

Dartmouth College

Dartmouth Digital Commons

Computer Science Technical Reports

Computer Science

1-1-1991

Privacy-Enhanced Electronic Mail

Matt Bishop

Dartmouth College

Follow this and additional works at: https://digitalcommons.dartmouth.edu/cs_tr



Part of the [Computer Sciences Commons](#)

Dartmouth Digital Commons Citation

Bishop, Matt, "Privacy-Enhanced Electronic Mail" (1991). Computer Science Technical Report PCS-TR91-150. https://digitalcommons.dartmouth.edu/cs_tr/55

This Technical Report is brought to you for free and open access by the Computer Science at Dartmouth Digital Commons. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of Dartmouth Digital Commons. For more information, please contact dartmouthdigitalcommons@groups.dartmouth.edu.

PRIVACY-ENHANCED ELECTRONIC MAIL

Matt Bishop

Technical Report PCS-TR91-150, Revision 3

Privacy-Enhanced Electronic Mail

*Matt Bishop*¹

ABSTRACT

The security of electronic mail sent through the Internet may be described in exactly three words: there is none. The Privacy and Security Research Group has recommended implementing mechanisms designed to provide security enhancements. The first set of mechanisms provides a protocol to provide privacy, integrity, and authentication for electronic mail; the second provides a certificate-based key management infrastructure to support key distribution throughout the internet, to support the first set of mechanisms. This paper describes these mechanisms, as well as the reasons behind their selection and how these mechanisms can be used to provide some measure of security in the exchange of electronic mail.

Index Terms: certificates, cryptography, electronic mail, Internet, privacy, security

1. Introduction

Probably the best-known use of computer networks is for the transmission of electronic mail. Much less widely known is the inherent lack of security in most mailing systems. Recently this has become the subject of much discussion, and some proposals have been made to enhance the security of electronic mail. The X.411 Recommendation [9], the Message Security Protocol MSP [39], and the protocols in RFC 1113 [26], RFC 1114 [23], and RFC 1115 [27] (which we shall call the *privacy-enhanced mail protocols*), all attempt to address this issue.

This last set of protocols is designed to add security-related enhancements to a very large-scale, existing electronic mail structure in a manner transparent to both users and administrators of the network. The most significant contribution of the privacy enhanced mail protocols is to integrate issues of cryptography, systems engineering, key management, user interfacing, and networking on a scale not attempted in practice before. The proposals are being implemented by various groups [3][5][15] and are currently draft Internet standards.

This paper describes the privacy enhanced mail protocols, which specify a set of protocols for sending electronic mail that provides privacy, integrity, and sender authenticity; under certain circumstances, it also provides non-repudiation. The next section presents some background on electronic mailing systems, some relevant aspects of cryptography, and discusses some of the constraints leading to design decisions. Following that, we show how to send a privacy-enhanced message. The fourth section discusses a supporting certificate-based key management architecture, and

gives an example of a mechanism used to support that protocol. The fifth and sixth sections consider an alternate key distribution mechanism and the use of mailing lists. Finally, we conclude by comparing these proposals with the X.411 Recommendation and MSP.

This description reflects the current state of those protocols rather than the state specified in the RFCs; however, differences will be noted.² Also, we have tied the statements in the text to sections in the RFCs so the reader interested in the reference description (again, except for the differences specified) can quickly and easily locate the relevant text.

We should note in passing that descriptions of earlier versions of these protocols have appeared in the literature [5][28][29]; unlike these, however, the focus of this paper is not only on the protocols themselves but also on the reasons that the choices involved in designing the proposals were made. We hope to provide a broader perspective on some of the design decisions, as well as placing them more completely in context than other work.

2. Background and Design Considerations

In this section we review some topics that play a role in enhancing the security of electronic mail. First, we look at a model of message transmittal systems, then we discuss some properties of cryptography, and finally we look at key distribution systems.

2.1. Message Handling System Model

Perhaps the most useful model of mail systems is the *Message Handling System Model* [8] (see Figure 1). It treats the mail system as being composed of a number of connected *Message Transport Agents* and corresponding *User Agents*. A *sender* or an *originator* creates an electronic message and invokes the user agent, which — on behalf of the user — submits that message to a message transport agent. This agent passes the message along to another message transport agent




Figure 1. The Message Handling System Model.

on another host, which passes it to yet another host, and so forth, until it reaches the message transport agent at the destination host. This agent passes the message on to a user agent, which saves the message. The *recipient* may then invoke the user agent to read the message, save it, reply to it, or perform other functions.

When a message is accepted by a message transport agent, it has a particular format the precise description of which depends upon the specific mail system protocol or protocols understood by the agent. When agents attempt to exchange messages, the format of the message must be known to both agents, or it must be translated to a form acceptable to both. For example, message transport agents handling electronic mail from the largest set of connected networks (called the Internet) expect messages to be in a format described in RFC-822 [13], which consists of a set of *header fields* followed by a blank line, and then the body of the message. The information in the header lines is used to forward the message from one message transfer agent to another using the Simple Mail Transfer Protocol described in RFC-821 [35]. Another very large network is that of UNIX³-based systems, which often use an alternate transfer protocol called UUCP [34] to transmit mail messages by telephone.

The existence of numerous message transport agents and protocols suggests that to be used widely, a protocol to enhance security of electronic mail should not require the redesign of existing message transport protocols or message handling systems. The difficulty of coordinating a global change of protocols throughout the Internet would be immense, especially since many system managers would delay installing the new software until they were sure it functioned correctly. Thus, compatibility with existing protocols and software would need to be maintained, so there would be two electronic mail protocols, not just one, extant.

Under this assumption, the enhancements must *not* reside at the message transport level, but instead must be at the user agent level. The security enhancements should be invisible to the message transport agent, and therefore the mail systems would be able to send both privacy-enhanced messages and regular messages.

The above suggestion actually has the force of a constraint since the privacy-enhanced mail proposals are intended for widespread use throughout the Internet, which is not a centrally-managed network. Instead, it is composed of a collection of autonomous networks, each of which has its own management structure. Any attempt to require sites to alter their message transport agents would be resisted.

The privacy, integrity, and authentication services cannot rely on software over which the user has no control; he or she must have some means of verifying that the software has not compromised the privacy or integrity of the message as well as the identity of the sender. Alterations could occur at three points: at the originator's computer, at the recipient's computer, or at any intermediate node. To ensure that the message could not be altered at the first two points would require considering very general issues of operating system security, and certainly require modification of most existing systems. Hence we assume that the sender and recipient can trust the software *above* the message transport agent level on each other's computer, but that they cannot trust software below that level.⁴ This again suggests putting the enhancements *above* the transport level at the end systems as well as using cryptography to protect the messages, since any controls built into the message transport agents can be subverted (by the system operators, if not by anyone else).

2.2. Cryptography for Authentication and Privacy

Protecting the privacy of messages on a network requires encryption whenever the network is physically beyond the control of the sender. Since he has no power to prevent another from listening in, he must rely on mechanisms that prevent disclosure and enable authentication *even when* a wiretapper is present; ideally, whether that individual is active or passive should be irrelevant.

The literature describing the protection of remote communications suggests using two keys [42]. The first is a one-time key, selected pseudorandomly, called the *session key* or the *data exchange key*; this key is used to encrypt the messages sent during the session. The second key is an *interchange key* associated with the user and/or recipient, and is used to encrypt the data exchange key; the data exchange key is transmitted when the session is begun and then retained for further use. This protocol has three advantages over using the interchange key as a data exchange key.

If the cryptosystem for the interchange key is symmetric (classical), it is often theoretically possible to derive this key given sufficient ciphertext, especially if the corresponding plaintext is known. For strong cryptosystems, such cryptanalysis requires large amounts of ciphertext (and, possibly, corresponding plaintext); hence, restricting the use of the interchange key to encrypting small amounts of data, and then using different data exchange keys to encrypt (possibly large) amounts of data only once, limits the amount of data a listener will have available to determine the key. Even if the data exchange key is compromised, only the single session will be known; the listener will have to derive a new data exchange key for the next session.

If the interchange key is that of an asymmetric (public key) cryptosystem, the encryption and corresponding decryption will be more expensive in time, space, or both than those for symmetric cryptosystems; but the short data exchange key can be encrypted using it without affecting overall performance very much. Hence the trade-off chosen for the privacy-enhanced mail protocols was to use a symmetric data encryption key and an asymmetric interchange key.

Sender authentication and protection of message integrity are provided by sending with a message its *digital signature*; this is a function which computes a value based on the contents of a message. The function must be easy to compute, and it must be computationally infeasible to find any two inputs which produce the same output. For example, in a public key cryptosystem, encrypting a message with a private key produces a digital signature as large as the message; anyone can validate the signature (as the corresponding public key is widely available), and it is computationally infeasible to generate a message which will produce the same digital signature.

In practice, it is also undesirable to double the size of a message to ensure integrity. For this reason, digital signatures are often produced by using a *manipulation detection code* (also called a *message integrity check*) to compute a small, fixed-size hash of the message, and then encrypting (signing) the hash. These codes, of which one-way (non-invertible) hash functions are examples, must be easy to compute but, as with digital signatures in general, it must be computationally infeasible to find any pair of different files which produce the same manipulation detection code [14].

Both these schemes also provide an integrity check for the message. One can determine the contents of the message, or what the message hashed to, when signed; if the message sent with the signature disagrees or hashes to a different value, then either the message or the signature was altered in transit and the integrity of the communication should be regarded as violated.

Finally, if no one other than the sender has access to the sender's private key, it would not be possible for the sender to disavow a signed message. Should the sender do so, the recipient can prove the message was sent by a party with the sender's private key; since this key should be known only to the sender, the sender would then have to demonstrate that someone had stolen this key. Note that this assumes a trusted system of some sort, because if that part of the sender's system which handles the private key could be compromised, the private key could be stolen. It also assumes that messages are timestamped, that keys are issued to users whose identities have been authenticated in some fashion (for example, by a public notary), and that all compromised keys are reported at once. (A note of terminology: the key to a classical cryptosystem is called a *secret key* in this paper, whereas the hidden key for a public key cryptosystem is called a *private key*.)

Using cryptography implies that there must be a mechanism for distributing cryptographic keys to communicating parties; we now examine the ways that can be done.

2.3. Key Distribution and Management

Creating and distributing interchange keys is a complex problem; one solution described in the literature is to create *key distribution centers*. If the interchange keys are keys for symmetric ciphers, the key distribution centers must be trusted. Given that the Internet is composed of many autonomous domains, which are themselves often composed of other autonomous subdomains, it is very unlikely that members of such an amalgamation would trust key distribution centers not under their control. Hence the privacy-enhanced mail protocols suggest using an asymmetric cryptosystem for interchange keys, and provide a standard for the management of those keys by encapsulating the required data in *certificates*.

