the MIC using the secret component of his public-key pair, and includes the resulting value in the message header in the X-MIC-Info field. The certificate of the originator is also included in the header in the X-Certificate field as described in RFC-1113, in order to facilitate validation in the absence of ubiquitous directory services. Upon receipt of a privacy enhanced message, a recipient validates the originator's certificate, extracts the public component from the certificate, and uses that value to recover (decrypt) the MIC. The recovered MIC is compared against the locally calculated MIC to verify the integrity and data origin authenticity of the message.

3. Architecture

3.1 Scope and Restrictions

The architecture described below is intended to provide a basis for managing public-key cryptosystem values in support of privacy enhanced electronic mail (see RFC-1113) in the Internet environment. The architecture describes procedures for ordering certificates from issuers, for generating and distributing certificates, and for "hot listing" of revoked certificates. Concurrent with the issuance of this RFC, RFC 1040 has been updated and reissued as RFC-1113 to describe the syntax and semantics of new or revised header fields used to transfer certificates, represent the DEK and MIC in this public-key context, and to segregate algorithm definitions into a separate RFC to facilitate the addition of other algorithms in the future. This RFC focuses on the management aspects of certificate-

based, public-key cryptography for privacy enhanced mail while RFC-1113 addresses representation and processing aspects of such mail, including changes required by this key management technology.

The proposed architecture imposes conventions for certification paths which are not strictly required by the X.509 recommendation nor by the technology itself. The decision to impose these conventions is based in part on constraints imposed by the status of the RSA cryptosystem within the U.S. as a patented algorithm, and in part on the need for an organization to assume operational responsibility for certificate management in the current (minimal) directory system infrastructure for electronic mail. Over time, we anticipate that some of these constraints, e.g., directory service availability, will change and the procedures specified in the RFC will be reviewed and modified as appropriate.

At this time, we propose a system in which user certificates represent the leaves in a shallow (usually two tier) certification hierarchy (tree). Organizations which act as issuers are represented by certificates higher in the tree. This convention minimizes the complexity of validating user certificates by limiting the length of "certification paths" and by making very explicit the relationship between a certificate issuer and a user. Note that only organizations may act as issuers in the proposed architecture; a user certificate may not appear in a certification path, except as the terminal node in the path. These conventions result in a certification hierarchy which is a compatible subset of that permitted under X.509, with respect to both syntax and semantics.

The RFC proposes that RSADSI act as a "co-issuer" of certificates on behalf of most organizations. This can be effected in a fashion which is "transparent" so that the organizations appear to be the issuers with regard to certificate formats and validation procedures. This is effected by having RSADSI generate and hold the secret components used to sign certificates on behalf of organizations. The motivation for RSADSI's role in certificate signing is twofold. First, it simplifies accounting controls in support of licensing, ensuring that RSADSI is paid for each certificate. Second, it contributes to the overall integrity of the system by establishing a uniform, high level of protection for the private-components used to sign certificates. If an organization were to sign certificates directly on behalf of its affiliated users, the organization would have to establish very stringent security and accounting mechanisms and enter into (elaborate) legal agreements with RSADSI in order to provide a comparable level of assurance. Requests by organizations to perform direct certificate signing will be considered on a case-by-case basis, but organizations are strongly urged to make use of the facilities proposed by this RFC.

Note that the risks associated with disclosure of an organization's secret component are different from those associated with disclosure of a user's secret component. The former component is used only to sign certificates, never to encrypt message traffic. Thus the exposure of an organization's secret component could result in the generation of forged certificates for users affiliated with that organization, but it would not affect privacy-enhanced messages which are protected using legitimate certificates. Also note that any certificates generated as a result of such a disclosure are readily traceable to the issuing authority which holds this component, e.g., RSADSI, due to the non-repudiation feature of the digital signature. The certificate registration and signing procedures established in this RFC would provide non-repudiable evidence of disclosure of an organization's secret component by RSADSI. Thus this RFC advocates use of RSADSI as a co-issuer for certificates until such time as technical security mechanisms are available to provide a similar, system-wide level of assurance for (distributed) certificate signing by organizations.

We identify two classes of exceptions to this certificate signing paradigm. First, the RSA algorithm is patented only within the U.S., and thus it is very likely that certificate signing by issuers will arise outside of the U.S., independent of RSADSI. Second, the research that led to the RSA algorithm was sponsored by the National Science Foundation, and thus the U.S. government retains royalty-free license rights to the algorithm. Thus the U.S. government may establish a certificate generation facilities for its affiliated users. A number of the procedures described in this document apply only to the use of RSADSI as a certificate co-issuer; all other certificate generation practices lie outside the scope of this RFC.

