# Deploying a New Hash Algorithm

Steven M. Bellovin
smb@cs.columbia.edu
Columbia University

Eric K. Rescorla
ekr@networkresonance.com
Network Resonance

## Abstract

The strength of hash functions such as MD5 and SHA-1 has been called into question as a result of recent discoveries. Regardless of whether or not it is necessary to move away from those now, it is clear that it will be necessary to do so in the not-too-distant future. This poses a number of challenges, especially for certificate-based protocols. We analyze a number of protocols, including S/MIME and TLS. All require protocol or implementation changes. We explain the necessary changes, show how the conversion can be done, and list what measures should be taken immediately.

## 1 Introduction

Nearly all major cryptographic protocols depend on the security of hash functions. However, this is increasingly looking like a brittle foundation: although a variety of hash functions are available, only MD5 [Riv92] and SHA-1 [Nat02] are in wide use. Both hash functions derive from MD4 [Riv90], which has long been known to be weak [Dob96, Dob98], thus leading to concerns that they might have common weaknesses.

These concerns were borne out in late 2004, when techniques for efficiently finding collisions in MD5 [WY05] and SHA-0 [BCJ$^+$05] were announced. Subsequently, Wang [WYY05] announced a technique for finding collisions in SHA-1 in $2^{69}$ operations,[1] rather than the $2^{80}$ for which it was designed, and Lenstra et al. [LWdW05] demonstrated a pair of X.509 certificates with the same distinguished name, different public keys, and identical signatures, though no extension is known which can generate such a pair with different distinguished names.

It should be emphasized at this point that none of these results have translated into demonstrable attacks on real-world protocols, though [LWdW05] comes uncomfortably close. However, it is clear that neither MD5 nor SHA-1 is as strong as its target security level and so need to be replaced. The possibility of new attacks lends some urgency to this transition.

[1] In a presentation delivered at the Rump Session of CRYPTO 2005, Shamir stated that Wang had improved the attack to $2^{63}$ operations.

It is clear that a transition to newer hash functions is necessary. The need is not immediate; however, it cannot be postponed indefinitely. Our analysis indicates that several major Internet protocols were not designed properly for such a transition. This paper presents our results.

Although we don't discuss the issue in detail, most of our work applies to deploying new signature algorithms as well. If the signature algorithm is linked to a particular hash function, as DSA is tied to SHA-1, the two would change together; beyond that, since signature algorithms are almost always applied to the output of hash functions, if there is no easy way to substitute a new hash algorithm there is almost certainly no way to substitute a new signature algorithm, either.

## 2 Background

### 2.1 Uses of Hash Functions

Hash functions are used for many different purposes. In this section, we outline their major uses.

#### 2.1.1 Digital Signature

The purpose for which cryptographic hash functions were originally designed is input preparation for digital signatures. Algorithms such as RSA are far too expensive to apply directly to each input block of almost any real message. Instead, the message is securely compressed using a hash function; the resulting "fingerprint" is the input to the digital signature algorithm.

An enemy thus has two ways to attack a digital signature algorithm: either the algorithm itself can be cryptanalyzed (i.e., by factoring the RSA modulus), or the hash function can be abused.

#### 2.1.2 Message Authentication Codes

Hash functions are also used for high-speed message authentication between parties who share a common secret. There are a number of ways in which this can be done; the most common is to use the HMAC [KBC97] framework:

$$H(K \oplus \text{opad}, H(K \oplus \text{ipad}, M))$$

where $H$ is the hash function, $K$ is the shared secret and $M$ is the message to be authenticated. Note that the double application of $H$ and the double occurrence of $K$ helps prevent many attacks; the enemy does not know and cannot control what the inputs are to the outer hash.

### 2.1.3 Pseudo-Random Functions

Hash functions are often used as pseudo-random functions. That is, they provide a deterministic mechanism for generating random-seeming bit streams from some input source without disclosing any information about the input. A typical use is generating cipher keying material after a Diffie-Hellman exchange. IKE [KM97] uses HMAC for this purpose, as does TLS [DA99].

### 2.1.4 Data Fingerprinting

As noted above, hash functions can be used to produce fingerprints of files or messages. Sometimes, instead of digitally signing these fingerprints, the values are stored separately from the data. This permits later detection of changes to the original data.

One system in which this is used is Tripwire [KS94a, KS94c, KS94b]. Tripwire is used as a host intrusion detection system. Critical system files are fingerprinted; at intervals thereafter, the stored fingerprints are compared to values newly-calculated on the running system. Any change to a file will cause its fingerprint to change.

## 2.2 Overview of Recent Hash Function Attacks

Conventionally, hash functions are designed to have three properties:

**Collision resistance** It is computationally infeasible to find $x, y, x \neq y$ such that $H(x) = H(y)$.

**Preimage resistance** Given an output value $y$, it is computationally infeasible to find $x$ such that $H(x) = y$.

**Second preimage resistance** Given an input $x'$, it is computationally infeasible to find $x$ such that $H(x) = H(x')$.

The current generation of attacks address collision resistance. MD5 is effectively dead from that perspective; SHA-1 is much weaker than it should be, though finding collisions is still impractical.

While not as devastating as failures of the other two properties, collision resistance is indeed a serious issue. Lucks and Daum have generated Postscript files that exploit the attack (see `http://th.informatik.uni-mannheim.de/people/lucks/HashCollisions`). They took advantage of a well-known property of hash functions:

$$H(x) = H(y) \Rightarrow H(x||\Sigma) = H(y||\Sigma)$$

where $\Sigma$ is an arbitrary string, provided that $x$ and $y$ are the same length.

First, they generated two Postscript prologues that contained a collision in what was, syntactically, a constant. This constant was assigned to a variable. To each of these files, they then appended a Postscript program that checked the value of this variable and displayed one of two letters. An attacker could persuade someone to digitally sign the first, harmless letter; this same signature would match the second, harmful letter. Note, however, that to a great degree this attack is enabled by the fact that users do not directly view the Postscript code and rather use an interpreter. Similar attacks can be demonstrated against such systems (e.g., HTML with JavaScript) even without the ability to find hash collisions [Res05] by exploiting conditional elements in the display system.

Collision-finding attacks do not rule out all uses of a hash function. In particular, the pseudo-random function properties are not affected at all. Furthermore, HMAC is probably safe, since the unknown component—the key—of the inner hash function makes it impossible to generate a collision at that stage; this in turn helps protect the outer hash.