The certificates consist of the user's public key as well as information identifying the associated user and a version number of some form so that, if a user has more than one public key, the sender can identify which one was used. Also a digital signature must be included to protect the integrity of the stored data; this is a hash encrypted using the private key of an issuer, so that anyone may verify the integrity of the data by recomputing the hash and comparing that value with the value obtained by decrypting the stored signature using the stored public key. Note that the key distribution center storing the certificates need not be trusted, since the encrypted hash value binds the user, version, and public key together, and any alterations can be detected by the sender before the key is used. For this reason, that center is referred to as a *directory*.

The privacy enhanced mail protocols separate the key distribution scheme from the message encoding scheme. This way, users operating in an environment where they could trust a central server to manage interchange could do so, whereas users in an environment without such a server could use a certificate-based key distribution mechanism.

3. Sending a Privacy-Enhanced Mail Message

The interoperability constraints described in the introduction suggest encapsulating the header fields related to the enhancements and the privacy-enhanced body of the message so that the privacy-enhanced mail message becomes the body of a regular mail message. This has two effects. First, some user agents do not allow users (or other processes) to add special header fields holding the extra information needed to send privacy-enhanced mail;⁵ this is not a problem if the


```

-----BEGIN PRIVACY-ENHANCED MESSAGE BOUNDARY-----
Proc-Type: 4, ENCRYPTED
DEK-Info: DES-CBC, 3729F9DC6300925A
Originator-ID: someone@somewhere.com::
Recipient-ID: someone@somewhere.com:ptf-kmc:3
Key-Info: DES-ECB, RSA-MD2,
          0F65D99570758593, AEE05B42181E5E261B301291D83DB8F1
Recipient-ID: towho@somewhere.else.com:ptf-kmc:4
Key-Info: DES-ECB, RSA-MD2,
          4B425C85F819327E, 0F5FACDB20FD89B9CA8636AE7E70BE74

iwf9eh8CjiF4A9OKf8ZPayOHjAZNLvWYxsGTIr8xQxfQ0uSN4GbnHHzJPXk6saOd
rTLE42WxcsYZsYYZaBfbpBR3GF3nwwIY100s3LZTUuwsfsGeIu/v2C+dfSlq/wIE
VesyVXPegFX1MSxZb5bv1J7WK9GE3G58TgwzuYNROPcF61lFq7RDWms61ji3ZyFy
DNkG/JlAeWlMgP8WpfJgVg==
-----END PRIVACY-ENHANCED MESSAGE BOUNDARY-----

```

Figure 2. Sample encapsulated portion of a message (symmetric interchange keys).

relevant header fields are encapsulated in the body. Secondly, as stated in the introduction, an important goal is to be compatible with current electronic mail processing, and adding new header fields requires that existing message transport agents be able to handle such fields. But if the enhancements are simply part of the body of a message, the message transport agents need know nothing about the new headers.

The encapsulation mechanism chosen is one widely used in electronic mail throughout the Internet [38]; that standard delimits encapsulated material between two delimiter lines. But, as privacy-enhanced mail allows multiple encapsulated bodies to be present either sequentially or with non-privacy enhanced messages interspersed, using the same delimiter to begin and end encapsulated text introduces ambiguity. So the beginning and ending delimiters must be different. Specifically, the encapsulated portion of each message is preceded by a line containing the delimiter

```

-----BEGIN PRIVACY-ENHANCED MESSAGE BOUNDARY-----

```

following which come the *encapsulated header fields*, a blank line, and the *encapsulated text*; a line containing the delimiter

```

-----END PRIVACY-ENHANCED MESSAGE BOUNDARY-----

```

indicates the end of the privacy-enhanced messages.⁶ Figure 2 and Figure 3⁷ show examples of encapsulated portions of privacy enhanced messages. Note that if the entire message is protected, a blank line must separate the first delimiter from the enclosing headers; this is required by [13].

3.1. Transforming the Encapsulated Body

In addition to encrypting, messages must be transformed into a canonical form so that the encrypted text can be decrypted, and the message integrity check recomputed, correctly despite the vagrations of intermediate mail agents. As this proposal is designed for the Internet, the transformed encapsulated message must be sent in a form acceptable to SMTP message transport agents; the relevant requirements are:⁹

1. all characters must be members of the 7-bit ASCII character set; the successful transmission of the eighth (high-order) bit is not guaranteed (and in practice often does not work);
2. text lines may be no more than 1000 characters long;
3. text lines must be delimited by a carriage return "<CR>" followed by a line feed "<LF>";
4. the character sequence "<CR><LF>.<CR><LF>", which is used by the message transport agents to indicate the end of a message, cannot appear within the body of the message.

Because most computer systems do not use a local representation which complies with this standard (and which in fact varies between computers), some transformation of mail messages is necessary to provide interoperability between hosts using different local representations. Since the encapsulated body of a privacy-enhanced mail message may be encrypted, resulting in characters where the high-order bit must be transmitted accurately for the message to be decrypted correctly, the transformation must take this into account to provide transparency to the underlying message transport agents.

Hence the encapsulated body undergoes a three step transformation: first it is put into a machine-independent character-oriented format, then an integrity check is generated and (if required) the message is encrypted, and finally the resulting bit stream is converted to a special set of printable characters suitable for processing by any (reasonable) message transport agent. Initially, of course, the message text is entered into the system, for example with a text editor; the representation of the characters is that used by the local computer system and is called the *local form*. The local form is then altered in three steps.

3.1.1. Canonicalization¹⁰

The message is converted into a canonical form to ensure that encryption, decryption, and integrity checking all are done on a consistent representation of the message despite the differences between end systems and the vagaries of intermediate mail agents. The representation chosen is to transform characters in the message into their ASCII representations, with the parity (high-order) bit cleared and with line (record) delimiters changed to <CR><LF>; note that this is almost the representation used by SMTP [35].¹¹ Although any consistent representation would do, this one was chosen because it is a common representation familiar to all SMTP message transport agents, and therefore software to do this is widely available throughout the Internet.

3.1.2. Authentication and Encryption¹²

An integrity check is then computed using one of a number of specified algorithms; note that this check can be verified on any type of destination computer because it is computed on the canonical representation of the message and not the local, machine dependent, representation. Currently, three algorithms are specified. The DEA-1-based Message Authentication Code (MAC) algorithm [18] is a well-known message integrity checksum algorithm, has been examined extensively for weaknesses (see for example [1]), and appears to be quite strong; however, it is only suitable when sending a message to *one* recipient and does not provide for non-repudiation [30]. Hence its use is strongly discouraged for messages sent to more than one party. An alternate algorithm, the RSA-MD2 Message Digest Algorithm[27], does not suffer from this weakness and while formal analyses of it are not available, it does not appear to have any exploitable cryptographic weaknesses. Its successor, the RSA-MD4 Message Digest Algorithm [36] also appears to be quite robust and very difficult to compromise so, in the absense of evidence to the contrary, it has also been defined as an acceptable algorithm.¹³

If privacy is required, the message will then be encrypted using an appropriate encryption algorithm with a data encryption key generated for the message.¹⁴ Currently, the only encryption algorithm allowed is the DES in cipher block chaining mode [16][17]. This U. S. government standard is in widespread use, has withstood the test of time, can be implemented very efficiently in either hardware or software, and all attacks known are very time-consuming.¹⁵ Padding is done by creating an octet containing the number of extra octets needed, and replicating it appropriately. For example, if the message required 3 octets of padding, the bits “030303” (in hex) would be added. Up to eight octets may be added to ensure the input can be unambiguously decrypted.¹⁶

header field	subfields	used ...	see section
Proc-Type	protocol, proc_type	always	3.1.2.
DEK-Info	enc_alg, init_vector	for encrypted messages	3.1.2.
Sender-ID	address:iss_auth:ver	always	3.2.1.
Recipient-ID	address:iss_auth:ver	always	3.2.1.
Key-Info	ik_use, mic_alg, dek, mic	with symmetric IKs	3.2.2.
Key-Info	ik_use, dek	with asymmetric IKs	3.2.3.
MIC-Info	ik_use, mic_alg, mic	with asymmetric IKs	3.2.3.
Certificate	enc_cert	with certificate-based KM	4.1.
Issuer-Certificate	enc_cert	with certificate-based KM	4.1.
CRL	crl_encoded	with certificate-based KM	4.2.5.

Figure 4. Summary of header fields; here, “KM” stands for “key management.” See the referenced sections for more information on each.

protocol described in	<i>protocol</i> subfield value
RFC-989	1
RFC-1040	2
RFC-1113	3
<i>next RFC</i>	4

Figure 5. Table of Protocol Numbers.

Figure 4 summarizes the encapsulated header fields that convey the information needed to decrypt, and check the integrity of, the message. All privacy-enhanced messages contain the encapsulated header field

Proc-Type: *protocol_version, proc_type*

which indicates the type of processing done. Its first subfield, *protocol*, is the number of the protocol used (see Figure 5). Its second subfield, *proc_type*, is MIC-ONLY if *no* part of the message is encrypted, ENCRYPTED if the encapsulated message is encrypted, or MIC-CLEAR if the message is unencrypted and not encoded as described in section 3.1.3.¹⁷ A second header field, required only for encrypted messages, conveys the message encryption algorithm and initialization vector:

DEK-Info: *encrypt_alg, init_vector*

The first subfield is a string indicating how the message was encrypted; as only the DES in cipher block chaining mode is defined for use, this string must be DES-CBC. The second subfield is the initialization vector. Although the initialization vector is usually encrypted to prevent spoofing by altering key bits (and also altering the decryption of the first block), the integrity check will detect such tampering; hence the initialization vector need not be encrypted and is represented as a string

0	A	8	I	16	Q	24	Y	32	g	40	o	48	w	56	4
1	B	9	J	17	R	25	Z	33	h	41	p	49	x	57	5
2	C	10	K	18	S	26	a	34	i	42	q	50	y	58	6
3	D	11	L	19	T	27	b	35	j	43	r	51	z	59	7
4	E	12	M	20	U	28	c	36	k	44	s	52	0	60	8
5	F	13	N	21	V	29	d	37	l	45	t	53	1	61	9
6	G	14	O	22	W	30	e	38	m	46	u	54	2	62	+
7	H	15	P	23	X	31	f	39	n	47	v	55	3	63	/

Figure 6. Printable Encoding Characters

of 16 hexadecimal digits. The integrity check and data encryption key will be transmitted after being encrypted with the interchange key, and we shall defer the details until the next section.¹⁸

3.1.3. Printable Encoding¹⁹

After the second step, the message text should be thought of as a bit stream. (The encapsulated headers are not part of this stream unless they are replicated within the message text for integrity checking purposes.) Note that even though there are an integral number of characters, all 8 bits may be significant. This form must be changed into another that meets the SMTP requirements stated above. The simplest way to do this is to expand the message.