This RFC specifies procedures by which users order certificates either directly from RSADSI or via a representative in an organization with which the user holds some affiliation (e.g., the user's employer or educational institution). Syntactic provisions are made which allow a recipient to determine, to some granularity, which identifying information contained in the certificate is vouched for by the certificate issuer. In particular, organizations will usually be vouching for the affiliation of a user with that organization and perhaps a user's role within the organization, in addition to the user's name. In other circumstances, as discussed in section 3.3.3, a certificate may indicate that an issuer vouches only for the user's name, implying that any other identifying information contained in the certificate may not have been validated by the issuer. These semantics are beyond the scope of X.509, but are not incompatible with that recommendation.

The key management architecture described in this RFC has been

designed to support privacy enhanced mail as defined in this RFC, RFC-1113, and their successors. Note that this infrastructure also supports X.400 mail security facilities (as per X.411) and thus paves the way for transition to the OSI/CCITT Message Handling System paradigm in the Internet in the future. The certificate issued to a user for the \$25 biennial fee will grant to the user identified by that certificate a license from RSADSI to employ the RSA algorithm for certificate validation and for encryption and decryption operations in this electronic mail context. No use of the algorithm outside the scope defined in this RFC is authorized by this license as of this time. Expansion of the license to other Internet security applications is possible but not yet authorized. The license granted by this fee does not authorize the sale of software or hardware incorporating the RSA algorithm; it is an end-user license, not a developer's license.

3.2 Relation to X.509 Architecture

CCITT 1988 Recommendation X.509, "The Directory - Authentication Framework", defines a framework for authentication of entities involved in a distributed directory service. Strong authentication, as defined in X.509, is accomplished with the use of public-key cryptosystems. Unforgeable certificates are generated by certification authorities; these authorities may be organized hierarchically, though such organization is not required by X.509. There is no implied mapping between a certification hierarchy and the naming hierarchy imposed by directory system naming attributes. The public-key certificate approach defined in X.509 has also been adopted in CCITT 1988 X.411 in support of the message handling application.

This RFC interprets the X.509 certificate mechanism to serve the needs of privacy-enhanced mail in the Internet environment. The certification hierarchy proposed in this RFC in support of privacy enhanced mail is intentionally a subset of that allowed under X.509. In large part constraints have been levied in order to simplify certificate validation in the absence of a widely available, user-level directory service. The certification hierarchy proposed here also embodies semantics which are not explicitly addressed by X.509, but which are consistent with X.509 precepts. The additional semantic constraints have been adopted to explicitly address questions of issuer "authority" which we feel are not well defined in X.509.

3.3 Entities' Roles and Responsibilities

One way to explain the architecture proposed by this RFC is to examine the various roles which are defined for various entities in

the architecture and to describe what is required of each entity in order for the proposed system to work properly. The following sections identify three different types of entities within this architecture: users and user agents, organizational notaries, and certification authorities. For each class of entity we describe the (electronic and paper) procedures which the entity must execute as part of the architecture and what responsibilities the entity assumes as a function of its role in the architecture. Note that the infrastructure described here applies to the situation wherein RSADSI acts as a co-issuer of certificates, sharing the role of certification authority as described later. Other certifying authority arrangements may employ different procedures and are not addressed by this RFC.

3.3.1 Users and User Agents

The term User Agent (UA) is taken from CCITT X.400 Message Handling Systems (MHS) Recommendations, which define it as follows: "In the context of message handling, the functional object, a component of MHS, by means of which a single direct user engages in message handling." UAs exchange messages by calling on a supporting Message Transfer Service (MTS).

A UA process supporting privacy-enhanced mail processing must protect the private component of its associated entity (ordinarily, a human user) from disclosure. We anticipate that a user will employ ancillary software (not otherwise associated with the UA) to generate his public/private component pair and to compute the (one-way) message hash required by the registration procedure. The public component, along with information that identifies the user, will be transferred to an organizational notary (see below) for inclusion in an order to an issuer. The process of generating public and private components is a local matter, but we anticipate Internet-wide distribution of software suitable for component-pair generation to facilitate the process. The mechanisms used to transfer the public component and the user identification information must preserve the integrity of both quantities and bind the two during this transfer.

This proposal establishes two ways in which a user may order a certificate, i.e., through the user's affiliation with an organization or directly through RSADSI. In either case, a user will be required to send a paper order to RSADSI on a form described in a subsequent RFC and containing the following information:

1. Distinguished Name elements (e.g., full legal name, organization name, etc.)
2. Postal address

3. Internet electronic mail address
4. A message hash function, binding the above information to the user's public component

Note that the user's public component is NOT transmitted via this paper path. In part the rationale here is that the public component consists of many (>100) digits and thus is prone to error if it is copied to and from a piece of paper. Instead, a message hash is computed on the identifying information and the public component and this (smaller) message hash value is transmitted along with the identifying information. Thus the public component is transferred only via an electronic path, as described below.