On the other hand, there is grave danger for many situations involving digital signatures or fingerprinting. If a would-be attacker can supply the message to be signed, that same attacker could have prepared two versions of the message, one innocuous and one harmful, while presenting only the former. The attacks work because the victim inspects the innocuous version and verifies that it is acceptable. In environments where victims do not carefully inspect data before it is hashed, collision attacks only modestly increase the threat level.

# 3 Overview of the Hash Transition Problem

Although the details of transition strategies for any given protocol may vary, there are many common elements. In this section, we provide an overview of the hash transition problem and the design goals that transition strategies should attempt to fulfill.

The hash transition problem is a special case of the general protocol transition problem. Whenever a new version of a protocol is rolled out, designers and implementors must figure out how to accomplish a smooth transition from old to new versions with a minimum level of disruption. In a typical protocol transitional environment, there are three types of agent:

**Old** Agents which only speak the older version.

**Switch-hitting** Agents which can speak both versions.

**New** Agents which can only speak the new version.

At the beginning of the transition, all agents are Old. At the end of the transition (at least theoretically), all agents are New. The purpose of a transition strategy is to accomplish the transition between these states with a minimum of disruption. This immediately implies some requirements for such a strategy.

Note carefully that the notion of an end to the transition is a theoretical one. In practice, transitions of this nature tend to persist for an arbitrarily long time, since old systems never quite die off. In the security field, this is especially bad, because it leaves open the door to downgrade attacks.

## 3.1 Backward Compatibility

The most basic requirement for a seamless transition is backward compatibility. Old agents and Switch-hitting units should be able to communicate, using the older version. Without backward-compatibility, users have an enormous disincentive to upgrade their software because it immediately cuts them off from most of the people they would communicate with. The general approach is to deploy Switch-hitting implementations until most implementations are Switch-hitting. Once this is accomplished, New implementations can be safely deployed.

In interactive protocols such as SSL/TLS or IKE, backward compatibility is generally accomplished by having the Switch-hitting peer recognize that it is speaking to an Old peer and fall back to the older version. In non-interactive protocols such as S/MIME, Switch-hitting implementations must transmit messages that can be read by Old implementations unless they know that the peer is Switch-hitting or New.

## 3.2 Newest Common Version

When two Switch-hitting clients communicate, they can either use the new or old versions of the protocol. Because the purpose of the transition is to deploy the newer version, it is desirable that they use that version where possible. With interactive protocols, this is straightforward; the peers detect that they both speak the new version and simply use it. With non-interactive protocols it is more difficult. The standard approach is for Switch-hitting implementations to start out speaking the older version but advertise that you support the newer version. This allows the receiver to cache that the peer is Switch-hitting and use the newer version from then on. Absent corner cases, the recipient of a new-version message can assume that the sender can read the new version.

## 3.3 Downgrade Protection

An additional requirement for security protocols is to defend against version/algorithm downgrade. Consider the situation where two peers each support two cryptographic algorithms, one of which is strong and one of which is weak. If an attacker can force the peers to use the weaker algorithm, he may be able to attack the communication. The classic example of this attack is "export" algorithms: all versions of SSL/TLS up to and including TLS 1.0 [DA99] included support for weak "exportable" algorithms. In SSLv2 [Hic95] it was possible to force two implementations to use the weak algorithms even though strong algorithms were available.

The attack is shown in Figure 1. The attacker intercepts the client's CLIENT-HELLO message and removes the offer of the strong algorithm (in this case, RC4-128). The server thinks that the client is only offering weak algorithms and so negotiates RC4-40, leaving the connection open to attack.

Both SSLv3/TLS and IKE include partial defenses against downgrade; the basic strategy is to exchange MACs over the entire handshake. If the MACs verify, the peers can have some assurance that the handshake has not been tampered with. These defenses work well against compromise of the symmetric ciphers and partial compromise of the digest/MAC algorithms, but are not complete defenses. For instance, an attacker who had broken a peer's public key would be able to impersonate that peer even if the peer subsequently generated a new, stronger key (assuming, of course, that the public key has not been revoked). As will become apparent in subsequent sections, we cannot always depend on these mechanisms.

## 3.4 Credentials versus Implementations

Another issue specific to security protocols is the separation of credentials and implementations. In typical public key-based systems, a peer's public keys is authenticated using certificates (in the case of the protocols being discussed here, PKIX [HPFS02] certificates). Certificates are a general credential and are not tied to any specific revision of a given security protocol. Moreover, the upgrade cycle for protocol implementations is uncoupled from the certificate issuance cycle. There is thus the potential for a situation in which the user has a security protocol implementation which understands SHA-256 but a certificate which was digested with SHA-1—or indeed, one certificate digested with SHA-1 and one with SHA-256. As we shall see, this can lead to situations in which certificate and protocol version capabilities need to be dealt with separately.

| Client | | Attacker | | Server |
|---|---|---|---|---|
| CLIENT-HELLO[RC4-128,RC4-40] | $\longrightarrow$ | CLIENT-HELLO[RC4-40] | $\longrightarrow$ | |
| | $\longleftarrow$ | SERVER-HELLO[RC4-40] | $\longleftarrow$ | SERVER-HELLO[RC4-40] |

Figure 1: SSLv2 downgrade attack

### 3.5 Outline

In the remainder of this paper, we consider the application of these principles to three major Internet security protocols: S/MIME (a store-and-forward protocol), TLS (a session-oriented protocol), and IPsec (which separates key exchange from transport). A longer version of this paper, with analysis of additional protocols and protocol design principles, appears in [BR05].

## 4 S/MIME

The first protocol we will consider is S/MIME [Ram04a, Ram04b, Hou04]. S/MIME is a standard message encryption and authentication protocol. In the most common modes, it uses public key cryptography (RSA [JK03] and DH [Res99]) for key establishment, symmetric cryptography for bulk encryption, and digital signatures (RSA and DSA [Nat00]) for message authentication/nonrepudiation. User public keys are transported/authenticated using PKIX [HPFS02] certificates.

There are five major types of S/MIME client:

1. Old clients.
2. Switch-hitting clients with only old certificates.
3. Switch-hitting clients with both types of certificate.
4. Switch-hitting clients with new certificates.
5. New clients with only new certificates.

The types of messages that each type of implementation should send are shown in Figure 2.

Note that this table assumes perfect information about the recipient's capabilities, which is not always the case. We now consider how to achieve interoperability in practice, which is a matter of trying to estimate the recipient's capabilities and create a message which they are most likely to be able to decode. For the remainder of this section, we focus on the behavior of Switch-hitting clients, since Old and New clients only have one possible behavior.

| | Receiver | | | | |
|---|---|---|---|---|---|
| **Sender** | Old | Switch/Old | Switch/Both | Switch/New | New |
| Old | Old | Old | Old | - | |
| S/O | Old | Old | Old | Old | - |
| S/B | Old | Either | Either | Either | New |
| S/N | - | New | New | New | New |
| New | - | New | New | New | New |

Figure 2: Interoperability table for S/MIME implementations

### 4.1 The Initial Message

The first case we consider is the case where a user is sending a message to someone with whom he has never communicated before. There are two possible sub-cases:

1. The sender does not have the recipient's certificate.
2. The sender has the recipient's certificate.

We consider each sub-case in turn.