The first thought is to group the bits into sets of 7, and transmit the 7-bit ASCII character corresponding to each. However, since some of those character sequences are special to SMTP (such as "<CR><LF>.<CR><LF>") and others to the privacy enhancements (such as "--", which occurs in the encapsulation delimiters), there would need to be an escape sequence. It is easier to ignore this issue by restricting the set of characters to those which are not special. The letters and digits fall into this category, as do numerous punctuation characters, but there are fewer than 128 (2^7) of them. Hence each bit stream is grouped into sets of 6 bits, each of which is mapped into a character in the alphabet shown in figure 6.²⁰

The final set of bits in each stream may contain 2 or 4 bits rather than 6. In the former case, the two meaningful bits are padded with four cleared bits, and two "=" characters are appended to indicate the last 4 bits are padding. In the latter case, the pad is two cleared bits, and one "=" is appended. Finally, the resulting character stream is split into lines of 64 printable characters (except, possibly, for the last line, which may contain fewer).²¹ This form is suitable for transmission by any SMTP-like message transport agent.

Grouping the stream into sets of 4 bits rather than 6 would make encoding faster and eliminate the special handling of the final set of bits. However, this would also double the length of the

encoded message. With most messages, network transmission times will dominate the time to encode the message, so the 6-bit encoding was chosen.

Because reading unencrypted, integrity-checked messages encoded in this format would require all user agents to be able to translate this printable encoding, the privacy-enhanced mail protocols allow this step to be skipped for such messages; the second field in the `Proc-Type` header field will be `MIC-CLEAR` for such messages. However, as `MIC-CLEAR` messages are not printably encoded, privacy-enhanced user agents should indicate that failure of the integrity checking for this type of message may result from differences in character representations between the end hosts and transformations by the message transport agents rather than an alteration of the contents of the message.²²

3.2. Transmission of Integrity Check and Data Encryption Key

We now turn to the transmission of the data encryption key and the message integrity check. The former must be encrypted using an interchange key to protect the privacy of the message, the latter to ensure authenticity and integrity. The mechanism used to do this depends on whether the interchange key is used with a classical cryptosystem or with a public key cryptosystem. In either case, the interchange key depends upon proper identification of the sender and recipient.

3.2.1. Sender, Recipient, and Interchange Key Information²³

The sender and recipient of each message are identified in the following encapsulated header fields:²⁴

```
Originator-ID: entity_id:issuing_authority:version
Recipient-ID:  entity_id:issuing_authority:version
```

These fields specify which interchange key was used. The first subfield contains the identity of the sender or receiver, the second the identity of the authority issuing the interchange key, and the third an indicator of the specific interchange key being used. Since interchange keys depend on both the sender and the recipient, each `Recipient-ID` field is associated with the last preceding `Sender-ID` line.

The first subfield is mandatory; it also requires that each sender and recipient be uniquely identifiable. Hence it assumes the form *user@host*, where *user* is unique to *host* and *host* is unique throughout the set of hosts using electronic mail. Any scheme guaranteeing this will work, but those sites which transmit into the Internet should use the fully qualified domain name for *host*,²⁵ that name will be processed in a case-insensitive manner.²⁶

The second subfield contains the unique name of the authority that issued the interchange key; it need not be unique among all entities, but must be unique over all issuing authorities.²⁷ If certificate-based asymmetric key management is used, then this name is to be the issuing authority's Distinguished Name encoded using the basic encoding rules of ASN.1 [7] and represented using the printable encoding form described in section 3.1.3.²⁸ This convention was chosen because it is also used in certificates to identify issuing authorities unambiguously. The nature of a Distinguished Name will be described later; however, as examples, in Figure 3 the issuing authorities of the certificates of the recipients are the Home Office of Somewhere, Inc. and Security of Elsewhere Corp., both in the United States.

The third subfield is some number or string selected by the issuing authority to disambiguate among the interchange keys issued by that authority to the sender/recipient pair. For certificate-based interchange key management schemes, this subfield must be the serial number of the certificate; for other schemes, the proposal recommends the use of timestamps, which not only provide uniqueness but also allow an expiration date to be prescribed (for example, two years after issue).²⁹

If the subfields contain redundant information, they may be omitted.³⁰ For example, if the interchange key is used in a symmetric cryptosystem, the contents of the last two subfields of both Originator-ID and Recipient-ID fields will be the same for each sender/recipient pair, and so are normally omitted from the Originator-ID field. In Figure 2, the field

Originator-ID: someone@somewhere.com::

identifies the sender as user *someone* at host *somewhere.com*, but leaves the issuing authority and version subfields blank. Additional relevant information is in the associated Recipient-ID line:

Recipient-ID: someone@somewhere.com:ptf-kmc:3

means that the recipient is *someone@somewhere.com*, that the interchange authority named *ptf-kmc* has issued an interchange key for *someone* to use when he sends messages to himself (since the last preceding Originator-ID field identified *someone@somewhere.com* as the sender), and that this key can be identified uniquely by the string 3. The second such field,

Recipient-ID: towho@somewhere.else.com:ptf-kmc:4

means that *ptf-kmc* has also issued an interchange key for *someone@somewhere.com* to use when sending messages to *towho@somewhere.else.com*, and that this key can be identified uniquely by the string 4.

Similarly, if the sender's public key is sent in the message as a certificate containing both the version number and issuing authority (using the certificates and the Certificate header field described in the next section), the last two subfields in the Originator-ID field may be omitted. In Figure 3, the certificate in the Certificate line is that of the sender in the preceding Originator-ID line, so the last two subfields of that field are omitted. But the Recipient-ID lines specify the issuer and serial number of the recipients' certificates, so that they can determine which of their public keys was used to encrypt the data encryption key.

3.2.2. Symmetric Interchange Keys

With symmetric cryptosystems, the encryption key is either the same as, or easily derivable from, the decryption key. For this reason, the *interchange key* is defined as the single key associated with both the sender and the recipient, and the message integrity check and data encryption key follow the identification of each sender and recipient.³¹

The message integrity check and data encryption key for each sender and recipient pair are given on lines of the form³²

Key-Info: *ik_use, mic_algorithm, dek, mic*

One such line normally follows each recipient name, and uses the interchange key associated with the last-preceding sender and that recipient. The first subfield identifies the algorithm which was used to encrypt the data exchange key; the proposal requires using either the DES in electronic code book mode (indicated by DES-ECB) [17] or the DES in encrypt-decrypt-encrypt mode (indicated by DES-EDE) [2]; these were chosen because they are cryptographically very strong, can be implemented efficiently in hardware or software, and are used by standards relied upon throughout different communities.³³ The second subfield identifies the algorithm used to generate the message integrity check; this subfield must be MAC (for the DEA-1-based algorithm), RSA-MD2 (for the RSA-MD2 Message Digest Algorithm) or RSA-MD4 (for the RSA-MD4 Message Digest Algorithm).³⁴ The third and fourth subfields are the data exchange key and the message integrity check encrypted using the algorithm given by the first subfield with the interchange key. Both are represented as strings of hexadecimal digits. For example, in Figure 2, the line

Key-Info: DES-ECB, RSA-MD2,
0F65D99570758593, AEE05B42181E5E261B301291D83DB8F1

indicates that the message integrity check was computed using the RSA-MD2 Message Digest Algorithm, and both it and the data exchange key were encrypted using the interchange key and

the DES algorithm in electronic code book form. The second Key-Info line in that figure indicates the same algorithms, but since another Recipient-ID line has occurred since the last Key-Info line, the data encryption key and message integrity check were encrypted with the interchange key for the sender/recipient pair *someone@somewhere.com* and *towho@somewhere.else.com* (and hence look quite different than in the preceding line).

3.2.3. Asymmetric Interchange Keys

With public key cryptosystems, encryption and decryption keys are different, and neither can be derived merely by knowing the other. This allows authentication and encryption to be done with two different keys, and so the *interchange key* is comprised of both the recipient's public key and the sender's private key, the former being used to encrypt the data exchange key (privacy), and the latter being used to encrypt the message integrity check (integrity/authenticity).³⁵ The message integrity check and data encryption keys are given by two different header fields because the parts of the interchange key differ. The message integrity check is given in header fields of the form³⁶

Key-Info: *ik_use, dek*

The first subfield is the name of the algorithm used to encrypt the data encryption key; since only one algorithm is currently defined for this purpose, the first subfield will always be RSA.³⁷ The second subfield is the data exchange key, encrypted using the recipient's public key and represented as a string of printable encoding characters using the encoding transformation. For example, in Figure 3 the header fields

```
Recipient-ID: someone@somewhere.com:MF8xCzAJBgNVBAYTA1VTMR0wGAYDV
QQKEw9Tb21ld2hlcmUsIEluYy4xHjAcBgNVBASTCkhvbWUgT2ZmaWNl:3
Key-Info: RSA,
1BLpvXR0UrUzYbkNpk0agV2IzUpk8tEjmF/zxB+bATMtPjCUWbz8Lr9wloXIkJHU
Recipient-ID: towho@somewhere.else.com:MF8xBjAJBgNVBAYTA1VTMR0wGA
YDVQQKEw9FbHNld2hlcmUgQ29ycC4xHjAcBgNVBASTCFNlY3VyaXR5:4
Key-Info: RSA,
NcUk2jHEUSoH1nvNSIWL9MLLrHB0eJzyhP+/fSStdW8okeEnv47jxe7SJ/iN72oh
```

show that the sender (*someone@somewhere.com*) has encrypted the data encryption key using RSA and his public key for the first recipient (himself), and using RSA and *towho@somewhere.else.com*'s public key for the second.

The message integrity check is transmitted similarly, in a header field of the form³⁸

MIC-Info: *ik_use, mic_algorithm, mic*

The first subfield identifies the algorithm used to compute the integrity check, and so must be either MAC (if the DEA-1-based algorithm was used), RSA-MD2 (if the RSA Message Digest Algorithm was used), or RSA-MD4 (if the newer RSA Message Digest Algorithm was used).³⁹ The second subfield identifies the algorithm used to encrypt the integrity check; again, currently this must be RSA. The final subfield is the message integrity check encrypted using the sender's private key and then represented as a string of printable encoding characters using the encoding transformation. If the message is encrypted, then the message integrity check must also be encrypted (using the same algorithm and data encryption key as the message); otherwise, if an encrypted message known to correspond to one of n possible plaintexts, were intercepted, then comparing the intercepted message's integrity check with those of the plaintexts would show precisely which one the intercepted message corresponded to.⁴⁰ Note that the information is associated with the sender and not the recipient, so this line should only occur after the Originator-ID field. In Figure 3, the header field

MIC-Info: RSA-MD2, RSA,

5rDqUcMlK1Z6720dcBWGGsDLpTpSCnptJ6UiRRGcDSvzrsoK+oNvqu6z7Xs5Xfz

and its location show that the message integrity check was computed using the RSA-MD2 Message Digest Algorithm, and encrypted using the RSA algorithm and *someone@somewhere.com*'s private key.

4. Key Management Infrastructure

If asymmetric interchange keys are used, a method for ensuring that correct public keys are made available must be provided; without such assurance, an active wiretapper could spoof the intended recipient and read the sender's supposedly private message, or could alter or forge messages from another sender. The privacy-enhanced mail protocol suggests using X.509 [11] compliant certificates containing an entity name, a public key, an issuer name, and a digital signature created by the issuer based on its public key and the information in the certificate. This section presents the architecture for managing those certificates.