If the user is not affiliated with an organization which has established its own "electronic notary" capability (an organization notary or "ON" as discussed in the next section), then this paper registration form must be notarized by a Notary Public. If the user is affiliated with an organization which has established one or more ONs, the paper registration form need not carry the endorsement of a Notary Public. Concurrent with the paper registration, the user must send the information outlined above, plus his public component, either to his ON, or directly to RSADSI if no appropriate ON is available to the user. Direct transmission to RSADSI of this information will be via electronic mail, using a representation described in a subsequent RFC. The paper registration must be accompanied by a check or money order for \$25 or an organization may establish some other billing arrangement with RSADSI. The maximum (and default) lifetime of a certificate ordered through this process is two years.

The transmission of ID information and public component from a user to his ON is a local matter, but we expect electronic mail will also be the preferred approach in many circumstances and we anticipate general distribution of software to support this process. Note that it is the responsibility of the user and his organization to ensure the integrity of this transfer by some means deemed adequately secure for the local computing and communication environment. There is no requirement for secrecy in conjunction with this information transfer, but the integrity of the information must be ensured.

3.3.2 Organizational Notaries

An organizational notary is an individual who acts as a clearinghouse for certificate orders originating within an administrative domain such as a corporation or a university. An ON represents an organization or organizational unit (in X.500 naming terms), and is assumed to have some independence from the users on whose behalf

certificates are ordered. An ON will be restricted through mechanisms implemented by the issuing authority, e.g., RSADSI, to ordering certificates properly associated with the domain of that ON. For example, an ON for BBN should not be able to order certificates for users affiliated with MIT or MITRE, nor vice versa. Similarly, if a corporation such as BBN were to establish ONs on a per-subsidary basis (corresponding to organization units in X.500 naming parlance), then an ON for the BBN Communications subsidiary should not be allowed to order a certificate for a user who claims affiliation with the BBN Software Products subsidiary.

It can be assumed that the set of ONs changes relatively slowly and that the number of ONs is relatively small in comparison with the number of users. Thus a more extensive, higher assurance process may reasonably be associated with ON accreditation than with per-user certificate ordering. Restrictions on the range of information which an ON is authorized to certify are established as part of this more elaborate registration process. The procedures by which organizations and organizational units are established in the RSADSI database, and by which ONs are registered, will be described in a subsequent RFC.

An ON is responsible for establishing the correctness and integrity of information incorporated in an order, and will generally vouch for (certify) the accuracy of identity information at a granularity finer than that provided by a Notary Public. We do not believe that it is feasible to enforce uniform standards for the user certification process across all ONs, but we anticipate that organizations will endeavor to maintain high standards in this process in recognition of the "visibility" associated with the identification data contained in certificates. An ON also may constrain the validity period of an ordered certificate, restricting it to less than the default two year interval imposed by the RSADSI license agreement.

An ON participates in the certificate ordering process by accepting and validating identification information from a user and forwarding this information to RSADSI. The ON accepts the electronic ordering information described above (Distinguished Name elements, mailing address, public component, and message hash computed on all of this data) from a user. (The representation for user-to-ON transmission of this data is a local matter, but we anticipate that the encoding specified for ON-to-RSADSI representation of this data will often be employed.) The ON sends an integrity-protected (as described in RFC-1113) electronic message to RSADSI, vouching for the correctness of the binding between the public component and the identification data. Thus, to support this function, each ON will hold a certificate as an individual user within the organization which he represents. RSADSI will maintain a database which identifies the

users who also act as ONs and the database will specify constraints on credentials which each ON is authorized to certify. The electronic mail representation for a user's certificate data in an ON message to RSADSI will be specified in a subsequent RFC.

3.3.3 Certification Authorities

In X.509 the term "certification authority" is defined as "an authority trusted by one or more users to create and assign certificates". This alternate expansion for the acronym "CA" is roughly equivalent to that contemplated as a "central authority" in RFC-1040 and RFC-1113. The only difference is that in X.509 there is no requirement that a CA be a distinguished entity or that a CA serve a large number of users, as envisioned in these RFCs. Rather, any user who holds a certificate can, in the X.509 context, act as a CA for any other user. As noted above, we have chosen to restrict the role of CA in this electronic mail environment to organizational entities, to simplify the certificate validation process, to impose semantics which support organizational affiliation as a basis for certification, and to facilitate license accountability.

In the proposed architecture, individuals who are affiliated with (registered) organizations will go through the process described above, in which they forward their certificate information to their ON for certification. The ON will, based on local procedures, verify the accuracy of the user's credentials and forward this information to RSADSI using privacy-enhanced mail to ensure the integrity and authenticity of the information. RSADSI will carry out the actual certificate generation process on behalf of the organization represented by the ON. Recall that it is the identity of the organization which the ON represents, not the ON's identity, which appears in the issuer field of the user certificate. Therefore it is the private component of the organization, not the ON, which is used to sign the user certificate.

