Recent Changes to Privacy Enhanced Electronic Mail

Matt Bishop

Department of Mathematics and Computer Science
Dartmouth College
6188 Bradley Hall
Hanover, NH 03755-3551

ABSTRACT

Privacy enhanced electronic mail is a set of protocols which provide confidentiality, authenticity, and integrity for electronic mail. A version of these protocols was released in August 1989, and revised two years later. Since then, several other changes were made to the protocols, many of them minor, but some major. This note describes these changes.

1. Introduction

A previous paper [1] described the privacy enhanced electronic mail protocols proposed in RFC 1113, RFC 1114, and RFC 1115 [2][3][4], along with changes made through June 1991. These protocols provide sender authenticity, message integrity and confidentiality, and under certain circumstances non-repudiation. Cryptography is a key tool for these protocols; first, the message is encrypted using one cryptosystem and key (the session key), and then that key and algorithm identifier are themselves encrypted by a second cryptosystem and key (the interchange key). The protocols provide support for both symmetric and asymmetric interchange cryptosystems.

Several improvements were made during the development of software that implements these protocols, and the design and discussion of an infrastructure to support a large scale key distribution mechanism [5][6][7]. This note discusses those made after June 1, 1991 (where [1] left off). The protocols are those as of April 1, 1992; no major changes are expected for some time.

2. Changes to the Message Format

The biggest change is in the Originator-ID and Recipient-ID fields. Previously, the same field was used for symmetric and asymmetric interchange keys:

Originator-ID: *entity_id*:*issuing_authority*:*version*

The support of grant NAG 2-628 from the NASA Ames Research Center, and of a Dartmouth College Fellowship, is gratefully acknowledged.

Recipient-ID: *entity_id*:*issuing_authority*:*version*

The first component of the field, *entity_id*, is the entity identifier; typically this is simply an electronic mailbox name. The *issuing_authority* is the issuer of the interchange key, and the *version* identifies the interchange key. Note that the last two subfields uniquely identify any interchange key.

With symmetric interchange keys, all components of the field are necessary. However, in the asymmetric case, as identity must somehow be bound to the public key used, the entity identifier subfield is at best redundant and can be at worst misleading; for as the correctness of the entity identifier is not assured, the recipient may be misled unless he checks that identifier with the identity bound to the public key. Hence, these subfields have been discarded for the asymmetric case. The result has been to replace Originator-ID and Recipient-ID with the following four fields:

Originator-ID-Asymmetric: *issuing_authority*,*version*
Originator-ID-Symmetric: *entity_id*,*issuing_authority*,*version*
Recipient-ID-Asymmetric: *issuing_authority*,*version*
Recipient-ID-Symmetric: *entity_id*,*issuing_authority*,*version*

Note the subfield delimiters have been changed to commas to be consistent with the other header fields.

One further simplification has been made. If a certificate-based key management scheme is used for asymmetric cryptosystems, the originator may include his or her certificate in a Certificate header field. As the contents of a Certificate header field is the certificate of the message originator, the Originator-ID-Asymmetric header is redundant. To make this more explicit, the Certificate header field has been renamed to Originator-Certificate, and if this field is present, the Originator-ID-Asymmetric header field may be omitted entirely. Table 1 summarizes these changes to the header field.

Table 1: Changes to Header Fields

header field	*what happened*
Originator-ID: *entity_id:issuing_authority:version*	deleted
Recipient-ID: *entity_id:issuing_authority:version*	deleted
Originator-ID-Asymmetric: *issuing_authority*,*version*	added
Originator-ID-Symmetric: *entity_id,issuing_authority*,*version*	added
Recipient-ID-Asymmetric: *issuing_authority*,*version*	added

Table 1: Changes to Header Fields

header field	*what happened*
Recipient-ID-Symmetric: *entity_id*,*issuing_authority*,*version*	added

The last change is a small one, noted here for completeness. In the case of asymmetric interchange keys, the certificate's serial number is put into the version subfield of the Originator-ID-Asymmetric and Recipient-ID-Asymmetric header fields. These certificate serial numbers are now represented in hexadecimal.

3. Changes to the Certificate Based Key Management Infrastructure

This section describes the changes in the semantics of the certificates and the relationship of entities within the certification hierarchy.