4.1. Overview of Certificate Management Architecture

The architecture envisions a set of trees with the inner nodes being organizations entitled to issue certificates. The root of each tree will be a *top-level certification authority*; each of the inner nodes will be a *certification authority*. The leaves are subjects to whom certificates are issued, and who are not authorized to issue certificates themselves. The next higher tier is composed of

organizations who may certify users; the next tier is composed of organizations who may certify other organizations, and so forth. In all cases a certifying authority may only certify those entities directly beneath it. The top-level certifying authority (or any certifying authority, for that matter) establishes its own procedures for and rules for determining what entities it will certify, how those entities request certification, and how the certification is done.⁴¹ This arrangement of certifying authorities within each tree is a subset of the certification hierarchy allowed under X.509.

The problem of transmitting and obtaining certificates for use with privacy-enhanced electronic mail in the absence of ubiquitous directories containing certificates is solved by including one or more instances of two optional encapsulated header fields in the privacy-enhanced mail message. The encapsulated header field⁴²

`Certificate: encoded_certificate`

contains the sender's certificate as a bit stream represented as a string of printable encoding characters using the encoding transformation. Normally only one such field would be present. Similarly, the encapsulated header field⁴³

`Issuer-Certificate: encoded_certificate`

is like the `Certificate` field, except that it contains the certificate of the issuer of a certificate. As many of these header fields as necessary to enable the recipient to validate the certificates in either the `Certificate` or in another `Issuer-Certificate` header field may be present. Note that this field is optional, and once directory servers containing certificates become widely available, will fade into disuse. Examples of both these fields are given in Figure 3.

We now examine the contents of certificates, the management of certificates, and how the two relate to one another.

4.2. Certificate Definition and Use

A *certificate* binds a public key to an issuing authority, subject (which may be a user or some other entity), and other information. Each contains a version number indicating which certificate format is used (currently the only format defined is this one, so the version number will be 0),⁴⁴ a subject identity, an issuer name, a serial number defined by the issuer (certifying authority) and unique to that certifying authority,⁴⁵ a validity period during which the certificate is valid,⁴⁶ the subject's public key, and a certificate signature binding all the above information together.⁴⁷ The signature is generated by computing a hash of the certificate and encrypting the hash with the certifying authority's private key. Currently the only signature algorithm is the `md2WithRSA` al-

<i>attribute</i>	<i>meaning</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>
C	country (ISO 3166 [20] encoding)	m*	m	o	o	o
S	state or province name	o	m	o	o	o
L	locality name (<i>e.g.</i> , city)	o	m	o	o	o
O	organization name	m*	-	m	m	o
OU	organizational unit name	o	-	o	o	o
CN	common name	-	m	m	-	m
T	title	-	-	o	m	-
PA	postal address (<i>e.g.</i> , street or box)	o	m	o	o	o

Figure 8. Summary of Distinguished Name attributes and their meanings. The columns labelled 1 through 5 indicate whether, for that kind of Distinguished Name, the attribute is mandatory (m), optional (o), or prohibited (-). The columns refer to the following kinds of Distinguished Names: (1) certifying authority; (2) residential person; (3) organizational person; (4) organizational role; (5) distribution list.

gorithm [33], which computes a message hash using the RSA Message Digest Algorithms RSA-MD2, and then encrypts it using the RSA encryption algorithm with the private key of the issuing authority.⁴⁸ An example certificate is shown in figure 7.

4.2.1. Identity and Distinguished Names

Both issuers and subjects must be identified unambiguously; the mechanism chosen to do this is the *Distinguished Name*,⁴⁹ which is a set of attributes containing a key and a value. An example of a typical Distinguished Name might be:

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

where “Matt Bishop” is the subject’s Common Name,⁵⁰ “Dept. of Math & CS” is the Organizational Unit name, “Dartmouth College” the Organizational name, and “US” the Country name.⁵¹ The allowed Distinguished Name attributes and what they signify are shown in Figure 8; precisely which attributes may be used depends on which kind of Distinguished Name is being used:⁵²

1. A *certifying authority* is an entity authorized to issue certificates, and the Distinguished Name of such an authority *always* appears in the issuer field. As such an authority is *never* an individual, it has no associated Common Name or Title.⁵³ Further, the Country in which the certifying Authority resides (or the Organization Name, if it is a multinational entity), must be present. Further identifying information (such as the Locality, Organizational Unit, and so forth) must be present if necessary to unambiguously identify the issuing authority.
2. A *residential person* is an entity not claiming affiliation with any organization, for example a person who obtains a certificate for private use. The Distinguished Name must unambiguously

```

30 82 01 94
30 82 01 32      { version, serial numbers (defaults 0) }
02 01 01        { serial number (1) }
30 0a           { signature information }
06 04 55 08 03 02 { algorithm (MD2 hash, RSA signature) }
02 02 02 78      { parameter (632) }
30 49           { issuer's distinguished name }
31 0b 30 09 06 03 55 04 06 13 02 55 53 { country }
13 02 55 53      { US }
31 1a 30 18 06 03 55 04 0a      { organization }
13 11 44 61 72 74 6d 6f 75 74 68 20 43 6f 6c 6c 65 67 65
{ Dartmouth College }
31 1e 30 1c 06 03 55 04 0b      { organizational unit }
13 15 44 65 70 74 2e 20 6f 66 20 4d 61 74 68 2e 20 26 20 43 2e 53 2e
{ Dept. of Math. & C.S. }
30 1a           { validity period }
17 0b 39 30 30 33 31 36 31 32 30 30 5a { from 3/16/90 at 12:00 GMT }
17 0b 39 32 30 33 31 36 31 32 30 30 5a { to 3/16/92 at 12:00 GMT }
30 5f           { subject's distinguished name }
31 0b 30 09 06 03 55 04 06      { country }
13 02 55 53      { US }
31 1a 30 18 06 03 55 04 0a      { organization }
13 11 44 61 72 74 6d 6f 75 74 68 20 43 6f 6c 6c 65 67 65
{ Dartmouth College }
31 1e 30 1c 06 03 55 04 0b      { organizational unit }
13 15 44 65 70 74 2e 20 6f 66 20 4d 61 74 68 2e 20 26 20 43 2e 53 2e
{ Dept. of Math. & C.S. }
31 14 30 12 06 03 55 04 03 13 0b { common name }
4d 61 74 74 20 42 69 73 68 6f 70 69 73 68 6f 70 { Matt Bishop }
30 59 30 0a      { subject's public key information }
06 04 55 08 03 02 { algorithm (RSA signature, MD2 hash) }
02 02 02 00      { length (512; here, the modulus size in bits) }
03 4b 00 30 48   { modulus }
02 41 00 cf 77 fd 00 aa e3 46 2c 66 ee 7c 2b fd de 2d 09 ed 2a 3e dd 3c
a1 93 98 68 f9 95 96 8e 17 9e a8 92 8c 4a 7c b8 4f 92 fe 02 7c ab
d5 09 9a ff 8d b8 1e f2 f4 80 b0 6d e8 30 fa 62 ca 09 90 f7 e5
{ 10866017727492218440042295274990800056806083243-
77396246750949131125699891639786361314357176306-
10202659399114672481140199313984031350744785286-
19036491708389 }
02 03 01 00 01 { exponent (65537) }
30 0a           { signature }
06 04 55 08 03 02 { algorithm (MD2 hash, RSA signature) }
02 02 02 78      { parameter (632) }
03 50 00 42 e3 50 db c2 1b da a1 cd 0d 2b 70 e6 66 5f 59 29 91 dc 28 76
df 2a ef a1 7a c2 7d 13 99 e0 ea e8 d4 7e 5a 18 70 0e 73 b8 c1 cc
68 98 27 a6 f2 aa f2 d5 0b 43 69 a0 64 16 92 8a 7c ed cd e3 45 b0
8e 54 a8 06 36 24 3e 04 f0 28 84 20 75 a1

```

Figure 7. An example of a certificate. The certificate is a bit stream encoded as described in section 3.1.3. For illustrative purposes, the printable encoding step of that section has been omitted, and the hex representation of the ASN.1 encoding is shown above. The signature may be checked using the issuer public key 31 86 38 84 a5 a5 6a 53 7a 8c 5c 79 80 1c 63 fd 98 51 13 36 ae 14 de 61 f8 e2 1b f0 81 a2 1c bd 02 68 6f 63 ea 2a 9f 8c c1 8d 1a 0d c6 f7 a1 0f 74 8d 9d 90 5b 0a 50 56 d0 57 93 bb 37 76 9a a9 2b ef a9 16 a8 dc 20 99 fb da 1a cb e9 91 ae 00 (hex); note that this public key is *not* RSA Data Security, Inc.'s public key, but merely one used to generate this example of a certificate.

ly identify the person by providing an address. No organizational attributes may be given.

3. An *organizational person* is an entity claiming affiliation with an organization, for example a professor whose certificate is issued by (or on behalf of) his college. In its Distinguished Name, the Common Name of the person and the Organization with which he is affiliated are required; other attributes are optional.
4. An *organizational role* is a position within an organization to which the certificate is tied; it may be filled by many people over time, but the certificate stays bound to the role and not the individual holding that position. Required attributes of this kind of Distinguished Name are the Organization and the Title of the role; other attributes except the Common Name are allowed. The Common Name is specifically disallowed to prevent confusion between this type of Distinguished Name and that of an organizational person.
5. A distribution list is a collection of users, and mail sent to it is forwarded to all those users. Its Distinguished Name requires the Common Name attribute to be the distinguished string `Distribution List` (to distinguish it from an individual), and allows any other attributes except Title to be present.

As another example, the Distinguished Name for the organizational role of postmaster might be:

`/C=US/O=USRA/OU=RIACS/OU= Network Systems Division/T=Postmaster/`

This name contains two hierarchical organizational unit names (“Network Systems Division” is a branch of “RIACS” which is in turn a research laboratory of “USRA”). If multiple organizational unit names are present, they are taken hierarchically, with the broadest unit first.⁵⁴

This information is encoded into a certificate using the ASN.1 [6][7] representation shown in X.509, Annex G, which is part of the international directory standard. This is the same format as is used for the encoding of the issuer identity in the `Originator-ID` and `Recipient-ID` lines described in section 3.2.1.

4.2.2. Certification Authorities

At the root of each certification hierarchy is a top-level certifying authority. This organization may issue certificates and it may authorize other organizations to issue certificates, imposing upon them whatever restrictions it wishes. It may also make agreements with other top-level certification authorities to allow interoperation across different hierarchies; these agreements are essen-

tially cross-certification, where each top-level certification authority generates and makes available certificates for the other.⁵⁵

Cross-certification agreements are not transitive; that is, if one top-level certifying authority cross-certifies a second which in turn cross-certifies a third, the first and third are *not* cross-certified. The reason for this restriction is that cross-certification implies satisfaction with and acceptance of legal and certification procedures of the cross-certified top-level certification authority, and that must be done with knowledge of those agreements, not by default. Hence, a certification path may have elements of at most two hierarchies.