In order to carry out this procedure RSADSI will serve as the repository for the private components associated with certificates representing organizations or organizational units (but not individuals). In effect the role of CA will be shared between the organizational notaries and RSADSI. This shared role will not be visible in the syntax of the certificates issued under this arrangement nor is it apparent from the validation procedure one applies to these certificates. In this sense, the role of RSADSI as the actual signer of certificates on behalf of organizations is transparent to this aspect of system operation.

If an organization were to carry out the certificate signing process locally, and thus hold the private component associated with its

organization certificate, it would need to contact RSADSI to discuss security safeguards, special legal agreements, etc. A number of requirements would be imposed on an organization if such an approach were pursued. The organization would be required to execute additional legal instruments with RSADSI, e.g., to ensure proper accounting for certificates generated by the organization. Special software will be required to support the certificate signing process, distinct from the software required for an ON. Stringent procedural, physical, personnel and computer security safeguards would be required to support this process, to maintain a relatively high level of security for the system as a whole. Thus, at this time, it is not recommended that organizations pursue this approach although local certificate generation is not expressly precluded by the proposed architecture.

RSADSI has offered to operate a service in which it serves as a CA for users who are not affiliated with any organization or who are affiliated with an organization which has not opted to establish an organizational notary. To distinguish certificates issued to such "non-affiliated" users the distinguished string "Notary" will appear as the organizational unit name of the issuer of the certificate. This convention will be employed throughout the system. Thus not only RSADSI but any other organization which elects to provide this type of service to non-affiliated users may do so in a standard fashion. Hence a corporation might issue a certificate with the "Notary" designation to students hired for the summer, to differentiate them from full-time employees. At least in the case of RSADSI, the standards for verifying user credentials that carry this designation will be well known and widely recognized (e.g., Notary Public endorsement).

To illustrate this convention, consider the following examples. Employees of RSADSI will hold certificates which indicate "RSADSI" as the organization in both the issuer field and the subject field, perhaps with no organizational unit specified. Certificates obtained directly from RSADSI, by user's who are not affiliated with any ON, will also indicate "RSADSI" as the organization and will specify "Notary" as an organizational unit in the issuer field. However, these latter certificates will carry some other designation for organization (and, optionally, organizational unit) in the subject field. Moreover, an organization designated in the subject field for such a certificate will not match any for which RSADSI has an ON registered (to avoid possible confusion).

In all cases described above, when a certificate is generated RSADSI will send a paper reply to the ordering user, including two message hash functions:

1. a message hash computed on the user's identifying information and public component (and sent to RSADSI in the registration process), to guarantee its integrity across the ordering process, and
2. a message hash computed on the public component of RSADSI, to provide independent authentication for this public component which is transmitted to the user via email (see below).

RSADSI will send to the user via electronic mail (not privacy enhanced) a copy of his certificate, a copy of the organization certificate identified in the issuer field of the user's certificate, and the public component used to validate certificates signed by RSADSI. The "issuer" certificate is included to simplify the validation process in the absence of a user-level directory system; its distribution via this procedure will probably be phased out in the future. Thus, as described in RFC-1113, the originator of a message is encouraged, though not required, to include his certificate, and that of its issuer, in the privacy enhanced message header (X-Issuer-Certificate) to ensure that each recipient can process the message using only the information contained in this header. The organization (organizational unit) identified in the subject field of the issuer certificate should correspond to that which the user claims affiliation (as declared in the subject field of his certificate). If there is no appropriate correspondence between these fields, recipients ought to be suspicious of the implied certification path. This relationship should hold except in the case of "non-affiliated" users for whom the "Notary" convention is employed.

In contrast, the issuer field of the issuer's certificate will specify "RSADSI" as the organization, i.e., RSADSI will certify all organizational certificates. This convention allows a recipient to validate any originator's certificate (within the RSADSI certification hierarchy) in just two steps. Even if an organization establishes a certification hierarchy involving organizational units, certificates corresponding to each unit can be certified both by RSADSI and by the organizational entity immediately superior to the unit in the hierarchy, so as to preserve this short certification path feature. First, the public component of RSADSI is employed to validate the issuer's certificate. Then the issuer's public component is extracted from that certificate and is used to validate the originator's certificate. The recipient then extracts the originator's public component for use in processing the X-Mic-Info field of the message (see and RFC-1113).