#### 4.1.1 Sending Without a Recipient Certificate

If the sender does not have access to the recipient's certificate, then he is subject to two limitations. First, he cannot encrypt because he does not have the public key to encrypt under. Second, he has no information about the recipient's capabilities, In particular, he cannot safely assume that the recipient's software will be able to process new hash functions.

**Choice of certificate** A sender with only one certificate must use that certificate. The difficulty comes when a sender has two certificates, one generated with an old hash function, and one with a new hash function. The possibilities, of course, are to use only one certificate or—because S/MIME allows multiple signatures—to use both. Any one-certificate strategy guarantees that some class of recipients will not be able to verify the message. Using both certificates preserves the possibility that the recipient can verify the message.

In order for this to work, however, recipients must be able to correctly verify messages with multiple signatures when one of them is unverifiable. Unfortunately, the S/MIME specification is fairly vague on this point. An unscientific poll of S/MIME implementors indicates that support for this option is spotty at best. Reported behaviors when one signature is good and the other unverifiable include:

- Verify only the first signature [Ram05].

- Generate an error [Hen05].
- Report success with warnings about unverifiable signatures [Gut05]

Only the final choice allows Switch-hitting implementations to guarantee interoperability for the messages they send. Another option would be to treat a message with any valid signature as valid, but we have not heard this behavior reported.

Because receiver behavior is unpredictable, senders must attempt to estimate what sorts of implementation receivers are likely to have. This probably means choosing interoperability with the most popular strategies as a default (which are currently the older, weak, algorithms) and allowing users the option to configure a new behavior. This is irritating in that it involves a manual step if the sender guesses wrong. However there are already a number of non-security scenarios in which users must retransmit unreadable messages (bad attachment formats, HTML-vs-ASCII text, etc.) so it's not totally foreign to users.

If only collision attacks are available, there is little security advantage to the sender in using stronger hash functions, as long as he controls the content being signed (see Section 4.4.3 for what can happen if he does not). A third party attacker cannot use collisions to attack a signed message. However, if preimage attacks are possible, then such an attacker can potentially use the signature on a signed message to create a new signed message with different content. See 4.4.1 for more on this topic.

**Choice of digest algorithm** Once the certificate has been chosen, the sender must choose a digest algorithm to digest the message before signing. This choice is made independently for each signature, so it is possible the message will be digested twice. In general, if the certificate being used was generated with one of the old algorithms (MD5, SHA-1), the message should be digested using SHA-1, which receivers are required to accept by section 2.1 of RFC 3851 [Ram04b]. This minimizes the chance that the recipient will not be able to verify the message signature. (MD5 should not be used at all for message digests, even if the certificate uses it.)

If the certificate being used was digested with a new hash algorithm, we recommend that the sender use the same algorithm to digest the message, on the grounds that if the recipient can use the digest algorithm to verify the certificate they can use it to verify the message. This runs the risk that the recipient will be using a separate toolkit to verify the certificate signature than they used to verify the message signature; however we are not aware of any S/MIME client that behaves in this way. This algorithm has the attractive property that it automatically works correctly with DSA, which can only sign SHA-1 digests. The only requirement is that the certificate is itself signed with DSA, as is standard practice.

#### 4.1.2 Sending With the Recipient's Certificate

The case where the sender has the recipient's certificate(s) is somewhat simpler. Although there is no guarantee, we believe that it is a reasonable assumption that implementations can verify their own certificates and therefore must implement whatever digest algorithm was used to create them. If the recipient has only one certificate, the sender should therefore use their certificate with the corresponding algorithm. If the recipient has multiple certificates, the sender should use the one created using the strongest algorithm. For the reasons indicated above, we do not recommend sending multiple certificates in this case.

The choice of which certificate to send would be simpler yet if the recipient's certificate indicated which algorithms it was capable of using. Although this not currently possible, the S/MIME working group is currently considering a considering a draft [San05] that would allow certificates to contain an SMIMECapabilities [Ram04b] extension for the owner of the certificate. This information could include information about allowed digest algorithms. However, because this extension is not included in current certificates or processed by current implementations, it is of limited value in promoting interoperability. Note that if the recipient's certificate was received via a signed message from the recipient, then it should contain SMIMECapabilities, which makes the problem easier, as discussed below.

### 4.2 Subsequent Messages

Once an S/MIME implementation has received a signed message from it is in a much better position to estimate the sender's capabilities. For clarity, say that Alice has received a signed message from Bob. With high probability Bob can verify signatures produced with whatever algorithm(s) it used to digest its own message. If this is a new (strong) algorithm then all is good and Alice should herself use that algorithm.

If Bob used an old (weak) algorithm, then Alice at least knows that she can communicate with Bob using that algorithm. However, it is still possible that Bob has a Switch-hitting implementation. Optimally, Alice would be able to detect this case and use a newer algorithm for her response. S/MIME has a standard way for Bob to signal this fact using the SMIMECapabilities signature attribute, which includes a (potentially partial) list of the algorithms that Bob supports. Bob can send a message using SHA-1 but include an SMIMECapabilities attribute indicating that he also supports SHA-512. If this attribute is included, it is always signed, thus preventing the introduction of a false attribute.

Thus, Bob can send a message using SHA-1 but include an SMIMECapabilities attribute indicating that he also supports SHA-512. Upon processing the signature, Alice can detect that Bob's implementation is Switch-hitting and respond with the stronger algorithm. Thus, we recommend that when Switch-hitting implementations send messages using weak algorithms they include an indication that they also support a stronger algorithm. There is no point in including such an indication if you are sending with the stronger algorithm, since that algorithm is preferred and a recipient which cannot process the stronger algorithm cannot verify that you also support the weak one.

If, on the other hand, Bob's message includes an SMIMECapabilities attribute saying that he does not support strong hash functions, Alice's system will be forced to use old ones. Presumably, her implementation will cache that information. It is important that such cache entries expire after some period of time, since Bob may upgrade his client and certificate.

Because the SMIMECapabilities attribute is part of the signerInfo element, it is not included in messages which are unsigned. However, if Alice receives an encrypted message from Bob, she knows that he was able to verify the certificate that he used to encrypt to her. Therefore, if she wishes to sign future messages she should digest using whatever algorithm was used to produce that certificate.