A certifying authority is one which may issue certificates. With three exceptions to be discussed in the next section, when a certificate is issued, the issuing certifying authority is vouching for the identity as embodied in the subject's Distinguished Name; that is, if an organizational affiliation is shown, such affiliation has been proved, and the public component as embodied within the certificate belongs to that subject. How such proof is given is up to the certifying authority, but may be affected by contracts or licenses with the top-level certifying authority.⁵⁶

A key component of acquiring this proof is the *organizational notary*.⁵⁷ The organizational notary is a clearinghouse for certificate orders within an administrative domain such as an organization or organizational unit, and is assumed to be somewhat independent of the users in that administrative domain. Only those users in the domain may order certificates through that organizational notary, who accepts and validates the information to be put in the certificate; the manner in which this is done is up to the organization, except for certain minimum requirements which the top-level certifying authority may establish as a condition of making the organization a certifying authority. If necessary, the notary may alter the period of validity (to comply with the policies of the certifying authority), the serial number (to ensure it is unique among all those issued by the certifying authority), or any other field in the certificate except the user's personal name (in the Distinguished Name) and the user's public key. Under normal circumstances, though, only the serial number and validity period are likely to be changed.

The use of an organizational notary has benefits both to the user and to the organization. First, it enables the organization to validate those to whom it issues certificates quickly and easily. Secondly, the user need not go to an outside party to prove his or her association with the organization; it can be handled in-house. Third, should the user leave the organization, his or her affiliation can be repudiated by revoking the certificate. The benefits to the certifying authority are also large, residing mainly in the ability to delegate the job of verifying the information provided by the

user. Since there will be far fewer organizational notaries than users, they can be screened more carefully.⁵⁸

A certifying authority's private key is used to sign certificates and is *never* used as a component of an interchange key, hence a certifying authority's certificate is useful only for validation purposes. This limits the damage should such a private key be compromised; while certificates could be forged (until the corresponding public key certificate were revoked), privacy-enhanced messages protected using certificates issued before the compromise by the compromised certifying authority would still be protected.⁵⁹ (Of course, an attacker could use the private key to generate forged certificates and date them before the compromise; hence, all certificates issued by the compromised certifying authority should be considered suspect when its private key is compromised.)

Finally, even though there may be other certifying authorities between the issuer and the top-level certifying authority, all certificate hierarchies follow the convention that each certifying authority's certificate is signed by the top-level certifying authority. This shortens the certificate validation procedure, described below.⁶⁰

4.2.3. Conventions and Special Certificates⁶¹

Under normal circumstances, a person obtaining a certificate with an associated organization in the subject name is presumed to be closely affiliated in some way with that organization. Under certain circumstances, such affiliation may not be desired, for example if the person is visiting for a limited period of time. In this case, the certifying authority may issue a *guest certificate* in which the final attribute of the issuing authority name is an Organizational Unit name with a value the distinguished string Guest. For example, if the certifying authority Dartmouth College issues a certificate to visiting professor Tom Jones, the issuing authority name would be

/C=US/O=Dartmouth College/OU=Guest/

and the subject name would be

/C=US/O=Dartmouth College/OU=Guest/OU=Dept. of Math & CS/CN=Tom Jones/

Note that a guest certificate affirms the identity of the user; it merely asserts a weak, rather than a strong, affiliation with the organization. Of course, the certifying authority cannot issue such a certificate for a member of another organization.

A similar type of certificate, the *notary certificate*, asserts that the user's identity has been proved to the certifying authority but that no organizational affiliations are present. This is most

<i>certificate</i>	<i>degree of affiliation</i>	<i>certainty of identity</i>
persona	none	none
notary	none	strong
guest	tenuous	strong
other	strong	strong

Figure 9. How the different types of certificates bind subject (user) identity and affiliation with organizations named in the subject's Distinguished Name. The entries indicate the representations being made by the certifying authority in each case; so, for example, in a notary certificate, the certifying authority has verified the user's identity but makes no claims about any organizational affiliations.

useful when certificates are issued to residential persons, and in fact notary certificates may only be used for that type of entity; they may never be used for an organizational person or role (nor certifying authority or distribution list). In this case, the final attribute of the issuing authority name is an Organizational Unit name with a value the distinguished string `Notary`. For example, if Dartmouth College issues a residential certificate to Tom Jones, the issuing authority name would be

/C=US/O=Dartmouth College/OU=Notary/

and the subject name might be

/C=US/S=New Hampshire/L=Hanover/CN=Tom Jones/

Under other circumstances, holders of certificates may wish to remain anonymous, but ensure that recipients of successive letters be aware that the letters came from the same sender, and to ensure integrity and/or privacy; this essentially preserves the anonymity currently provided by the ability to use arbitrary electronic mail names. One possible use may be to report a security problem anonymously. A certifying authority may issue *persona certificates* in which it explicitly is not vouching for the identity of the user. The final attribute of the issuing authority name in such a certificate is an Organizational Unit name with a value the distinguished string `Persona`. To continue the example, if Certificates, Inc. issued a persona certificate, the issuing authority name would be

/C=US/O=Certificates, Inc./OU=Persona/

In such a certificate, the subject name should be considered arbitrary and in no way reflective of the identity of the user to whom the certificate was issued.

Figure 9 summarizes the representations of the certifying authority with respect to subject identity and organizational affiliation for these types of certificates.

4.2.4. Certificate Validation

Checking the validity of a certificate is straightforward. As noted above, each certificate contains an entity name, a public key, an issuer name, a digital signature created by the issuer based on its private key and the information in the certificate, and some other information. To validate that the certificate has not been altered, one need only obtain the issuer's certificate, extract the public key, and use that to decrypt the digital signature on the certificate in question. If this differs from the (locally computed) hash of the certificate in question, that certificate is bogus. The certificate of the issuer may be checked using the same procedure once the certificate of the top-level certifying authority (which will be widely publicized) is obtained.⁶²

Note that the convention of having the top-level certification authority certify all other certification authorities in its domain means that within a domain, validation requires checking two certificates. If the original certificate came from another domain which is cross-certified by the local domain's top-level certification authority, then a third step is necessary: the certificate of the remote top-level certification authority must be checked by using the local top-level certification authority's certificate. So in the worst case, at most three certificates must be validated. Keeping the validation path so short was a primary consideration in defining the convention that the top-level certification authority certifies all certification authorities in its domain.

4.2.5. Certificate Revocation

Each certificate is issued for a limited period of time, and – more importantly – may be invalidated or compromised in a number of ways. For example, the entity name might change, or the private key associated with the public key in the certificate may be revealed. The international standard X.509 places the responsibility for maintaining time-stamped lists of revoked certificates, as well as revoked certificates representing certifying authorities, upon the certifying authority issuing the certificates. So the serial number of the compromised certificate is added to a revoked certificate list maintained by the certifying authority issuing the certificate. This list is signed by the private key of the certifying authority, is dated, and contains the date at which the next revoked certificate list will be issued. Although an addition to the format prescribed by X.509, this date enables a site to determine if the revoked list is the most recent one from that certifying authority. Note that this requires a new list to be issued at the stated time even if no certificates have been revoked; however, it also allows a site to determine if the list is out of date.⁶³

Top-level certifying authorities are required to establish a database which maintains certificate revocation lists for all certification authorities in its domain, and that is accessible through electronic mail. Further, all certifying authorities must transmit their current list to the top-level certifying authority as well as to users and user agents within their domain.

Transmission of the lists may be done in a number of ways. The currently-defined mechanism is to use privacy-enhanced mail to propagate these certificate revocation lists. Such a letter contains three types of header fields: a `Certificate` field containing the certificate of the certifying authority issuing the list, any number of `Issuer-Certificate` fields, and a `CRL` field:

`CRL: crl_encoded`

containing the certificate revocation list `crl_encoded` encoded using the printable encoding defined in section 3.1.3. Future mechanisms may include the use of a USENET newsgroup.

4.3. Example of a Certifying Authority Hierarchy: RSA Data Security Inc.

The use of this certification mechanism requires a public key cryptosystem which must meet several conditions:

1. it must provide both privacy and authentication;
2. it must be (relatively) efficient to implement in either software or hardware;
3. it must expand the encrypted information as little as possible;
4. it must be (or at least be believed to be) cryptographically strong.

In addition to meeting all these requirements, the RSA cryptographic algorithm [37] is also the primary algorithm recommended for use in international standards requiring (or recommending) use of a public key cryptosystem. So the initial asymmetric cryptographic algorithm defined in the proposal is the RSA algorithm.

The basic requirements for the RSA algorithm in this context are that the modulus size vary between 508 and 1024 bits (or approximately 8.4×10^{152} and 1.8×10^{309}),⁶⁴ and that the public exponent be either 3 or $2^{16}+1$.⁶⁵ The former requirement is intended to make determining the private key acceptably hard, and the latter to allow the public key to be the modulus alone rather than the modulus and the public exponent; the algorithm identifier will indicate which of the two is being used. The public exponent $2^{16}+1$ was selected because it allows relatively efficient processing⁶⁶ and it is recommended by X.509, Annex G. The exponent 3 was chosen because exponentiation is even faster. Finally, if the message integrity check must be padded, it is to be padded on the left

with the ASN.1 encoding of the identifier for the integrity checking algorithm and then with zeroes, and if the data encryption keys must be padded, they are to be padded with zeroes unless the message is addressed to multiple recipients and the public exponent is 3 [31]. In this specific case, a pseudorandom 64-bit quantity is to be generated for each recipient and as many copies as will fit are to be placed to the left of the data encryption key before encrypting.⁶⁷

The RSA algorithm is covered within the United States by patents administered by RSA Data Security, Inc. (These patents do not apply outside the United States or to the United States government.) Currently RSA Data Security, Inc. plans to act as a top-level certifying authority, and the certificates it issues include a license to use the RSA algorithm for certificate validation and encryption and decryption operations to send privacy-enhanced electronic mail.⁶⁸

4.3.1. Issuing Organizational Certificates

A top-level certifying authority must provide for the certification of lower-level certifying authorities, and the approach taken by RSA Data Security, Inc. indicates the care that this process requires to ensure private keys are not compromised. There are two scenarios possible, one in which the organization signs its own certificates and the other in which RSA Data Security, Inc. signs certificates on behalf of the organization. In either case, the organization is the certifying authority. In either case, these organizations must first establish with RSA Data Security, Inc. an appropriate pair of keys, called "organization keys," to sign certificates issued by the organization. Either RSA Data Security, Inc. will hold the private key and sign certificates for that organization, or the organization will hold the key and use it to generate and sign the certificates it issues. The method for obtaining a certificate for oneself depends on where the organization private key is held.