3.1. Overview

There are two major differences between the certificate management infrastructure presented here and the earlier one. First, the current structure is that of a rooted tree hierarchy; the previous version used a forest of rooted trees. This eliminates the need for, and complications produced by, cross-certification. Secondly, the trustworthiness of the data in the certificates issued by top-level certifying authorities could have been difficult to learn; now, the notion of a policy for issuing certificates is more formalized, as the (equivalent of) top-level certification authorities, called *policy certification authorities*, now must publish the precise criteria they use to determine whether to accredit a certificate issuing authority.

The latter comment is particularly important given the complexity of assurance across certification hierarchies. First, the two top-level certification authorities would have to execute a legal agreement in which each accepted the other's mechanism for verifying identity as sufficiently strong. Second, the transitivity of such acceptance could lead to one top level certifying authority unknowingly placing trust in entities which it considers untrustworthy. To prevent this, each top-level certifying authority would need a separate agreement with each top-level certification authority it cross-certified. A cumbersome technical mechanism was added to attempt to enforce this policy – and with the new structure, that mechanism becomes superfluous and can now be discarded.

The new structure, as before, is that of a hierarchy, and a subset of the X.509 certification infrastructure. However, rather than a forest of cross-certified certification hierarchies, the revised

certification architecture now has a single rooted tree. The root node for all certifications will be an *internet policy registration authority,* which will be set up as an arm of the Internet Society, an international, non-profit entity. Directly under the internet policy registration authority will be multiple policy certification authorities, which will in turn certify other authorities subject to a set of policies determined by the policy certification authority. A certifying authority may register with any, or many, policy certification authorities, provided the certifying authority agrees to abide by whatever certification policies the policy certification authorities require, and provided each certificate issued contains a distinct subject public key.

As before, individuals obtain certificates from certification authorities, which in turn may obtain certificates (used to sign certificates) from other certification authorities. However, all these certificates must be issued following a clearly-stated and easily-obtainable policy set by a certification authority superior in the hierarchy; this policy-setting entity is called a *policy certification authority.*

This architecture is much like the directory naming hierarchy of X.500, with two important differences. First, the root of the X.500 hierarchy is not instantiated, whereas the root of this hierarchy is (namely, the internet policy registration authority); and secondly, the policy certification authorities themselves represent roots of subtrees. Bearing these differences in mind, the certification infrastructure could easily be used with OSI.

This section describes the function of the entities in the certification hierarchy, and what they do to certify other entities. We begin at the lower level with ordinary certification authorities, then move to policy certification authorities and finally up to the internet policy registration authority. We close with a few comments on the procedures for validating a certificate and some requirements of conforming software.

3.2. User Agents

Users are the leaves of the certification hierarchy, and processes acting on their behalf (called *user agents*) handle certificates for use in authenticating and protecting the confidentiality of electronic mail. Most critical is the handling of the private key involved in the process, for if this key is exposed the user is open to attack. How the private key is protected is unspecified; for example, it could be on a smart card, or it could be stored encrypted using a symmetric key system such as the Data Encryption Standard. Similarly, particular care must be paid to the generation of public and private keys; if these are poorly chosen, or are generated using a cryptographically weak pseu-

dorandom number generator, an attacker may be able to compromise messages. Special-purpose hardware or software such as the certificate meter described in [1] may be used, but the standards do not require any particular technique. The standard does specify that the holder of the private key need never disclose it; indeed, this preserves the integrity of the privacy enhanced electronic mail mechanisms.

Registration of users is a matter between the user and the relevant certification authority, and may be constrained in any way that the certification authority feels necessary or proper. A user must provide a Distinguished Name and a public key to be put into the user's certificate. The certification authority will validate the user's identity appropriately for the policy being followed (the following section describes sample policies and their effect on this validation step), and then generates the certificate (unless the user has done so already) and signs it using the private key associated with the certificate authorizing that certification authority to issue certificates.

Certificate management is also a local matter; for example, whether a query is sent to a database for a certificate every time a letter is received, or whether certificates are requested once and then cached locally, is not specified. However, the chosen mechanism must be able to determine if a certificate has been revoked by locating and checking the relevant certificate revocation list.