## 4.3 Diffie-Hellman Key Agreement

RFC 2631 [Res99] specified a method for Diffie-Hellman (DH) key agreement in which SHA-1 is used as a PRF (pseudo-random function) to compute key encryption keys from the DH shared secret. There is no room for negotiation here: the standard specifies SHA-1, and a new algorithm identifier would need to be defined for DH with the newer hash function. It is not currently known how to attack SHA-1 when used in this way. If such an attack were to be found, implementations would need to convert to a new digest algorithm and use it every time they used DH key agreement. Although this does not guarantee interoperability, the alternative is worse: encrypting data with an algorithm known to be insecure.

## 4.4 Attacks

In this section, we consider the problem of protecting Switch-hitting implementations during the transition period when it is impractical to turn off support for the old algorithms. There are three basic scenarios:

- The attacker does not have a valid certificate and private key for either peer.
- The attacker has acquired a valid (but false) certificate and knows the private key.
- The attacker is one of the communicating parties.

### 4.4.1 Attacks Without a Valid Certificate

If the attacker does not have a valid certificate, then his ability to mount attacks, even on older digest algorithms, is fairly minimal unless he can compute preimages.[2] Clearly, an attacker who can compute preimages can undetectably modify messages in transit. In this case, the only defense is to stop using the affected algorithm. Note that senders cannot prevent this attack by multiply signing their messages; S/MIME multiple signatures are parallel and independent, so the attacker can simply strip the strong signature. Indeed, as a general matter, not sending messages signed with old algorithms is not a complete defense against preimage attacks. Because S/MIME messages are generally not securely timestamped, an attacker can potentially attack any signed message, even a historical one, so the increase in security exposure by continuing to send messages with old digests may not be that large. Rather, receivers must stop accepting an algorithm where computing preimages is possible.

### 4.4.2 Attacks Using a Valid Certificate

If the attacker has a certificate with a valid signature containing the identity of one of the peers—for instance obtained using an improved version of the Lenstra construction—he can impersonate that peer. This would allow him to forge messages that appear to be from that peer. It may also allow him to convince the other peer to encrypt messages using his fake certificates. The only certain countermeasure here is to stop accepting the compromised algorithm. One partial workaround would be for the victim to refuse to accept certificates dated after the time when the algorithm was compromised. This is a defense against collision attacks, but if the attacker can generate 2nd preimages, then he can forge a certificate with an arbitrary date and bypass this countermeasure. Another partial workaround is to store copies of previously used peer certificates (as with SSH [Ylo96, Ylo05]), thus reducing the window of exposure to the first exchange of messages.[3]

[2] An attacker who can compute preimages is likely to be able to forge certificates. However, it is possible that an attacker could compute preimages but without fine enough control to forge a specific certificate.

[3] Note that it's common to store a digest of the certificate rather than the certificate itself. This obviously leaves one open to preimage attacks if the attacker can manage to get a certificate with the same digest (not easy, because he must also simultaneously attack the CA's digesting process which covers different data). If a digest is being stored, it might be wise to store a keyed hash using some locally known key instead.)

#### 4.4.3 The Attacker is One of the Communicating Parties

If it is easy to find collisions in a hash, then being one of the communicating parties—or at least in a position to substantially control the message contents—confers substantial advantage to the attacker. In particular, it allows him to cheat in contexts where an S/MIME signature is to be verified by a third party. The basic scenario is described in Section 2.2: two versions of a document are prepared, one innocuous and one malicious. One or both of the parties signs the innocuous version and then the attacker convinces the third party that the victim signed the malicious version. This attack can be mounted regardless of which party does the actual signing. The key is for the attacker to be allowed to prepare the document to be signed, since the colliding pair must be generated together.

In order to mount this attack on a Switch-hitting peer, the attacker must represent that he only supports the broken algorithm, thus forcing the signature to be performed using that algorithm. However, since supporting only old algorithms is a legitimate configuration, this is extremely easy to achieve. The victim has the choice of using that algorithm or not communicating at all.

This attack is extremely difficult to defend against in standard systems. Bob can defened against being conned by preparing the final document version and inserting enough randomness near the beginning (e.g., in a dummy field) to make it infeasible for Alice to have generated a collision.[4] However, this is complex and not supported by typical application software. Moreover, Alice should be suspicious of this request, since it allows Bob to mount a collision attack himself. A more general defense is for the parties to jointly agree on random values once the document content is fixed, but this is even more complex for ordinary users.[5] S/MIME implementations could of course do this automatically, but if one is willing to modify implementations it is easier to simply add strong algorithms.

We stress that this attack is very real and very practical if MD5 is used.

Because defense against this attack is difficult, in contexts when users are signing messages that might be verified by a third party, it is better to simply insist on using a strong algorithm. Similarly, third parties should be extremely suspicious when they are asked to rely on signatures that use weak algorithms, especially MD5. Note that as with the Lucks/Daum attack, close inspection of such messages generally will reveal their unusual structure and so this attack can only be mounted when the documents in question will be subject to only casual (or automatic) scrutiny.

[4] From a security perspective this is inferior to randomized hashing [HK05] but doesn't require changing the S/MIME implementation on either side.

[5] Kelsey and Kohno presented a "Herding" attack at the CRYPTO '05 rump session that allows cheating in this scenario, but the effort level ($2^{87}$ for MD5, $2^{108}$ for SHA-1) far exceeds that of ordinary collision finding.

# 5 TLS

TLS [DA99] is a standard channel security protocol which lives above the transport layer (where the OSI session layer sits). Originally designed for Web security [Res00], it is now widely used for other application protocols including SIP [RSC$^+$02] and SMTP [Kle01]. The most common TLS deployment involves an anonymous client connecting to a server an using the server's certificate and public RSA key for key exchange.

There are five major places digest algorithms are used in TLS:

- In the per-record MAC.
- In the certificates used by client and server.
- In the digitally-signed element.
- In the PRF used to make keying material.
- In the Finished message

TLS contains an extensive framework for algorithm negotiation, using the concept of "cipher suites". A cipher suite consists of a triple specifying the key establishment mechanism, the symmetric encryption algorithm used to encrypt traffic, and the message digest used to provide traffic message integrity. For instance, the cipher suite TLS_RSA_WITH_RC4_128_MD5 indicates RSA key exchange, encryption with RC4-128, and message integrity with a MAC based on MD5 (in TLS this is HMAC-MD5 [KBC97].)