4.3.1.1. RSA Data Security, Inc. Holds the Organization's Private Key

If RSA Data Security, Inc. is to sign a certificate on behalf of the organization, the user must first generate a public key and private key pair (possibly by using special-purpose software), and include the public key in a certificate. He then constructs a prototype privacy-enhanced electronic mail message with the certificate signed using his own private key, and sends this to his organizational notary. The organizational notary accepts and validates the information in the electronic letter, and then forwards the electronic message to RSA Data Security, Inc. by privacy-enhanced electronic mail; this letter indicates that the organizational notary vouches for the correctness and integrity of the information, and authorizes RSA Data Security Inc. to sign a certificate on behalf of the organization.

When RSA Data Security, Inc. receives the electronic and paper letters and its fee, and is convinced all is in order, it issues a certificate signed with the private key of the organization.⁶⁹ This certificate will be valid for two years, at which point a new certificate must be acquired.

Not surprisingly, the fee required when RSA Data Security, Inc. signs certificates on behalf of the organization is the most controversial point of the scheme; while the fee seems small, if a computer installation has several thousand users and wishes to obtain certificates for all, the cost grows rather quickly. Possible solutions include having the users buy their own certificates (which is reasonable since those who do not can still use regular electronic mail), or buying certificates only for those for whom the installation deems them a necessity.

4.3.1.2. The Organization Holds its Organization Private Key

Some organizations may object to the loss of autonomy entailed by giving their private key to RSA Data Security, Inc., or may find the fee for the certificate generation service not to be cost-effective. An alternate arrangement involves attaching to a workstation (or other computer) a device known as a *certificate signature unit*.⁷⁰ This tamperproof device is like postage meters but signs certificates. Organizations using these devices are their own *issuing authority*.

Figure 10 summarizes the protocols that the certificate signature unit uses. They were designed with two goals in mind: first, that the issuing authority's (organization's) private key never be transmitted to anyone; and second, that RSA Data Security, Inc. be able to control the number




Figure 10. Protocols for the certificate signature unit. The CSU exchanges public information about the issuing authority with RSA Data Security, Inc., using a channel on which authentication only is done. The issuing authority's private key is encrypted using a symmetric cryptosystem, once with the CSU's secret key and then with the IA's secret key, which is stored on a removable key storage device. The doubly-encrypted private key is then stored on a second removable key storage device. Both must be inserted into the CSU before any certificates may be signed.

of certificates issued (for royalty purposes). To ensure these two goals, the messages sent over the communications channel between the certificate signature unit and RSA Data Security, Inc. are authenticated (so they cannot be altered in transit) but not encrypted (so the issuing authority can be convinced that the certificate signature unit is not sending the issuing authority's private key to RSA Data Security, Inc.); hence, this channel is called the *authenticated channel*. To prevent the issuing authority private key from having to be entered manually or downloaded from another computer (doing either would ensure the key would be written down or stored in a memory from which it could be read), it is encrypted and stored on a removable key storage device that can be inserted (and removed) from the signature unit. When the key is stored, it passes over a communications channel secured by symmetric cryptography (and called the *secure channel*); associated with the unit is the certificate signature unit secret key, used to encrypt the issuing authority private key before it is stored on the removable key storage device. Note that this binds the device to the particular unit. Physical security practises dictate that this information be protected in such a way that the loss of the removable key storage device will not compromise the organization, so before being stored on the device, the encrypted issuing authority private key is again encrypted using a secret key associated with the issuing authority. This key is also stored on another removable key storage device. Note that both removable key storage devices must physically be inserted into the certificate signature unit before any certificates may be signed, and the device containing the issuing authority's private key must be present when each certificate is signed.

When an organization buys a certificate signature unit, the meter comes with the *certificate signature unit private key* (used for the authenticated channel), the *certificate signature unit secret key* (used for the secure channel), and RSA Data Security, Inc.'s public key stored in tamperproof memory. The organization must now generate an *issuing authority private key* and the associated public key. First, the unit generates an *issuing authority secret key* (for the secure channel) and saves it on one removable key storage device. It then generates an RSA key pair (to be used as the issuing authority private and public keys), signs the new public key with the certificate signature unit private key, and returns the result to the workstation it is attached to. It then encrypts the new private key with the certificate signature unit secret key and the issuing authority secret key, and stores the result on a second removable key storage device. The organization can verify the meter generated the issuing authority's key pair by using the generation unit's public key to validate the issuing authority's public key), and at no time does the organization itself (or RSA Data Security,

Inc.) have access to its issuing authority private key except through using the removable key storage devices. This prevents anyone from issuing certificates without proper authorization.

That authorization is provided both by the organization, upon insertion of both removable key storage devices into the certificate meter, and by RSA Data Security, Inc., which sends the issuing authority an authorization message good for some limited number⁷¹ of certificates. These authorization messages are transmitted to the certificate signature unit as they arrive. To issue a certificate, the organization first ensures the unsigned certificate resides on the attached workstation. The issuing authority private key is supplied by inserting the key storage devices into the meter, which then decrypts the private key. Then for each user, the issuing authority sends the meter an unsigned certificate; the meter checks that the serial number of the certificate sequentially follows the serial number of the previous certificate, that the certificate is in the correct ASN.1 encoding, that the subject belongs to the organization issuing the certificate, and that the Distinguished Name of the organization is the same as in the authorization message; if all these conditions hold, the unit will sign the certificate with the organization private key and return the result to the issuing authority.

In addition to charging for the certificate meter, RSA Data Security, Inc. will charge a much smaller royalty per certificate than if RSA Data Security, Inc. generated the certificate itself. The reason for the difference in cost lies in the administrative overhead for RSA Data Security, Inc. With the certificate signature unit, that organization need only supply the software and box once. However, if RSA Data Security Inc. manages an organization's private key and generates certificates on behalf of that organization, it will have overhead of extra personnel and equipment to run the software and to provide adequate management and protection for those keys and software. Hence the expense of having RSA Data Security, Inc. create the certificates is greater than using the certificate signature unit.

5. Implementation Requirements⁷²

Although the privacy-enhanced electronic mail proposals do not mandate details of implementation of the protocols, they do place certain constraints on what conforming implementations must do. The two most important restrictions are upon the processing of nested privacy-enhanced messages and error handling.

The proposals require conforming implementations to handle both recursively nested messages and serially nested messages; so, for example, a privacy-enhanced message may be placed

in the body of a second message and forwarded, and that message may itself be put into the body of a third message and forwarded, and so on. Similarly, a set of privacy-enhanced messages may be placed one after the other inside the body of a privacy-enhanced message. Precisely how the implementation handles such cases is up to the implementor, and one conformant prototype implementation simply requires the user to run a text editor on the message to extract each enclosed privacy-enhanced message and store it in its own file; each file may then be checked (and decrypted, if necessary). But some mechanism for doing this must be available and usable with the privacy-enhanced user agent.

More interesting are the requirements for handling errors. Syntax errors in privacy-enhanced messages should be flagged, as should messages which yield an incorrect message integrity checksum. If the message is not of the type MIC-CLEAR, the latter case raises a very serious user interface issue: should the user be shown a message which may be useful even in its altered state? For example, if a character was deleted due to network transmission problems, the rest of the message may give the user enough indications to be able to reconstruct the sender and contact him for more information. On the other hand, if the entire message is forged, the user can place no reliance on it at all. The privacy-enhanced mail proposal strikes a balance between not showing the letter at all and showing the letter with a warning that the user may ignore (or not notice): it requires that the user be informed of the error and *then* be forced to ask that the relevant message be shown. This positive indication is the best technique to ensure the user is aware of the problem and accepts the risks of reading the message (and acting on it, if he does).

In any case, a user agent conforming to the requirements of the privacy-enhanced mail protocols must display the full distinguished name of the sender; if the message is encrypted, the full distinguished name of each recipient must also be shown. Further, if any certification path involves cross-certification of top-level certifying authorities, this must be made apparent to the user, as must the use of any persona certificates. The precise mechanism is up to the implementor.

Finally, an error may occur because a privacy-enhanced letter is mis-addressed. In this case, most message transport agents will return the letter to the originator. However, if the letter is encrypted, the sender could not read the contents, because he has none of the recipients' private keys. So, in order to enable the sender to process any encrypted message which is returned as undeliverable, all privacy-enhanced mail implementations must allow the sender to insert a Key-Info header field after the Originator-ID header field. This line may then be used to decrypt the message should it be returned.

6. A Symmetric Key Distribution Mechanism

Although the privacy-enhanced mail proposal recommends interchange keys be distributed using certificates, alternate key distribution methods may be used with it. In particular, given a trusted third party (server), symmetric cryptography may be used.

The simplest such configuration [32] works by having each user establish a secret key known only to her and to the key distribution center. When one user wishes to send privacy-enhanced electronic mail to another, the sender obtains a data exchange key, encrypts the message in the normal manner, then encrypts the data exchange key and message integrity check using her interchange key, and transmits them to the trusted server. The server decrypts the two and encrypts them using its own key, along with the sender's identity. These are returned to the sender, who then inserts them into her message and sends the message. The recipient relays the information in the Key-Info and Sender fields to the trusted server, which decrypts the contents using its key and then encrypts them using the recipient's interchange key. When the recipient gets the result from the server, he can decrypt the message in the usual way.

The main problem with this scheme is that, in addition to trusting the user agents, the key distribution center must also be trusted; if it is not, the scheme is worthless. Since the proposal is intended as an Internet standard, it cannot realistically expect an entity on the internet to trust servers *not* under its control. Hence, the key distribution proposal in support of privacy-enhanced electronic mail uses a public key cryptosystem and a mechanism that does not require trusting a single host.

Notice that trusting a certifying authority and trusting a server of the sort described here are very different. If the server is compromised, any of its subscribers can be impersonated and existing privacy-enhanced electronic letters read, as well as all letters sent after the compromise. However, if a certifying authority is compromised (by someone obtaining the secret keys it signs certificates with) the only damage is that new certificates can be forged, users impersonated, and electronic mail sent using those certificates can be read; electronic mail sent with certificates existing at the time of compromise is still protected, as are all letters sent before the compromise. In short, compromising a trusted server compromises everything; compromising a certifying authority does not compromise previously sent messages.

A second problem with trusted servers springs from the use of multiple servers in different domains introduces latencies in negotiating exchange keys. If someone at Dartmouth were to send

a letter to someone at Purdue, the Dartmouth and Purdue servers would have to coordinate the selection of the interchange key. This would have to be done in a very timely fashion to prevent delays in the message transmission; this suggests replication of data on different servers, once again raising the issue of trusting a server not under local control. The alternative is to accept that encrypted mail might be delayed due to reasons beyond the control of the message transport agents.

We should also note that this server does more than provide an authentication service: it provides a *digital signature service*. An authentication service confirms the identity of the user, and may provide credentials attesting to the authentication. However, a digital signature service not only authenticates the user but also binds the user's identity to the message being sent. An authentication tool would therefore be quite unsuitable for this scheme; in particular, the functionality of the authentication server *Kerberos* [40] would need to be radically expanded to make it suitable.