Because propagation of certificates is essential to correct functioning of privacy enhanced electronic mail, in the absence of ubiquitous directory services, certificates may be enclosed in privacy enhanced electronic mail letters using either of two header fields:

Originator-Certificate: *certificate*

contains the certificate of the originator of the letter, and

Issuer-Certificate: *certificate*

contains the certificate of an issuer of another certificate. Only one Originator-Certificate field is allowed, but multiple Issuer-Certificate fields may be present, and need not be in any order. It is intended that these header fields be used to send a (full or partial) certification path; in any case, a user agent must allow the originator to force the agent to enclose a full certification path.

3.3. Certification Authorities

A certification authority is an entity authorized to sign certificates for another entity, which may be either a user or another certification authority. The single general requirement is that the

Distinguished Name of the certification authority issuing the certificate must be a part of the Distinguished name of the subject whose certificate is being signed. For example, if

/C=U/O=Dartmouth College/

is the Distinguished Name[1] of Dartmouth College, and the College is a certifying authority and is signing Matt Bishop's certificate, Matt Bishop's Distinguished Name must have

/C=US/O=Dartmouth College/

as the country and organization attributes. Thus, either of the Distinguished Names

/C=US/O=Dartmouth College/OU=Dept. of Math & CS/CN=Matt Bishop/

or

/C=US/O=Dartmouth College/CN=Matt Bishop/

is acceptable, but

/C=US/O=Research Institute for Advanced Computer Science/CN=Matt Bishop/

is not.

A certifying authority is authorized to issue certificates by a policy certifying authority, and agrees to abide by any policies set by the policy certification authority. The precise mechanism for obtaining this authorization is described later; one relevant point is that each certifying authority must have a unique Distinguished Name. Further, if the same certifying authority is certified by more than one policy certification authority, each certificate must have a distinct public key. This allows a recipient to determine which policy certifying authority issued the certificate.

If a certifying authority issues a certificate to another, subordinate, certifying authority, it is recommended that the policy certifying authority sign that certificate. Then the policy certification authority's certificate can be reached by two validations, and the root of the certification hierarchy by three, from any leaf in the hierarchy; this prevents the validation path from becoming too long.

Certifying authorities are responsible for maintaining a list of certificates that have been canceled before they expire. Once put onto such a Certificate Revocation List, a certificate must remain there until its validity interval expires (after which it could no longer be mistaken for a valid

1. The attributes which could be used in a Distinguished Name were previously restricted to some combination of the country (C), the state or province name (S), the locality name (L), the organization name (O), the organizational unit name (OU), a common name (CN), a title (T), and a postal address (PA); see [1], figure 8, for details. The restriction has been relaxed to simply require that the attribute be on a list maintained by the IANA; also, policy certification authorities are free to impose additional restrictions for their certifying authorities. It is, however, recommended that only those attributes with printable representations be used.

certificate). Each certifying authority decides how often to issue a Certificate Revocation List, and how long each list is valid, provided what the certifying authority does is consistent with any policy requirements (such as bounds for issuance) of its policy certification authority. Each certificate revocation list must be sent to the relevant policy certification authority.

The format of a certification revocation list is based on the X.509 specification, with some changes to make their use simpler in this particular certification hierarchy. Like X.509 certification revocation list, they are digitally signed by the issuer, and contain the issuer's Distinguished Name, a list of revoked certificates (by serial number and time and date at which the revocation was acknowledged by the listing certification authority), and the date of issue. Unlike X.509 certificate revocation lists, all certificates on each list come from the same certifying authority (so no entries need contain the issuer's Distinguished Name). Also, again unlike X.509 lists, each certification revocation list contains the date when the next certification revocation list will be issued, to ensure that the entity obtaining the certification revocation list can tell at a glance if it is the most recent one. This requires the certification authority to issue a new version of the certification revocation list by the scheduled date even if it has not changed from the last date of issue.

3.3.1. Organizational Certification Authorities

These certification authorities issue certificates to entities that are somehow connected with an organization for which the certifying authority has been given authority to issue certificates; for example, these certifying authorities could issue certificates to employees, to organization roles (such as a generic certificate for the company treasurer), and so on. An example of a subject Distinguished Name is:

/C=US/O=Dartmouth College/OU=Dept. of Math & CS/CN=Matt Bishop/

and of an issuer Distinguished Name is:

/C=US/O=Dartmouth College/

Issuing such a certificate implies the affiliation of the entity with the named organization is accurate; so, the subject's Distinguished Name must be subordinate to the issuer's Distinguished Name. In the example, note Matt Bishop's Distinguished Name is indeed subordinate to the issuer's.