Unfortunately, this mechanism is only useful for negotiating the record MAC. Although there is a mechanism for negotiating client certificate type, it does not include digest algorithm and the other algorithms cannot be negotiated. Indeed, the PRF, ServerKeyExchange, and ClientVerify messages are not parametrized, but rather are specified directly in the standard. In order to accomodate newer digest algorithms in these cases we must extend TLS.

## 5.1 MAC Functions

Negotiating the MAC in TLS is straightforward. Each cipher suite specifies the digest function function to be used as the basis for the MAC. So, in principle all that needs to be done is to define a new set of cipher suites with stronger hash algorithms. Note that because TLS uses HMAC, the current collision-only attacks most likely do not represent a threat, thus making this a low priority upgrade.

## 5.2 Server Certificates

The most important element of TLS to upgrade is the server certificate. Because certificates are automatically verified, they are the cryptographic technique most threatened by current digest attacks. TLS client certificates are rare; by contrast, virtually every TLS server has a certificate.

We assume that during the transition period, each server will have two certificates, one created with an old hash (typically SHA-1 or MD5) and one created with a new hash. The client can then indicate to the server that it can process the new certificate. There are two potential techniques for doing this: an overloaded cipher suite and a TLS extension [BWNH+03]. The TLS extension approach is probably superior in that it preserves protocol cleanliness—the hash functions in the TLS cipher suite offers do not refer to the certificate. Moreover, there are performance reasons for the client to prefer to use the older hash algorithms for MAC functions: SHA-1 is much faster than SHA-256, and the MAC functions do not need to be upgraded immediately.

Note that this does not address the problem of DSA, which, as noted previously, cannot be used with any algorithm other than SHA-1. The cleanest solution for DSA is simply to define a new set of cipher suites that specify a newer version of DSA (e.g., DSA2). This allows a client to simultaneously offer RSA with multiple algorithms but DSA with only SHA-1. If the newer version of DSA allows algorithm flexibility then the extension could extend to negotiating that algorithm as well.

## 5.3 Client Certificates

TLS client certificates are much less commonly used; where they are used they are often self-signed, although the US government is now issuing client certificates for establishing user identities. However, in the case where client authentication is used, it is desirable to have a way for the server to indicate which hashes it would like the client to use. This is a fairly simply protocol engineering matter with two obvious alternatives:

- Add new values to the certificate_types field of the CertificateRequest message. For instance, an rsa_sign_sha256 type could be created.

- Use extension values.

Each of these approaches has advantages. The CertificateRequest approach keeps all the information about the certificates that the client should produce together. Unfortunately, it creates the risk of combinatoric explosion of certificate_types values. Currently, four code points (rsa_sign, dss_sign, rsa_fixed_dh, and dss_fixed_dh) out of a possible 256 are defined. Every new hash function added thus potentially creates four new code points, and more if additional signature algorithms are defined.

The alternative approach is for the server to use an extension (most likely in response to the client's extension) indicating which hash algorithms it accepts. This is less elegant, but removes the combinatoric explosion problem. Neither approach is superior from a security perspective.

## 5.4 The Digitally-Signed Element

There are two places in TLS where data is explicitly digitally signed: the CertificateVerify and the ServerKeyExchange. In both places, the signature is accomplished using the digitally-signed element. When the signature algorithm is DSA, the input is as expected—a SHA-1 digest of the data to be signed. However, when the signature algorithm is RSA, the input is something unusual: the MD5 and SHA-1 digests of the input are concatenated and fed directly into the RSA signature algorithm with PKCS#1 padding, but without DigestInfo wrapping. This is not a negotiatiable algorithm but rather is wired into the specification.

This unusual construction raises the question of what the target construction should be. The original rationale for the dual hash construction was to provide security in the face of compromise of either hash. However, in practice this has been partially undercut by the common heritage of SHA-1 and MD5. A practical attack on SHA-1 could potentially extend to compromising the MD5/SHA-1 pair. The general feeling in the TLS community is that a single negotiated digest would be a better choice.

The best choice here is probably to have the digitally-signed element use the same algorithm as was used to sign the certificate of the party doing the signing (the client for the CertificateVerify and the server for the ServerKeyExchange). This avoids the creation of a new negotiable option, thus reducing protocol complexity. In principle this could lead to interoperability problems if the certificate system has different capabilities than the TLS implementation. However, we're skeptical that the number of real implementations with this problem would be large enough to justify the additional complexity.

There are two different ways to roll out this change. The first is to simply decree that new cipher suites (e.g., one that used SHA-256) use their hash to produce the digitally-signed element. This produces an inconsistency in that the older cipher suites would still be using the combined hash construct. However, there is not a security problem with this strategy since those cipher suites all use SHA-1 or MD5 for their MAC in any case, and the current construction is no weaker than SHA-1 or HMAC alone.

The second approach would be to simply change this rule in the next version of TLS.[6] This would be a more principled approach but has the drawback that TLS is otherwise extremely stable and that new versions have historically taken a very long time to produce even when the revisions were minor. Therefore, this would be the slower approach. However, given the relatively low security threat posed by the current attacks and the likely catastrophic nature of any hash compromise that would allow attacking the digitally-signed element, the level of urgency is relatively low. Thus, while either roll-out strategy is probably acceptable, we prefer the new version strategy as a matter of protocol cleanliness.

## 5.5 PRFs

TLS uses a hash function-based PRF to create the keying material from the PreMaster Secret and Master Secret. It is also used to compute the Finished messages which are used to secure the TLS negotiation against downgrade attack. Compromise of the PRF might potentially allow an attacker to determine the keying material or mount a downgrade attack.

The TLS PRF is actually two PRFs, both based on HMAC, with one using MD5 and the other using SHA-1. Like the digitally-signed element, the TLS PRF is explicitly specified in the standard and not negotiable.[7] This construction, while somewhat over-complex, is provably secure under the assumption that either HMAC-SHA1 or HMAC-MD5 are secure pseudorandom functions [Kra03]. Because the current attacks do not affect the security of HMAC, upgrading the PRF is a low-priority task. However, we briefly consider methods here.

The two basic methods for negotiating the PRF algorithm are to use the negotiated cipher suite or to create a new extension. In the first case, whatever digest algorithm was negotiated for the cipher suite would also be used as the basis for the PRF. This has the obvious drawback that it ties TLS to the basic HMAC-X structure of the PRF. If this construction were found to be insecure (despite the proofs of security), then it would not be possible to negotiate a new construction. By contrast, while using an extension adds complexity it would allow substitution of the construction without creating a new version of TLS.