The electronic representation used for transmission of the data items described above (between an ON and RSADSI) will be contained in a

subsequent RFC. To verify that the registration process has been successfully completed and to prepare for exchange of privacy-enhanced electronic mail, the user should perform the following steps:

1. extract the RSADSI public component, the issuer's certificate and the user's certificate from the message
2. compute the message hash on the RSADSI public component and compare the result to the corresponding message hash that was included in the paper receipt
3. use the RSADSI public component to validate the signature on the issuer's certificate (RSADSI will be the issuer of this certificate)
4. extract the organization public component from the validated issuer's certificate and use this public component to validate the user certificate
5. extract the identification information and public component from the user's certificate, compute the message hash on it and compare the result to the corresponding message hash value transmitted via the paper receipt

For a user whose order was processed via an ON, successful completion of these steps demonstrates that the certificate issued to him matches that which he requested and which was certified by his ON. It also demonstrates that he possesses the (correct) public component for RSADSI and for the issuer of his certificate. For a user whose order was placed directly with RSADSI, this process demonstrates that his certificate order was properly processed by RSADSI and that he possesses the valid issuer certificate for the RSADSI Notary. The user can use the RSADSI public component to validate organizational certificates for organizations other than his own. He can employ the public component associated with his own organization to validate certificates issued to other users in his organization.

3.3.3.1 Interoperation Across Certification Hierarchy Boundaries

In order to accommodate interoperation with other certification authorities, e.g., foreign or U.S. government CAs, two conventions will be adopted. First, all certifying authorities must agree to "cross-certify" one another, i.e., each must be willing to sign a certificate in which the issuer is that certifying authority and the subject is another certifying authority. Thus, RSADSI might generate a certificate in which it is identified as the issuer and a certifying authority for the U.S. government is identified as the

subject. Conversely, that U.S. government certifying authority would generate a certificate in which it is the issuer and RSADSI is the subject. This cross-certification of certificates for "top-level" CAs establishes a basis for "lower level" (e.g., organization and user) certificate validation across the hierarchy boundaries. This avoids the need for users in one certification hierarchy to engage in some "out-of-band" procedure to acquire a public-key for use in validating certificates from a different certification hierarchy.

The second convention is that more than one X-Issuer-Certificate field may appear in a privacy-enhanced mail header. Multiple issuer certificates can be included so that a recipient can more easily validate an originator's certificate when originator and recipient are not part of a common CA hierarchy. Thus, for example, if an originator served by the RSADSI certification hierarchy sends a message to a recipient served by a U.S. government hierarchy, the originator could (optionally) include an X-Issuer-Certificate field containing a certificate issued by the U.S. government CA for RSADSI. In this fashion the recipient could employ his public component for the U.S. government CA to validate this certificate for RSADSI, from which he would extract the RSADSI public component to validate the certificate for the originator's organization, from which he would extract the public component required to validate the originator's certificate. Thus, more steps can be required to validate certificates when certification hierarchy boundaries are crossed, but the same basic procedure is employed. Remember that caching of certificates by UAs can significantly reduce the effort required to process messages and so these examples should be viewed as "worse case" scenarios.

3.3.3.2 Certificate Revocation

X.509 states that it is a CA's responsibility to maintain:

1. a time-stamped list of the certificates it issued which have been revoked
2. a time-stamped list of revoked certificates representing other CAs

There are two primary reasons for a CA to revoke a certificate, i.e., suspected compromise of a secret component (invalidating the corresponding public component) or change of user affiliation (invalidating the Distinguished Name). As described in X.509, "hot listing" is one means of propagating information relative to certificate revocation, though it is not a perfect mechanism. In particular, an X.509 Revoked Certificate List (RCL) indicates only the age of the information contained in it; it does not provide any

basis for determining if the list is the most current RCL available from a given CA. To help address this concern, the proposed architecture establishes a format for an RCL in which not only the date of issue, but also the next scheduled date of issue is specified. This is a deviation from the format specified in X.509.

Adopting this convention, when the next scheduled issue date arrives a CA must issue a new RCL, even if there are no changes in the list of entries. In this fashion each CA can independently establish and advertise the frequency with which RCLs are issued by that CA. Note that this does not preclude RCL issuance on a more frequent basis, e.g., in case of some emergency, but no Internet-wide mechanisms are architected for alerting users that such an unscheduled issuance has taken place. This scheduled RCL issuance convention allows users (UAs) to determine whether a given RCL is "out of date," a facility not available from the standard RCL format.

A recent (draft) version of the X.509 recommendation calls for each RCL to contain the serial numbers of certificates which have been revoked by the CA administering that list, i.e., the CA that is identified as the issuer for the corresponding revoked certificates. Upon receipt of a RCL, a UA should compare the entries against any cached certificate information, deleting cache entries which match RCL entries. (Recall that the certificate serial numbers are unique only for each issuer, so care must be exercised in effecting this cache search.) The UA should also retain the RCL to screen incoming messages to detect use of revoked certificates carried in these message headers. More specific details for processing RCL are beyond the scope of this RFC as they are a function of local certificate management techniques.