Each certifying authority must ensure the subject's Distinguished Name is unique among all the subjects it has certified. Since the Distinguished Name of the certification authority is unique among all certification authorities, and the subject's Distinguished Name is subordinate to that of

the certifying authority, the certifying authority need only ensure it has not previously assigned the additional attributes in the subject's Distinguished Name.

3.3.2. Residential Certification Authorities

These certification authorities issue certificates for entities not affiliated with an organization. They claim that the subject is accurately and uniquely identified by the subject Distinguished Name in the certificate, but the certificate is not issued to that subject as affiliated with an organization, hence the name (residence). It is possible that at some point in the future civil authorities may assume the task of issuing these certificates, but until that time, a policy certification authority is free to establish or certify its own residential certification authority. An example of a subject Distinguished Name is:

/C=US/SP=New Hampshire/L=Hanover/PA=1 Tuck Drive/CN=James Freeman/

and of an issuer Distinguished Name is:

/C=US/SP=New Hampshire/L=Hanover/

Again, note that the issuer's Distinguished Name is superior to the subject's Distinguished Name.

This brings up a subtle point. Different policy certification authorities may establish residential certification authorities covering the same locality but using different policies to decide when to issue certificates. But these residential certifying authorities would all have Distinguished Names based on the same locality, which conflicts with the requirement that all certifying authorities have unique Distinguished Names. The approach taken is to allow non-unique certifying authority Distinguished Names when the certifying authorities are subordinate to different policy certification authorities in this case only. Further, the certification authorities must have different public keys, and must coordinate issuing certificates to ensure unique serial numbers. While there is some danger of confusion, the certificate of the certifying authority contains the name of the accrediting policy certification authority, so during validation of a residential certificate the relevant policy certification authority will be known. This is an interim measure; it is anticipated that civil authorities will take over the issuing of residential certificates at some future time, in which case no ambiguity in Distinguished Names will remain.

However, when certifying a residential user, the residential certification authority must ensure the Distinguished Name is unique. To do this, it sends the internet policy registration authority an entry containing the hash computed on the canonical ASN.1 encoding of the subject's Distin-

guished Name, the user's public key, and the Distinguished Name of the registering certifying authority. The internet policy registration authority then checks its residential use database using the hash and the public key as the lookup key. If the first two are unique, the internet certifying authority informs the residential certifying authority that the user's Distinguished Name is indeed unique and can be registered. The record is also entered into the internet policy registration authority's residential database with an additional field containing the time and date of entry and the Distinguished Name of the issuer. If there is a record with the first two fields the same as the residential user's Distinguished Name, that record is returned. As that record contains the Distinguished Name of the issuer, the residential certifying authority will know which other certifying authority issued the entry, and the two can resolve the matter.

3.3.3. Persona Certifying Authorities

These certification authorities certify that the entity to whom the certificate was issued is uniquely identified, but specifically state that the values of the attributes of the Distinguished Name have no relationship to the subject to which the Distinguished Name is bound. This policy provides a facility for anonymity. If two messages come from the holder of a certificate issued by a persona certification authority, then the recipient can determine that, indeed, both messages came from the same source and were unaltered in transit. However, the recipient cannot tell from the certificate who or what that source is.

Even though the binding of the Distinguished Name is expressly disclaimed, the certification authority must ensure that the Distinguished Name is unique. Further, the certification authority must maintain enough information to allow the holder of the certificate and relevant authorities (and no-one else) to revoke it. The certification authority decides what data is required for this; it need not specifically identify the certificate holder. An example of a subject Distinguished Name is:

/C=US/O=Pseudonymics, Inc./CN=Tiberius Claudius Drusus Nero Germanicus/

and of an issuer Distinguished Name is:

/C=US/O=Pseudonymics, Inc./

Note that there is no indication this certificate was issued by a persona certifying authority; the only way to determine this would be to examine the policy certification authority's policy statement. Because of the potential for confusion with certificates for which the values of the attributes of the Distinguished Name describe the subject to which the Distinguished Name is bound, these certifi-

cates are issued as if by an organizational certifying authority rather than a residential certifying authority.