We are skeptical that this increased flexibility justifies the added complexity of defining a new extension. In view of the security proofs for HMAC and its wide use in TLS, it seems likely that any attack on HMAC would imply compromise of the underlying digest function and result in the compromise of key elements of the system (message MACs, certificates, etc.), thus necessitating a new revision of TLS in any case. It would be straightforward to revise the PRF at that time.

PRFs have similar roll-out issues to those described in Section 5.4. As with the digitally-signed element, we recommend that the transition to a negotiated PRF occur in a future version of TLS.

## 5.6 The Finished Message

The TLS Finished message is computed by computing the TLS PRF over the master secret and the concatenation of two digests over the handshake messages, one using MD5 and one using SHA-1. The same considerations apply here as in the PRF. The hash itself is unkeyed although both sides contribute random nonces. This design modestly reduces memory requirements on the client and server. HMAC-based MACs digest the key before the data; however, the MAC key (the master secret) is not known until after the ClientKeyExchange message. Thus, the client and server cannot start computing an HMAC immediately and must instead store the pre-ClientKeyExchange messages (about of 2-5k of data). There is a potential risk in this design in that keyed hashes are harder to attack than simple hashes. However, because the attacker cannot control the client messages and can only slightly influence the server's messages (by modifying the client messages in flight to produce a different negotiation result) the ability to create collisions is insufficient to mount this attack.

The obvious approach to transition is to replace the pair of hashes with the negotiated hash function used for the message MAC. However, note that this requires both sides to store the handshake messages until the MAC algorithm is decided (in the ServerHello). This requires a modest change in TLS implementation behavior and a slight increase in storage requirements. An alternative design would be to replace the "digest then PRF" construction with a MAC directly over the handshake messages. This would have only slightly higher storage requirements and be modestly more secure in the event of preimage attacks on the underlying hash function. We consider either approach adequate, though we believe that the security considerations outweigh the memory issue and therefore recommend transitioning to a simple MAC over the messages.

## 5.7 Attacks

As with S/MIME, we consider the problem of protecting Switch-hitting implementations during the transition period. The general form of the attack is for the enemy to force one or both sides to believe that the other side is an old implementation and convince them to use weaker algorithms, thus rendering them susceptible to attack.

[6] This would be TLS 1.2 as TLS 1.1 [DR05] has just been approved.

[7] This has already been an issue with the proposed GOST cipher suite [CL04], which for regulatory reasons must use the GOST digest function in the PRF

We can divide these attacks broadly into two categories. In the first, the attacker has obtained a valid certificate for one side of the connection (most likely the server) and knows the corresponding private key. In the first case, where the attacker has a valid (but fake) certificate, no complete defense is possible other than turning off the old algorithm. The attacker can simply intercept the connection and use its certificate. As with S/MIME, partial defenses including rejecting newer certificates signed with weak algorithms and SSH-style fingerprint comparison.

If the attacker does not have a valid certificate, he must attack the negotiation more indirectly. However, because the negotiation is protected by a MAC computed using the PRF, the attacker must be able to predict PRF output in order to predict the key used for the PRF. As argued in Section 5.5, this would require a very serious break of HMAC and most likely that the attacker can compute preimages, making a direct attack on certificates possible.

# 6 IPsec

IPsec[KA98c] is composed of two major pieces: the per-packet protection mechanisms, ESP [MD98] and AH [KA98a], and the key exchange algorithm, IKE [HC98]. The two pieces have very different dependencies on hash algorithms.

A revised IPsec specification is currently being prepared. For our purposes, the most important change is the replacement of IKE by a substantially different version [Kau04]. However, because the IPsec WG opted to retain the basic elements of IKE (except for Aggressive Mode) in IKEv2, our analysis is largely the same. Differences are noted as necessary.

Although IPsec can provide general security, realistically it is generally restricted to VPNs. This implies that each VPN gateway knows its clients, and has perhaps issued their certificates. The gateway thus has the ability to refrain from using new algorithms until it has issued new certificates to its clients. More generally, with IPsec each party often knows who the other party is, and what credentials it will present; the presence of new hash functions in a peer's certificate are thus a signal for what certificate it should use. In some cases, a gateway may be configured to trust all clients presenting a certificate from a particular CA or group of CAs; in such cases, the signaling mechanisms described below can be used.

## 6.1 AH and ESP

AH [KA98a], the authentication header, provides authentication only. The usual algorithm is HMAC [KA98a, MG98a, MG98b, KA98b] with either MD5 or SHA-1. As noted, the use of today's hash functions within HMAC is not believed to be risky; as such, no changes are needed to AH. That said, HMAC-MD5 has been deprecated for use with IPsec.

ESP [KA98b], the Encapsulating Security Protocol, provides confidentiality and/or integrity protection. As with AH, the standard integrity algorithms are based on HMAC, and thus require no changes.

## 6.2 IKE

IKE [HC98, MSST98] is an extremely complex protocol, with many different variants. Authentication can be via public key technology, in which case certificates and hash functions are used, or shared secrets. In addition, hash functions are always used as PRFs and for integrity protection. We look at each of these issues separately.[8]

IKE has two phases. In the first phase, the two parties authenticate themselves to each other using potentially-expensive mechanisms. In addition, during Phase 1 they negotiate hash algorithms, authentication methods, and PRFs for use in Phase 2. Phase 2 is used to set up actual IPsec security associations (SAs); at this time, algorithms for such associations are negotiated. Because there is full negotiation via a protected channel, there are no compatibility issues with Phase 2; accordingly, we will not discuss it further.

For IKEv2 [Kau04] and Main Mode of IKEv1, the first set of messages in IKE contain security association proposals for use during Phase 2 and the remainder of Phase 1. This permits early negotiation of hash functions and PRFs. A party that has implemented new hash functions can, of course, specify them at this point. The IKE initiator transmits an ordered list of the algorithms it considers acceptable; the responder selects one from that list. There is no requirement that they be ordered by strength. Note, though, that there are no Transform types defined for hash or signature algorithms. Furthermore, it is possible to have valid SA messages that don't mention any hash functions at all; both the PRF and Integrity algorithms have AES-CBC variants. Accordingly, new Transform types are necessary.

The situation is much more complex if Aggressive Mode (IKEv1 only) is used, since the messages exchanged differ greatly. There are four different variants, depending on how the exchange is authenticated. In all of them, however, the initial SA message is combined with other parts of the key exchange. There is thus no opportunity for prenegotiation of hash function capabilities, and hence no graceful upgrade path.

If Phase 1 is authenticated with digital signatures, the responder sees the initiator's SA proposal before perform-

[8] In addition to MD5 and SHA-1, [HC98] says that the TIGER [AB96] hash function "should" be supported. To our knowledge, this is very rarely done in practice.

ing any public key operations. It thus knows the initiator's capabilities, and hence which of its certificates it may use. As with TLS, downgrade attacks are prevented by later hashing the SA proposals.