7. Forwarding Messages and Mailing Lists

Privacy-enhanced electronic mail can be forwarded, but if asymmetric interchange keys are used there is a subtlety. If the message is of type MIC-ONLY or MIC-CLEAR, the message integrity checksum in the MIC-Info header field is decrypted using the sender's public component, and re-encrypted using the recipient's private component; the message may then be forwarded; as an alternative, it can simply be placed within the body of another privacy-enhanced mail message. If the message is of type ENCRYPT, however, it cannot be forwarded encrypted. The only possible approach (the recipient decrypts the data encryption key, uses that key and the originator's public component to decrypt the message integrity checksum, encrypts the message integrity checksum using his private component and the data encryption key, and then encrypts the data encryption key using the public component of the party to whom the message is to be forwarded) fails because the originator is now lost; the message appears to have originated at the forwarder. Note also that in this case the message cannot simply be enclosed in another privacy-enhanced message, as the decryption requires the private key of the forwarder.

If symmetric interchange keys are used, the procedure is similar, except that the transformations involve agreed-upon secret keys, and in neither case can the message be forwarded simply by enclosing it in another privacy-enhanced mail message.

Sending a privacy-enhanced message to multiple recipients is straightforward and needs not be elaborated upon. However, the case when those multiple recipients are part of a mailing list (in which a letter sent to a single address is *exploded* or forwarded to multiple recipients) does.

In some cases, the membership of a mailing list may not be known or available to a sender for a variety of reasons. The host on which the alias is exploded may not be willing to reveal that information. The mailing list may itself contain mailing lists (for example, the mailing list *csnet-forum@relay.cs.net* has as an address *csnet-forum@dartmouth.edu*, which is itself a mailing list). In this case there must be an interchange key associated with the list; then all letters sent to the list are simply forwarded to each member of the list as described above.⁷³

If the sender can determine the membership of the mailing list, then a separate list interchange key is unnecessary; the sender can simply insert the destinations into the letter using multiple *Recipient-ID* and *Key-Info* fields. The message would be encrypted using the same data exchange key for all recipients and hence the encryption and integrity check need be done just once; however, the information in each *Key-Info* field would be encrypted using the associated recipient's public key.

If the certificate-based key distribution mechanism described above is used and the interchange keys are RSA keys with public exponent 3, note that even though there is a single recipient address there are multiple recipients. Hence the data exchange key should be padded with a 64-bit pseudorandom quantity, as discussed above.

8. Conclusion

The above two protocols satisfy the constraints and recommendations presented in the background section. Encapsulating the privacy-enhanced message renders it invisible to the message transport agents, so no transport-level protocols need to be changed; the message may be processed by a special program and then included in the body of a conventional electronic mail message, so only that part of the user agent providing the privacy enhancements need be trusted with special security information, and conventional (non-enhanced) electronic mail is unaffected by the presence (or absence) of these enhancements. Integrity and authenticity are assured so long as the interchange keys are not compromised, and if the message is encrypted it will also be private. Finally, since the protocol for privacy-enhanced mail allows the use of any key management system, it is flexible enough to be used in a multitude of environments.

All this suggests an obvious architecture for implementing privacy-enhanced electronic mail. A special program to enhance messages runs above the user agent, taking as input a message and producing as output the transformed result. The user then includes this in a conventional mail message using any user agent desired. The recipient reverses the process by extracting the message

from his mail message using his user agent, and then passing the encapsulated part to a second program which reverses the transformation and produces a local representation of the initial message.

The advent of personal computers and workstations has led many facilities to use central servers as mail hosts, so users can download their messages from their platform to the central host for transmission [12]. The above architecture fits into this scheme nicely; the privacy-enhancing software need reside only on the user's platform. This was no small consideration in the design of the protocol.

Both the X.411 Recommendation and MSP were designed for different requirements, and do not satisfy the same constraints as the privacy-enhanced mail protocols. Specifically, the X.411 Recommendation requires each message transport agent to be able to parse the headers containing the security parameters; this would require altering existing message transport agents. By way of contrast, the Message Security Protocol is similar to privacy-enhanced mail, except that it uses the X.400 protocols rather than SMTP as its basis. [21]

The reasons for selecting a patented cryptosystem as the public key cryptosystem to be used in the privacy-enhanced mail protocol have been explained in section 4.3. It is perfectly possible to use some other public key cryptosystem (or cryptosystems) to generate keys; however, factors such as ciphertext expansion (if the size of the interchange key or the size of the encrypted data encryption key becomes several thousand bits, the overhead would become prohibitive) and the need for software may constrain this option. Further, if only one public key system is used, it must provide both authentication and privacy, and if more than one system is used (for example, one to provide secrecy and another to provide authenticity), then the key distribution and management scheme must manage two sets of keys per user.

We must emphasize that the privacy-enhanced electronic mail protocol and the certificate-based key management protocol are distinct; one is free to adopt the first without using the second. Indeed, [26] specifically describes protocols to be used with key management schemes other than certificate-based schemes, and states that "the message processing procedures can also be used with symmetric key management."⁷⁴ Organizations that decide not to use the public key approach may substitute their own key management scheme; however, to be compatible with other implementations, all implementations of privacy-enhanced electronic mail should support the certificate-based approach.⁷⁵ This also means that the certificate-based key management protocols may be used in contexts other than privacy-enhanced mail [4]; since its infrastructure is similar to that of X.400, it can be used to transition to that, and related, standards.

More research in cryptography would aid in the maturation of this proposal. Specifically, one-way hash functions such as RSA-MD2 that can be used to compute a message integrity check are very few; more are needed. A public key cryptosystem as strong as RSA would allow the use of interchange keys not encumbered by licensing. If such a cryptosystem could be implemented as efficiently as the DES, it could be used to encrypt the message as well.

Finally, this proposal does not address issues in network and system security, such as the development of trusted software, routing controls, replays, and access controls. While all are important to the sending of electronic mail, they have much wider applications, and provide a fertile field for research and development.

Acknowledgments. This paper presents a set of protocols developed during a series of meetings of the IAB's Privacy and Security Research Group, and stated in a series of RFCs. Thanks are due to David Balenson, Curt Barker, Jim Bidzos, Danny Cohen, Tom Daniel, Charles Fox, Morrie Gasser, Stephen Kent, John Laws, John Linn, Steve Lipner, Ralph Merkle, Dan Nessel, Mike Padlipsky, Ken Rossen, Rob Shirey, Miles Smid, Dave Solo, Steve Walker, and Steve Wilbur. Special thanks to Burton Kaliski, Jr., of RSA Data Security, Inc., for information on how RSA Data Security Inc. would be handling the certification process, and to him, Ken Rossen, and Dave Solo on the certificate postage meters, and to David Balenson for generating the sample message for symmetric interchange keys. Also, thanks to Donald Johnson, Burton Kaliski, Jr., John Linn, Evi Nemeth, Dan Nessel, J. Shallit, and especially David Balenson for their constructive comments on an earlier draft, to Joan Feigenbaum of AT&T Bell Laboratories for encouraging me to write this paper, and to the participants at the DIMACS Workshop on Distributed Computing and Cryptography, without whose probing questions (and spirited discussion) many of the parts of this paper would be a good deal less complete and more confusing.

References

- [1] S. Akl, "On the Security of Compressed Encodings," *Advances in Cryptology: Proceedings of Crypto 83*, Plenum Press, New York, NY (1984) pp. 209-230.
- [2] ANSI X9.17-1985, *American National Standard Financial Institution Key Management (Wholesale)*, American Bankers Association (Apr. 1985).
- [3] W. Barker, P. Cochrane, and M. Branstad, "Embedding Cryptography into a Trusted Mach System," *Proceedings of the Fourth Aerospace Computer Security Applications Conference*, p. 379-383 (Dec. 1988).

- [4] M. Bishop, "An Authentication Mechanism for USENET," *Winter 1991 USENIX Proceedings* pp. 281-287 (Jan. 1991).
- [5] T. Casey and S. Wilbur, "Privacy Enhanced Electronic Mail," *Proceedings of the Fourth Aerospace Computer Security Applications Conference*, pp. 16-21 (Dec. 1988).
- [6] CCITT Recommendation X.208, *Specification of Abstract Syntax Notation One (ASN.1)* (1988).
- [7] CCITT Recommendation X.209, *Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1)* (1988)
- [8] CCITT Recommendation X.400, *Message Handling System Model* (1984).
- [9] CCITT Recommendation X.411, *Message Handling Systems: Message Transfer System: Abstract Service Definition and Procedures* (1988).
- [10] CCITT Recommendation X.500, *The Directory - Overview of Concepts, Models, and Services* (1987).
- [11] CCITT Recommendation X.509, *The Directory - Authentication Framework* (1987).
- [12] D. Clark and M. Lambert, *PCMAIL: A Distributed Mail System for Personal Computers*, RFC-993 (Dec. 1986).
- [13] D. Crocker, *Standard for the Format of ARPA Internet Text Messages*, RFC-822 (Aug. 1982).
- [14] D. Denning, *Cryptography and Data Security*, Addison-Wesley, Reading, MA (1982).
- [15] D. Dern, "The Trusted Mail System," *ConneXions – The Interoperability Report* 4(2) p. 8 (Feb. 1990).
- [16] Federal Information Processing Standards Publication 46, *Data Encryption Standard* (Jan. 1977).
- [17] Federal Information Processing Standards Publication 81, *DES Modes of Operation* (Dec. 1980).
- [18] Federal Information Processing Standards Publication 113, *Computer Data Authentication* (May 1985).
- [19] Federal Information Processing Standards Publication 146, *Government Open Systems Interconnection Profile (GOSIP)* (Apr. 1989).

- [20] ISO 3166, *Codes for the Representation of Names of Countries* (1987).
- [21] R. Housley, "Electronic Messaging Security: A Comparison of Three Approaches," *Proceedings of the Fifth Annual Computer Security Applications Conference*, p. 29 (Dec. 1989).
- [22] B. Kaliski Jr., *private communication* (Mar. 1990)
- [23] S. Kent and J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part II -- Certificate-Based Key Management*, RFC-1114 (Aug. 1989).
- [24] S. Kille, *Mapping Between X.400(1988) / ISO 10021 and RFC 822*, RFC-1148 ().
- [25] A. Lenstra and M. Manasse, "Factoring by Electronic Mail," to appear in *Advances in Cryptology – EUROCRYPT '89*, Springer-Verlag Berlin (1990).
- [26] J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part I -- Message Encipherment and Authentication Procedures*, RFC-1113 (Aug. 1989).
- [27] J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part III -- Algorithms, Modes, and Identifiers*, RFC-1115 (Aug. 1989).
- [28] J. Linn and S. Kent, "Electronic Mail Privacy Enhancement," *AIAA/ASIS/DODCI Second Aerospace Computer Security Conference*, pp. 40-44 (Dec. 1986).
- [29] J. Linn and S. Kent, "Privacy for DARPA-Internet Mail," *Proceedings of the Twelfth National Computer Security Conference*, pp. 215-229 (Oct. 1989).
- [30] C. Mitchell, D. Rush, and M. Walker, "A Remark on Hash Functions for Message Authentication," *Computers and Security* 8(1) pp. 55-58 (Feb. 1989).
- [31] J. Moore, "Protocol Failures in Cryptosystems," *Proceedings of the IEEE*, 76(5) pp. 597 (May 1988).
- [32] R. Needham and M. Schroeder, "Using Encryption for Authentication in Large Networks of Computers," *Communications of the ACM*, 21(12), pp. 993-999 (Dec. 1978).
- [33] NIST Special Publication 500-183, *Stable Implementation Agreements for Open Systems Interconnection Protocols*, Version 4, Edition 1, Part 11 (Dec. 1990).
- [34] D. Nowitz, and M. Lesk, *A Dial-Up Network of UNIX Systems*, document SMM:21 in *UNIX System Manager's Manual*, Computer Systems Research Group, University of California, Berkeley, CA 94720 (Apr. 1986). Reprinted by the USENIX Association.