3.4. Policy Certification Authority

The two aspects which distinguish policy certification authorities from other certification authorities are their policy statement and their proximity to the root of the certification hierarchy.

A policy certification authority issues certificates to subordinate certification authorities, which have agreed to follow the policy of the policy certification authority.

The policy specifies the identity of the policy certification authority, the scope of the policy, considerations of security and privacy, policies and procedures for certification authorities to certify other certification authorities, how revoked certificates are to be handled, what attributes Distinguished Names may (or must) have, and any other relevant matters (such as any fee to be charged).

The specification of the identity of the policy certification authority gives that entity's Distinguished Name, its electronic mail address, its real name, postal address, and telephone number, the date that the policy becomes effective, and how long the policy will remain in effect.

Perhaps the most descriptive part of the policy is its scope: what community or communities is the policy certification authority trying to serve? Will it certify organizational, residential and persona certification authorities, or only some of them?

Given the use of the certification mechanism to provide privacy in electronic mail, the inclusion of a security and privacy section is very appropriate. This section describes the technical and procedural security measures to generate and protect key pairs. If the certification authorities must meet special security requirements, those will be given here. Further, the policy certifying authority must explain how information collected during the certifying of a certifying authority will be protected If the policy certifying authority will act as a residential or persona certifying authority, it must describe how it will protect the information obtained in the process of certifying those individuals.

The certification policy describes the policies and procedures for certifying other certifying authorities. In particular, how does the policy apply transitively to entities certified by certifying authorities? How must an entity's claimed identity be verified before that entity can become a certifying authority? Does the policy certifying authority's policy dictate any maximum lifetime of va-

lidity for any of its certificates, or for those of its subordinate certifying authorities? If multiple entities claim the same Distinguished Name, how will the conflict be resolved?

When registering a Distinguished Name, a policy certifying authority must determine if that name is unique. To do this, it sends the internet policy registration authority an entry containing the hash computed on the canonical ASN.1 encoding of the certifying authority's Distinguished Name, the certifying authority's public key, and its own Distinguished Name. The internet policy registration authority then checks its certification authority database, using the hash and the public component as the lookup key. If the first two are unique, the internet policy registration authority informs the policy certification authority that the certification authority's Distinguished Name is indeed unique and the policy certification authority can register the certification authority; the record is also entered into the certification authority database with an addition field containing the time and date of entry as determined by the internet policy registration authority. If there is a record with the first two fields the same as the certification authority's entry, the internet policy registration authority returns that record. As that record contains the Distinguished Name of the issuer, the policy certification authority will know which other policy certification authority already issued the entry, and the two can resolve the matter.

In the policy statement, the policy certifying authority must also describe how it will handle revoked certificates. They will be added to a Certificate Revocation List, but how frequently will the policy certifying authority and its subordinate certifying authorities issue certificate revocation lists? Also, will the certificate revocation lists be available in ways other than by electronic mail queries? This is needed as the policy certifying authority will receive copies of all certificate revocation lists sent to the internet certifying authority by its subordinate certifying authorities.

The policy must also specify any constraints on the semantics and conventions of Distinguished Names.

Because business documents (such as fees, legal agreements, etc.) will change more frequently than the policy itself, and may be changed to accommodate treasured clients of the policy certification authority, those documents should not be a part of the policy statement itself; but the policy statement should refer to them. As the policy will be digitally signed by the internet policy registration authority, it should be seen as a static document applicable to all subordinates of the policy certification authority; fee schedules, legal agreements, and other business matters tend to change more rapidly and can be altered on a case-by-case basis.

Finally, any additional relevant information that the policy certification authority wishes to include may be included.

3.5. Internet Policy Registration Authority

The internet policy registration authority certifies all policy certification authorities. To be certified, a policy certification authority must agree to abide by the internet policy registration authority's policies, which this section describes.

Each policy certification authority must file a description of its proposed certification policy with the internet policy registration authority. If accepted, this policy will be released as an informational RFC and will be available electronically from (and digitally signed by) the internet policy registration authority. The internet policy registration authority will also check that the policy certification authority's supplied Distinguished Name is unique and legitimate within the context of the certification system. If so, the internet policy registration authority will then digitally sign the proposed certification policy to ensure it is not changed, and then issue a certificate to the policy certification authority. These two acts together provide the policy certification authority with authority to operate under the internet policy registration authority. Also, the policy certification authority must sign a legal agreement, pay a fee, and represent that it has obtained licenses to use any algorithms used to sign certificates and certificate revocation lists.