If public key encryption or revised public key encryption is used to authenticate the Phase 1 exchange, the situation is more complex. The initiator may use a hash function in its first message; it must also encrypt certain values with the responder's public key. Before doing this, however, the initiator must have the responder's certificate. It can thus use the same heuristic we have discussed earlier: if the certificate uses a new hash function, the client should do so as well.

Finally, shared secret mode can be used for initial authentication. In this case, there are no certificates; however, hash functions are used. The initiator does not need to employ any until it has seen the responder's SA; there is thus no problem negotiating newer hash functions if available.

## 6.3 Hash Functions

Section 4 of [HC98] requires that all exchanges start by negotiating certain SA parameters. One such parameter is which hash function should be used by IKE.

The negotiated hash function is used for several purposes. If no PRF is negotiated, the selected hash function is used via HMAC. In addition, a hash of a certificate is sometimes transmitted, to indicate which one is in use.

We suggest overloading this message for signaling what hash functions can be used in certificates. That is, if a new, strong hash function appears in the initiator's SA proposal, the responder can assume that the initiator will accept that hash function in certificates. Similarly, if the responder specifies such a function, the initiator will know the same about the responder. This is not the cleanest way of conveying this information; however, the code impact should be minimal.

The alternative would be to add another payload to the SA proposal message. We suspect that this would cause more interoperability problems; however, it would be cleaner, and would be the only way to signal support for new signature algorithms.

Given the limited direct use of hash functions in IKE, there is arguably no need to upgrade them. As far as is known, there is no need to use HMAC-512. However, the performance impact is minimal, and the ability to signal is quite important.

### 6.3.1 MAC Functions

IKE uses HMAC for authentication. As before, HMAC is believed to be resistant to collision attacks. There is thus no need for enhanced MAC functions.

### 6.3.2 PRFs

PRFs are use in IKE for key generation. Two types are in use, HMAC with MD5 or SHA-1 and AES-XCBC-PRF-128 [Hof04]. There is thus no need to change behavior here.

## 6.4 Attacks

Given that all uses of hash functions except for certificate exchanges are matters for negotiation, and given that IKE already uses protection against downgrade attacks, we restrict our attention to certificate exchanges. The question must be addressed for both initiator and responder certificates. We first consider Main Mode and IKEv2.

If the SA signaling described in Section 6.3 is used, the initiator will have a clear indication of whether or not the responder supports new hash functions. An attacker who has somehow created a fake responder certificate could tamper with the SA response; however, this will be detected as a downgrade attack. It is thus not possible to confuse the initiator. Similarly, if the attacker tampered with the SA proposal, the responder might believe that the initiator only supported old hash functions; again, this is easily detected.

Aggressive mode is much more complex, because of its many variants. As outlined earlier, though, the cases reduce to preprovisioning or downgrade protection.

There is one more interesting situation to consider: opportunistic encryption [RR05]. With opportunistic encryption, there is no prior knowledge of a peer's identity, let alone capabilities. Fortunately, [RR05] requires use of Main Mode, where a full SA negotiation is done beforehand.

# 7 Conclusions

It is clear that new hash functions or new methods of employing hash functions are necessary. However, as we have demonstrated, neither the specifications nor implementations are ready for the transition. We have presented an analysis of transition strategies for S/MIME, TLS, and IPsec; we have also analyzed DNSsec in [BR05]. We strongly urge the analysis of other protocols that use hash functions. Prominent candidates include OpenPGP [CDFT98], and Secure Shell [Ylo96, Ylo05].

For the protocols we analyzed, we present recommendations to implementors and the IETF. These changes are necessary *to prepare* for the transition. We suggest that they be made as quickly as possible, to provide maximum secure interoperability when new hash functions are ready.

In a number of protocols, users need to have a choice of which hash functions to offer or accept. We urge implementors to make this easily configurable, both by end users and system administrator.

When protocol upgrades are being designed, consideration should be given to signature algorithm agility as well. In most cases, the signaling will have to be done in the same place as for hash functions. However, some of the overloading we suggest is inappropriate for signature algorithms. For example, Section 6.3 suggests using the appearance of a new hash algorithm in the SA proposal as a signal that one party supports a new hash algorithm in one context, and hence presumably in another. There is no obvious way to extend this to, say, support of ECC signatures. (The growing popularity of ECC may require these changes sooner than would be required by the current attacks on hash functions.)

**S/MIME**

Implementors of S/MIME should ensure that their product handles multiple signatures properly. In particular, programs should report success with one signature while warning about unverifiable signatures.

MD5 should *never* be used for digests, since all conforming implementations already support SHA-1.

The IETF should develop a method for indicating digest function capabilities in certificates, CA vendors should implement it, and new certificates should contain explicit statements about hash functions supported.

**TLS**

The IETF should define a TLS extension by which clients can signal support for newer certificates.

The IETF should pick one of the two suggested alternatives for supporting client side certificates properly.

The IETF should consider making the PRF depend on the MAC algorithm in a future version of TLS.

The definition of the digitally-signed element should be amended to support new hash functions.

The definition of the Finished message should be amended to support new hash functions.

**IPsec**

The IETF should amend the IKE and IKEv2 specifications to describe signaling via the SA hash proposal, or via an explicit new field in the SA exchange.

**DSA**

DSA presents a special problem, since it may only be employed with SHA-1. NIST needs to clarify this situation, either by defining DSA-2 or by describing how DSA can be used with randomized hashes or truncated longer hashes.

The problems we have described here are symptomatic of a more general problem. Most security protocols allow for algorithm negotiation at some level. However, it is clear that this has never been thoroughly tested. Virtually all of the protocols we have examined have some wired-in assumptions about a common base of hash functions. It is a truism in programming that unexercised code paths are likely to be buggy. The same is true in cryptographic protocol design.

# 8 Acknowledgments

The authors would like to thank Russ Housley for his detailed review of this paper. We would also like to thank Alex Alten, Steve Burnett, Hal Finney, Eu-Jin Goh, Peter Gutmann, Steve Henson, Paul Hoffman, Michael Howard, Charlie Kaufman, John Kelsey, Ben Laurie, Arjen Lenstra, Blake Ramsdell, Hovav Shacham, and Benne de Weger for their advice and comments.

# References

[AB96] R. Anderson and E. Biham. Tiger: A fast new hash function. In *IWFSE: International Workshop on Fast Software Encryption, LNCS*, 1996.