In the architecture defined by this RFC, a RCL will be maintained for each CA (organization or organizational unit), signed using the private component of that organization (and thus verifiable using the public component of that organization as extracted from its certificate). The RSADSI Notary organizational unit is included in this collection of RCLs. CAs operated under the auspices of the U.S. government or foreign CAs are requested to provide RCLs conforming to these conventions, at least until such time as X.509 RCLs provide equivalent functionality, in support of interoperability with the Internet community. An additional, "top level" RCL, will be maintained by RSAD-SI, and should be maintained by other "top level" CAs, for revoked organizational certificates.

The hot listing procedure (expect for this top level RCL) will be effected by having an ON from each organization transmit to RSADSI a list of the serial numbers of users within his organization, to be hot listed. This list will be transmitted using privacy-enhanced

mail to ensure authenticity and integrity and will employ representation conventions to be provided in a subsequent RFC. RSADSI will format the RCL, sign it using the private component of the organization, and transmit it to the ON for dissemination, using a representation defined in a subsequent RFC. Means for dissemination of RCLs, both within the administrative domain of a CA and across domain boundaries, are not specified by this proposal. However, it is anticipated that each hot list will also be available via network information center databases, directory servers, etc.

The following ASN.1 syntax, derived from X.509, defines the format of RCLs for use in the Internet privacy enhanced email environment. See the ASN.1 definition of certificates (later in this RFC or in X.509, Annex G) for comparison.

```
revokedCertificateList ::= SIGNED SEQUENCE {
    signature      AlgorithmIdentifier,
    issuer         Name,
    list           SEQUENCE RCLEntry,
    lastUpdate     UTCTime,
    nextUpdate     UTCTime}

RCLEntry          ::= SEQUENCE {
    subject         CertificateSerialNumber,
    revocationDate  UTCTime}
```

3.4 Certificate Definition and Usage

3.4.1 Contents and Use

A certificate contains the following contents:

1. version
2. serial number
3. certificate signature (and associated algorithm identifier)
4. issuer name
5. validity period
6. subject name
7. subject public component (and associated algorithm identifier)

This section discusses the interpretation and use of each of these certificate elements.

3.4.1.1 Version Number

The version number field is intended to facilitate orderly changes in certificate formats over time. The initial version number for certificates is zero (0).

3.4.1.2 Serial Number

The serial number field provides a short form, unique identifier for each certificate generated by an issuer. The serial number is used in RCLs to identify revoked certificates instead of including entire certificates. Thus each certificate generated by an issuer must contain a unique serial number. It is suggested that these numbers be issued as a compact, monotonic increasing sequence.

3.4.1.3 Subject Name

A certificate provides a representation of its subject's identity and organizational affiliation in the form of a Distinguished Name. The fundamental binding ensured by the privacy enhancement mechanisms is that between public-key and the user identity. CCITT Recommendation X.500 defines the concept of Distinguished Name.

Version 2 of the U.S. Government Open Systems Interconnection Profile (GOSIP) specifies maximum sizes for O/R Name attributes. Since most of these attributes also appear in Distinguished Names, we have adopted the O/R Name attribute size constraints specified in GOSIP and noted below. Using these size constraints yields a maximum Distinguished Name length (exclusive of ASN encoding) of two-hundred fifty-nine (259) characters, based on the required and optional attributes described below for subject names. The following attributes are required in subject Distinguished Names for purposes of this RFC:

1. Country Name in standard encoding (e.g., the two-character Printable String "US" assigned by ISO 3166 as the identifier for the United States of America, the string "GB" assigned as the identifier for the United Kingdom, or the string "NQ" assigned as the identifier for Dronning Maud Land). Maximum ASCII character length of three (3).
2. Organizational Name (e.g., the Printable String "Bolt Beranek and Newman, Inc."). Maximum ASCII character length of sixty-four (64).
3. Personal Name (e.g., the X.402/X.411 structured Printable String encoding for the name John Linn). Maximum ASCII character length of sixty-four (64).

The following attributes are optional in subject Distinguished Names for purposes of this RFC:

1. Organizational Unit Name(s) (e.g., the Printable String "BBN Communications Corporation") A hierarchy of up to four organizational unit names may be provided; the least significant member of the hierarchy is represented first. Each of these attributes has a maximum ASCII character length of thirty-two (32), for a total of one-hundred and twenty-eight (128) characters if all four are present.

3.4.1.4 Issuer Name

A certificate provides a representation of its issuer's identity, in the form of a Distinguished Name. The issuer identification is needed in order to determine the appropriate issuer public component to use in performing certificate validation. The following attributes are required in issuer Distinguished Names for purposes of this RFC:

1. Country Name (e.g., encoding for "US")
2. Organizational Name

The following attributes are optional in issuer Distinguished Names for purposes of this RFC:

1. Organizational Unit Name(s). (A hierarchy of up to four organizational unit names may be provided; the least significant member of the hierarchy is represented first.) If the issuer is vouching for the user identity in the Notary capacity described above, then exactly one instance of this field must be present and it must consist of the string "Notary".