The internet policy registration authority maintains two databases of Distinguished Names, one for certification authorities and one for residential users. The databases have the same format; each entry has four components: a hash computed on the canonical ASN.1 encoding of the entity's Distinguished Name, the entity's public key, the Distinguished Name of the registering certifying authority (which may or may not be a policy certification authority), and the date and time at which the entry was created by the internet policy registration authority. This database is used to detect duplicate certification authority Distinguished Names, or duplicate residential user Distinguished Names. Note that Distinguished Names are *not* stored in the databases (only their hashes are); this provides a measure of privacy to those entities that do not wish their Distinguished Names to be made known to the internet policy registration authority.

Finally, the internet policy registration authority policy requires each policy certification authority to provide access to a database containing certificate revocation lists for the certification hierarchy. Access is to be robust and via electronic mail. The format of the database (that is, whether it is maintained at the policy certification authority whether it is a distributed database, and so

forth) are up to the various policy and internet certification authorities; but any user must be able to retrieve any certificate revocation list from that database, and any certification authority must be able to submit such a list.

4. Other Changes

Some changes have been made to the validation procedure and implementation requirements.

4.1. Validation procedure

Some new steps have been added to the validation procedure. As each certificate in the validation path is obtained, its validity interval must be checked, and if any is invalid, the user must be warned (although the precise form of the warning is up to the implementation). Each certificate must also be checked against the relevant certificate revocation lists, and the user must again be warned if either a revoked certificate is found or a certificate revocation list cannot be found. In either of these cases, local policy may bar further processing of the message. Finally, for each certificate the issuer of which is not a policy certification authority or the internet policy registration authority, the subject Distinguished Name must be subordinate to the issuer Distinguished Name to ensure that the certification hierarchy is consistent with the protocol architecture. Any certificate failing this test is invalid and is to be rejected.

4.2. Implementation Requirements

Protocols consist of both syntax and processing specifications, and the privacy enhanced mail protocols require some special processing to prevent users from accepting as verified messages which in fact are not verified. This prevents implementations from using procedures that might undermine the security services provided by the privacy-enhanced electronic mail protocols.

First, the ultimate recipient of the privacy-enhanced electronic mail message must be able to associate a Distinguished Name with the originator of the message based on the certification hierarchy and not merely on unauthenticated identification information (for example, in the message headers). Note there is no requirement that the originator's Distinguished Name be sufficient to identify the sender outside the sphere of electronic mail; whether that can be done depends on the type of certificates involved. If a certificate comes from a persona certifying authority, such identification may be impossible.

Second, the originator must be able to determine the recipient's Distinguished Name from his or her identity. Any technique binding the recipient identity to the Distinguished Name – such as displaying recipient Distinguished Names when a message is submitted – will suffice; the goal is to ensure the recipient whose certificate is used is the recipient that the originator intends.

Finally, any user must be able to display the full certification path for any certificate in the certification hierarchy on demand.

5. Changes to the Algorithms, Modes, and Identifiers

The privacy-enhanced electronic mail protocols use the same basic algorithms as described in [1], except that RSA-MD4 has been replaced by RSA-MD5 [8]; however, some ancillary components have changed.

Previously, when asymmetric interchange keys were used, data encryption keys and message integrity checksums were padded differently, and the padding used for message integrity checksums varied in different contexts. The method of padding has now been made uniform, and uses a representation proposed for RSA keys in general [9].

The data encryption key is placed at the end of a block with the same number of bits as the RSA modulus. This block begins with the octets 0x00 and 0x02 (to indicate the quantity is a key). The octet directly before the key is 0x00 (to indicate the end of the padding). The remaining octets are set pseudorandomly, except that no 0x00 octets are allowed. This block can then be encrypted with the recipient's public key.

The message integrity check is handled like the data encryption key, except that the block begins with the octet 0x01 and the padding is octets containing 0xFF; as the message integrity check will have a fixed number of bits (the precise number depending on the algorithm used), there is no danger of confusing the actual message integrity check with the padding.