- [35] J. Postel, *Simple Mail Transfer Protocol*, RFC-821 (Aug. 1982).
- [36] R. Rivest, *The MD4 Message Digest Algorithm*, RFC-1186 (Oct. 1990).
- [37] R. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," *Communications of the ACM*, **21**(2) pp. 120-126 (Feb. 1978).
- [38] M. Rose and E. Stefferud, *Proposed Standard for Message Encapsulation*, RFC-934 (Jan. 1985).
- [39] SDNS Protocol and Signaling Working Group, *SDNS Message Security Protocol*, SDN.701 Revision 1.5 (Aug. 1989)
- [40] J. Steiner, C. Neuman, and J. Schiller, "Kerberos: An Authentication Service for Open Network Systems," *USENIX Conference Proceedings*, pp. 191-202 (Winter 1988).
- [41] *UNIX User's Manual Reference Guide*, Computer Systems Research Group, University of California, Berkeley, CA 94720 (Apr. 1986). Reprinted by the USENIX Association.
- [42] V. Voydock and S. Kent, "Security Mechanisms in High-Level Network Protocols," *Computing Surveys* **15**(2) pp. 135-171 (June 1983).

Footnotes

- 1 Author's affiliation: Department of Mathematics and Computer Science, Dartmouth College, Hanover, NH 03755. Participation in the series of meetings leading to these proposals was supported by grant NCC2-397 from the National Aeronautics and Space Administration to the Research Institute for Advanced Computer Science, by grants NAG2-480 and NAG 2-628 from the National Aeronautics and Space Administration to Dartmouth College, and by a Burke Award from Dartmouth College. A preliminary version of this paper was presented at the DIMACS Workshop on Distributed Systems and Cryptography at Princeton, NJ, Oct. 4-6, 1989.
- 2 Footnotes will describe the specification in the appropriate RFC, and indicate the reason for the change. Revised RFCs have been submitted, and will be released soon; changes have been made to the existing RFCs are presented in the text and annotated in the footnotes.
- 3 UNIX is a Registered Trademark of AT&T Bell Laboratories.
- 4 Strictly speaking, this is a poor assumption; however, it serves to separate the issues involved in the security of a computer and its attendant software from the security enhancements needed to protect and authenticate an electronic mail message in transit.
- 5 The Mail(1) program in Berkeley UNIX is a good example of this [41].
- 6 This differs from [26], §4.4, which uses the same line to begin and end an encapsulation.
- 7 Note that these follow the revisions to the RFCs; in particular, in [23] and [26], all encapsulated header fields began with X- to indicate they were experimental and non-standard, as required by [13]. Since they are now part of a draft standard and no longer experimental, the X- will be dropped. Also, the version number has changed, and the Sender-ID header is now Originator-ID.
- 8 [26], §4.4, suggests that under some circumstances replicating header fields for authentication purposes in the body is possible. This has been dropped.
- 9 [35], §4.5.
- 10 [26], §4.3.2.2.
- 11 The only difference between this and the SMTP representation is the dot-stuffing transformation ([35], §4.5.2), in which lines consisting only of a period "." have a second period added. This is unnecessary since the purpose of the transformation is simply to ensure a common rep-

resentation of each character.

- 12 [26], §4.3.2.3.
- 13 Other integrity check algorithms may be added later; see [27], §1 and §4. For example, RSA-MD4 was added after [27] was issued.
- 14 [26], §4.6.1.1. This section also provided for partial encryption, in which only portions of the message were encrypted. Partial encryption has since been dropped from the protocol.
- 15 Again, other message encryption algorithms may be added later; see [27], §1 and §2.1.3.
- 16 This differs from [26], §4.3.2.3, which required padding with octets of all bits set; the high-order bit being set disambiguated the padding from the message. However, as the message integrity checksum (which contains arbitrary binary data) is also encrypted using the DES in CBC mode, another form of padding must be defined for it. The new padding can be used for both the message and the message integrity checksum.
- 17 [26], 4.6.1.1. That section provides only for the field values MIC-ONLY and ENCRYPTED.
- 18 [26], §4.6.1.2.
- 19 [26], §4.3.2.4.
- 20 This is actually a subset of the International Alphabet IA5; the elements of this subset are represented identically in IA5 and ASCII. See [26], p. 13.
- 21 [26], §4.3.2.4 also provided that, if the message were partially encrypted, an asterisk "*" was to be placed before and after the character sequences corresponding to cleartext regions. As partial encryption has been dropped, so has this character.
- 22 This processing mode is new and not described in [26], 4.6.1.1.
- 23 [26], §4.6.2.1, §4.6.4.1.
- 24 In [23], the Originator-ID header field was called the Sender-ID header field, but the name was changed to distinguish it from the more common Sender field ([13], §4.4.2).
- 25 [26], §5.2.1.1, §6.1.
- 26 This has been added to conform to the normal way of processing host names in the Internet.
- 27 [26], §5.2.1.3.
- 28 The common name attribute (CN) should be omitted, as the issuing authority is always an organization or an organizational unit. Originally, any unique name could be used, but in an in-

ternet without a central administrator, this could lead to ambiguities. Also, in [26], the Distinguished Name had to be written using the method in [24] before being put into ASN.1; this requirement has been dropped as unnecessary.

29 [26], §5.2.1.3. Originally any disambiguating string could be used for any scheme, but since
certificates are issued with a serial number unique to each issuing authority, it is simplest to
make the identifying string identify the precise certificate. Note that certificates contain expi-
ration times, so those are as readily available as if they were given as the identifying string.

30 [26], §4.6.2.1.

31 [26], §5.2.

32 [26], §4.6.4.2.1.

33 [27], §1, §2.1.1, and §2.1.2. Other data exchange key encryption algorithms may be added lat-
er.

34 [27], §1, §4. Again, other message integrity check algorithms may be added later.

35 [26], §5.2.

36 [26], §4.6.4.2.2.

37 Other data exchange key encryption algorithms may be added later; see [27], §1 and 3.1.

38 [26], §4.6.3.1.

39 Other message integrity check algorithms may be added later; see [27], §1, §4.1, and 4.2.

40 This changes [23], §4.6.2.3, in which the signed message integrity check is never encrypted.

41 [23], §3.1.

42 [26], §4.6.2.1.

43 [26], §4.6.3.1.

44 [23], §3.4.1.1.

45 [23], §3.4.1.2.

46 [23], §3.4.1.5.

47 [23], §3.4.1

48 Revisions to separate the signature algorithms from the hash algorithms; this particular signa-
ture algorithm is the same as given in [26], §3.1.4.7; [27], §4.2, but in those RFCs it is defined

as a hash algorithm followed by an encryption algorithm. As other equally strong (or stronger) signature functions are found that are computationally as efficient as this one, they may be added to the list of acceptable hash algorithms. A signature algorithm using RSA-MD4 was not added as it is quite new, and so the protocol authors were less comfortable using it to protect the interchange keys embodied in the certificates.

- 49 The Distinguished Name may be written using the full T.61 character set, as described in X.500 [10]; however, some alternate representations of characters which software cannot print would be represented as a backslash followed by the character's octal representation.
- 50 Originally this was to be a structured Personal Name component (as defined in X.400 [8]) but was changed to conform to X.500.
- 51 For compatibility with the U. S. Government Open Systems Interconnection Profile [19], [23], §3.4.1.3 limited the number and length of each field in the Distinguished Name. This has since been changed to conform to the limits in X.520.
- 52 Note that [23] allows any Distinguished Name attributes to be used, unless expressly prohibited (and this is done only in the case of certifying authorities, which cannot have the Common Name attribute). However, allowing a certificate with an organization in the subject's Distinguished Name to be issued to an affiliated user would be very misleading unless the certificate were examined carefully, so the proposal was revised to eliminate possible confusion.
- 53 [23], §3.4.1.4.
- 54 [23] does not specify the order of organizational units, but they are to be written most significant member first.
- 55 [23], §3.3.3.1.
- 56 [23], §3.3.1.
- 57 [23], §3.3.2.
- 58 [23], §3.3.3.
- 59 [23], §3.1.
- 60 [23] does not require top-level certifying authorities to sign certificates of all certifying authorities in their domain.
- 61 [23] combines the guest and notary certificates, calling them notary certificates. The distinction was made to clarify whether or not organizational affiliation as well as identity was being

vouched for.

62 [23], §3.4.2.

63 [23], §3.3.3.2.

64 The international standard [11] does not recommend a key length but suggests that “a value... of 512 bits be adopted initially, but subject to *further study*” (emphasis in original). Originally, these limits were 320 and 632 bits (about 2×10^{97} and 2×10^{191} , respectively); these limits on the modulus size were chosen so that the software implementing the RSA cryptographic algorithms would be potentially exportable from the United States. However, other characteristics of software implementing these protocols also affects exportability, and given the success that Lenstra and Manasse have had in factoring numbers of around 100 digits [25], it was deemed prudent to increase the modulus size to a minimum of 512 bits.

65 [27], §3.1. The proposal also states that the public key is to consist of more than 100 digits ([23], §3.3.1); this means that if the exponent is 3, the modulus must be at least 100 digits. As noted in the previous footnote, the modulus should actually be much larger.

66 At most 17 multiplications are required to exponentiate to that power.

67 [31] discusses the theory behind this. The pseudorandom quantity concatenated with the exponent should contain at least half as many bits as the modulus, hence the change to [27], §3.1, which only required one 64-bit pseudorandom quantity for padding..

68 [23], §1.

69 [23], §3.3.3.

70 [22]; this mechanism for issuing, and paying for, certificates is not in [23], nor is it yet in place; BBN Communications is designing and building a prototype certificate generation unit.

71 The exact number has not yet been decided.

72 These details are not prescribed by either [23] or [26], but have since been adopted to ensure the user receives at least minimal information when something is amiss.

73 3[26], §4.5.

74 [26], p. 10.

75 [26], §4.2.