As noted earlier, only organizations are allowed as issuers in the proposed authentication hierarchy. Hence the Distinguished Name for an issuer should always be that of an organization, not a user, and thus no Personal Name field may be included in the Distinguished Name of an issuer.

3.4.1.5 Validity Period

A certificate carries a pair of time specifiers, indicating the start and end of the time period over which a certificate is intended to be used. No message should ever be prepared for transmission with a non-current certificate, but recipients should be prepared to receive messages processed using recently-expired certificates. This fact results from the unpredictable (and sometimes substantial)

transmission delay of the staged-delivery electronic mail environment. The default and maximum validity period for certificates issued in this system will be two years.

3.4.1.6 Subject Public Component

A certificate carries the public component of its associated entity, as well as an indication of the algorithm with which the public component is to be used. For purposes of this RFC, the algorithm identifier will indicate use of the RSA algorithm, as specified in RFC-1115. Note that in this context, a user's public component is actually the modulus employed in RSA algorithm calculations. A "universal" (public) exponent is employed in conjunction with the modulus to complete the system. Two choices of exponents are recommended for use in this context and are described in section 3.4.3. Modulus size will be permitted to vary between 320 and 632 bits.

3.4.1.7 Certificate Signature

A certificate carries a signature algorithm identifier and a signature, applied to the certificate by its issuer. The signature is validated by the user of a certificate, in order to determine that the integrity of its contents have not been compromised subsequent to generation by a CA. An encrypted, one-way hash will be employed as the signature algorithm. Hash functions suitable for use in this context are notoriously difficult to design and tend to be computationally intensive. Initially we have adopted a hash function developed by RSADSI and which exhibits performance roughly equivalent to the DES (in software). This same function has been selected for use in other contexts in this system where a hash function (message hash algorithm) is required, e.g., MIC for multicast messages. In the future we expect other one-way hash functions will be added to the list of algorithms designated for this purpose.

3.4.2 Validation Conventions

Validating a certificate involves verifying that the signature affixed to the certificate is valid, i.e., that the hash value computed on the certificate contents matches the value that results from decrypting the signature field using the public component of the issuer. In order to perform this operation the user must possess the public component of the issuer, either via some integrity-assured channel, or by extracting it from another (validated) certificate. In the proposed architecture this recursive operation is terminated quickly by adopting the convention that RSADSI will certify the certificates of all organizations or organizational units which act as issuers for end users. (Additional validation steps may be

required for certificates issued by other CAs as described in section 3.3.3.1.)

Certification means that RSADSI will sign certificates in which the subject is the organization or organizational unit and for which RSADSI is the issuer, thus implying that RSADSI vouches for the credentials of the subject. This is an appropriate construct since each ON representing an organization or organizational unit must have registered with RSADSI via a procedure more rigorous than individual user registration. This does not preclude an organizational unit from also holding a certificate in which the "parent" organization (or organizational unit) is the issuer. Both certificates are appropriate and permitted in the X.509 framework. However, in order to facilitate the validation process in an environment where user-level directory services are generally not available, we will (at this time) adopt this certification convention.

The public component needed to validate certificates signed by RSADSI (in its role as a CA for issuers) is transmitted to each user as part of the registration process (using electronic mail with independent, postal confirmation via a message hash). Thus a user will be able to validate any user certificate (from the RSADSI hierarchy) in at most two steps. Consider the situation in which a user receives a privacy enhanced message from an originator with whom the recipient has never previously corresponded. Based on the certification convention described above, the recipient can use the RSADSI public component to validate the issuer's certificate contained in the X-Issuer-Certificate field. (We recommend that, initially, the originator include his organization's certificate in this optional field so that the recipient need not access a server or cache for this public component.) Using the issuer's public component (extracted from this certificate), the recipient can validate the originator's certificate contained in the X-Certificate field of the header.

Having performed this certificate validation process, the recipient can extract the originator's public component and use it to decrypt the content of the X-MIC-Info field and thus verify the data origin authenticity and integrity of the message. Of course, implementations of privacy enhanced mail should cache validated public components (acquired from incoming mail or via the message from a user registration process) to speed up this process. If a message arrives from an originator whose public component is held in the recipient's cache, the recipient can immediately employ that public component without the need for the certificate validation process described here. Also note that the arithmetic required for certificate validation is considerably faster than that involved in digitally signing a certificate, so as to minimize the computational burden on users.

A separate issue associated with validation of certificates is a semantic one, i.e., is the entity identified in the issuer field appropriate to vouch for the identifying information in the subject field. This is a topic outside the scope of X.509, but one which must be addressed in any viable system. The hierarchy proposed in this RFC is designed to address this issue. In most cases a user will claim, as part of his identifying information, affiliation with some organization and that organization will have the means and responsibility for verifying this identifying information. In such circumstances one should expect an obvious relationship between the Distinguished Name components in the issuer and subject fields.