The padding schemes are consistent across all uses; in particular, the message integrity check computed on the contents of a certificate before signing is to be padded as described above; so is the integrity check of a privacy-enhanced electronic mail message, and the integrity check of a certificate revocation list. This simplifies the structure of that part of the protocols used to check integrity.

Two other relatively minor changes have been made, mainly for clarity and ease of use. When asymmetric interchange keys are used with certificates, the encryption component must be

either $2^{16}+1$ or 3, with a choice of 3 encouraged to enable rapid certificate validation. Finally, when symmetric interchange keys are used, implementations must support the use of the DES in electronic code book mode to encrypt data encryption keys; however, support of the DES in encode-decode-encode mode is optional.

6. Conclusion

The certificate management infrastructure is the most intricate aspect of the privacy-enhanced electronic mail protocols, because it must be acceptable to a large segment of the networking community as well as be able to sustain the needs of electronic mail. In the attempt to balance these two needs, the infrastructure was substantially revised because the earlier version posed several non-technical problems that directly affected the security services provided by privacy-enhanced electronic mail. The current scheme does not suffer from those drawbacks and provides an equally high level of assurance at the technical level.

One interesting point is that the protocol assumes a single root node, the internet policy registration authority. Steve Crocker has proposed a simple generalization to the above mode, in which rather than only one root, there is a list of root nodes. Whenever a new root node is added to the list of root nodes, its public key is available and so any certificate issued in its hierarchy may be validated. This augmentation adds a number of advantages. First, it supports the stated protocol precisely (just make the internet policy registration authority the only root node). Second, it allows for multiple certification hierarchies, and allows the reader to determine under which hierarchy a certificate was issued. Third, it can very simply be used to handle a root node changing its key; just put the root's new certificate into the list. Finally, cross-certification of hierarchies is trivial, since simply exchanging root certificates and placing each in the other's list of roots effects the cross-certification. Whether or not this generalization will be adopted in the future remains to be seen.

Similarly, several minor changes were made to simplify the privacy-enhanced message envelope by eliminating extraneous information from the headers, and even allowing unnecessary headers to be omitted. Other changes made the handling of data used by the cryptographic algorithms consistent across a variety of uses. These enhancements make the protocols more logical, simpler to use, and more compatible with other standards.

Acknowledgments: The most current technical description of the architecture explained here are available as internet draft documents, and this document draws heavily on those drafts. Thanks go

to Steve Kent, Vint Cerf, and Steve Crocker for urging this update be written (and being very patient waiting for it!), to Steve Kent for clarifying some details, to Steve Crocker for sharing his thoughts on the future, to Dave Balenson for showing me an advance copy of the draft of the identifiers, algorithms, and modes document, and describing the changes that had been made, and to the anonymous referees whose comments helped improve this paper. Of course, the author takes credit only for writing this document; its contents are the result of work done by the members of the Internet Research Task Force's Privacy and Security Research Group, and the Privacy-Enhanced Electronic Mail Working Group of the Internet Engineering Task Force; my thanks to all the members of those groups also.

7. References

[1] M. Bishop, "Privacy-Enhanced Electronic Mail," *Journal of Internetworking: Research and Experience* **2** pp. 199-233 (1992).

[2] J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part I -- Message Encipherment and Authentication Procedures*, RFC-1113 (Aug. 1989)

[3] S. Kent and J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part II -- Certificate-Based Key Management*, RFC-1114 (Aug. 1989).

[4] J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part III -- Algorithms, Modes, and Identifiers*, RFC-1115 (Aug. 1989)

[5] S. Kent, *Privacy Enhancement for Internet Electronic Mail: Part II -- Certificate-Based Key Management*, Internet Draft (RFC in progress) (Mar. 1992).

[6] J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part I -- Message Encipherment and Authentication Procedures*, Internet Draft (RFC in progress) (Aug. 1991)

[7] D. Balenson, *Privacy Enhancement for Internet Electronic Mail: Part III -- Algorithms, Modes, and Identifiers*, Internet Draft (RFC in progress) (Apr. 1992)

[8] R. Rivest, The MD5 Message Digest Algorithm, RFC 1321 (Apr. 1992)

[9] RSA Data Security, *PKCS #1: RSA Encryption Standard*, Version 1.4 (June 1991)