For example, if the subject field of a certificate identified an individual as affiliated with the "Widget Systems Division" (Organizational Unit Name) of "Compudigicorp" (Organizational Name), one would expect the issuer field to specify "Compudigicorp" as the Organizational Name and, if an Organizational Unit Name were present, it should be "Widget Systems Division." If the issuer's certificate indicated "Compudigicorp" as the subject (with no Organizational Unit specified), then the issuer should be "RSADSI." If the issuer's certificate indicated "Widget Systems Division" as Organizational Unit and "Compudigicorp" as Organization in the subject field, then the issuer could be either "RSADSI" (due to the direct certification convention described earlier) or "Compudigicorp" (if the organization elected to distribute this intermediate level certificate). In the later case, the certificate path would involve an additional step using the certificate in which "Compudigicorp" is the subject and "RSADSI" is the issuer. One should be suspicious if the validation path does not indicate a subset relationship for the subject and issuer Distinguished Names in the certification path, expect where cross-certification is employed to cross CA boundaries.

It is a local matter whether the message system presents a human user with the certification path used to validate a certificate associated with incoming, privacy-enhanced mail. We note that a visual display of the Distinguished Names involved in that path is one means of providing the user with the necessary information. We recommend, however, that certificate validation software incorporate checks and alert the user whenever the expected certification path relationships are not present. The rationale here is that regular display of certification path data will likely be ignored by users, whereas automated checking with a warning provision is a more effective means of alerting users to possible certification path anomalies. We urge developers to provide facilities of this sort.

3.4.3 Relation with X.509 Certificate Specification

An X.509 certificate can be viewed as two components: contents and an

encrypted hash. The encrypted hash is formed and processed as follows:

1. X, the hash, is computed as a function of the certificate contents
2. the hash is signed by raising X to the power e (modulo n)
3. the hash's signature is validated by raising the result of step 2 to the power d (modulo n), yielding X, which is compared with the result computed as a function of certificate contents.

Annex C to X.509 suggests the use of Fermat number F4 (65537 decimal, $1 + 2^{2^{16}}$) as a fixed value for e which allows relatively efficient authentication processing, i.e., at most seventeen (17) multiplications are required to effect exponentiation). As an alternative one can employ three (3) as the value for e, yielding even faster exponentiation, but some precautions must be observed (see RFC-1115). Users of the algorithm select values for d (a secret quantity) and n (a non-secret quantity) given this fixed value for e. As noted earlier, this RFC proposes that either three (3) or F4 be employed as universal encryption exponents, with the choice specified in the algorithm identifier. In particular, use of an exponent value of three (3) for certificate validation is encouraged, to permit rapid certificate validation. Given these conventions, a user's public component, and thus the quantity represented in his certificate, is actually the modulus (n) employed in this computation (and in the computations used to protect the DEK and MSGHASH, as described in RFC-1113). A user's private component is the exponent (d) cited above.

The X.509 certificate format is defined (in X.509, Annex G) by the following ASN.1 syntax:

```
Certificate ::= SIGNED SEQUENCE{
    version [0]      Version DEFAULT v1988,
    serialNumber     CertificateSerialNumber,
    signature        AlgorithmIdentifier,
    issuer           Name,
    validity         Validity,
    subject          Name,
    subjectPublicKeyInfo SubjectPublicKeyInfo}

Version ::= INTEGER {v1988(0)}

CertificateSerialNumber ::= INTEGER
```

```
Validity ::= SEQUENCE{
    notBefore      UTCTime,
    notAfter       UTCTime}

SubjectPublicKeyInfo ::= SEQUENCE{
    algorithm       AlgorithmIdentifier,
    subjectPublicKey BIT STRING}

AlgorithmIdentifier ::= SEQUENCE{
    algorithm       OBJECT IDENTIFIER,
    parameters     ANY DEFINED BY algorithm OPTIONAL}
```

All components of this structure are well defined by ASN.1 syntax defined in the 1988 X.400 and X.500 Series Recommendations, except for the AlgorithmIdentifier. An algorithm identifier for RSA is contained in Annex H of X.509 but is unofficial. RFC-1115 will provide detailed syntax and values for this field.

NOTES:

- [1] CCITT Recommendation X.411 (1988), "Message Handling Systems: Message Transfer System: Abstract Service Definition and Procedures".
- [2] CCITT Recommendation X.509 (1988), "The Directory Authentication Framework".

Authors' Addresses

Steve Kent
BBN Communications
50 Moulton Street
Cambridge, MA 02138

Phone: (617) 873-3988

EMail: kent@BBN.COM

John Linn
Secure Systems
Digital Equipment Corporation
85 Swanson Road, BXB1-2/D04
Boxborough, MA 01719-1326

Phone: 508-264-5491

EMail: Linn@ultra.enet.dec.com