[BCJ+05] E. Biham, R. Chen, A. Joux, P. Carribault, C. Lemuet, and W. Jalby. Collisions of SHA-0 and Reduced SHA-1. In *Proceedings of Eurocrypt '05*, 2005.

[BR05] Steven M. Bellovin and Eric K. Rescorla. Deploying a new hash algorithm. Technical Report CUCS-036-05, Dept. of Computer Science, Columbia University, October 2005.

[BWNH+03] S. Blake-Wilson, M. Nystrom, D. Hopwood, J. Mikkelsen, and T. Wright. Transport Layer Security (TLS) Extensions. RFC 3546, June 2003.

[CDFT98] J. Callas, L. Donnerhacke, H. Finney, and R. Thayer. OpenPGP Message Format. RFC 2440, November 1998.

[CL04] G. Chudov and S. Leontiev. Addition of GOST ciphersuites to Transport Layer Security (TLS), May 2004. draft-chudov-cryptopro-cptls-01.txt.

[DA99] T. Dierks and C. Allen. The TLS Protocol Version 1.0. RFC 2246, January 1999.

[Dob96] H. Dobbertin. Cryptanalysis of MD4 (Third Workshop on Cryptographic Algoruthms, Cambridge 1996). *Lecture Notes in Computer Science*, pages 55–72, 1996.

[Dob98] H. Dobbertin. The First Two Rounds of MD4 are Not One-Way. *Lecture Notes in Computer Science*, 1372, 1998.

[DR05] T. Dierks and E. Rescorla. The TLS Protocol: Version 1.1, Jun 2005. draft-ietf-tls-rfc2246-bis-13.txt.

[Gut05] P. Gutmann. Personal communication, 2005.

[HC98] D. Harkins and D. Carrel. The Internet Key Exchange (IKE). RFC 2409, November 1998.

[Hen05] S. Henson. Personal communication, 2005.

[Hic95] K. Hickman. The SSL Protocol, February 1995. `http://www.netscape.com/eng/security/SSL_2.html`.

[HK05] S. Halevi and H. Krawczyk. Strengthening Digital Signatures via Randomized Hashing, May 2005. draft-irtf-cfrg-rhash-00.txt.

[Hof04] P. Hoffman. The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE). RFC 3664, January 2004.

[Hou04] R. Housley. Cryptographic Message Syntax (CMS). RFC 3852, July 2004.

[HPFS02] R. Housley, W. Polk, W. Ford, and D. Solo. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 3280, April 2002.

[JK03] J. Jonsson and Burton S. Kaliski. Public-key cryptography standards (PKCS) #1: RSA cryptography specifications version 2.1. RFC 3447, Internet Engineering Task Force, February 2003.

[KA98a] S. Kent and R. Atkinson. IP Authentication Header. RFC 2402, November 1998.

[KA98b] S. Kent and R. Atkinson. IP Encapsulating Security Payload (ESP). RFC 2406, November 1998.

[KA98c] S. Kent and R. Atkinson. Security Architecture for the Internet Protocol. RFC 2401, November 1998.

[Kau04] C. Kaufman. Internet Key Exchange (IKEv2) Protocol, Sep 2004. draft-ietf-ipsec-ikev2-17.txt.

[KBC97] H. Krawczyk, M. Bellare, and R. Canetti. HMAC: Keyed-Hashing for Message Authentication. RFC 2104, February 1997.

[Kle01] J. Klensin. Simple Mail Transfer Protocol. RFC 2821, April 2001.

[KM97] D. Kristol and L. Montulli. HTTP State Management Mechanism. RFC 2109, February 1997.

[Kra03] H. Krawczyk. SIGMA: The 'SIGn-and-MAc' Approach to Authenticaticated Diffie-Hellman and its Use in the IKE Protocol, June 2003. `http://www.ee.technion.ac.il/~hugo/sigma.ps`.

[KS94a] Gene Kim and Eugene H. Spafford. The design and implementation of Tripwire: A file system integrity checker. In *Proceedings of the 2nd ACM Conference on Computer and Communications Security*, November 1994.

[KS94b] Gene Kim and Eugene H. Spafford. Experiences with Tripwire: Using integrity checkers for intrusion detection. In *Proceedings of Systems Administration, Networking, and Security III*, 1994.

[KS94c] Gene Kim and Eugene H. Spafford. Writing, supporting, and evalutaing tripwire: A publically available security tool. In *Proceedings of the Usenix* UNIX *Applications Development Symposium*, 1994.

[LWdW05] A. Lenstra, X. Wang, and B. de Weger. Colliding X.509 Certificates. In *Proceedings of ACISP*, 2005. To appear. Online: `http://eprint.iacr.org/2005/067`.

[MD98] C. Madson and N. Doraswamy. The ESP DES-CBC Cipher Algorithm With Explicit IV. RFC 2405, November 1998.

[MG98a] C. Madson and R. Glenn. The Use of HMAC-MD5-96 within ESP and AH. RFC 2403, November 1998.

[MG98b] C. Madson and R. Glenn. The Use of HMAC-SHA-1-96 within ESP and AH. RFC 2404, November 1998.

[MSST98] D. Maughan, M. Schertler, M. Schneider, and J. Turner. Internet Security Association and Key Management Protocol (ISAKMP). RFC 2408, November 1998.

[Nat00] National Institute of Standards and Technology, U.S. Department of Commerce. Digital Signature Standard, 2000. FIPS PUB 186-2.

[Nat02] National Institute of Standards and Technology, U.S. Department of Commerce. Secure Hash Standard, 2002. FIPS PUB 180-2.

[Ram04a] B. Ramsdell. Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Certificate Handling. RFC 3850, July 2004.

[Ram04b] B. Ramsdell. Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification. RFC 3851, July 2004.

[Ram05] B. Ramsdell. Personal communication, 2005.

[Res99] E. Rescorla. Diffie-Hellman Key Agreement Method. RFC 2631, June 1999.

[Res00] E. Rescorla. HTTP Over TLS. RFC 2818, May 2000.

[Res05] E. Rescorla. MD5 Collisions in PostScript Files", June 2005. `http://www.educatedguesswork.org/movabletype/archives/2005/06/md5_collisions.html`.

[Riv90] R.L. Rivest. MD4 Message Digest Algorithm. RFC 1186, October 1990.

[Riv92] R. Rivest. The MD5 Message-Digest Algorithm . RFC 1321, April 1992.

[RR05] M. Richardson and D. Redelmeier. Opportunistic encryption using the Internet Key Exchange (IKE), 2005. draft-richardson-ipsec-opportunistic-17.txt.

[RSC+02] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. SIP: Session Initiation Protocol. RFC 3261, June 2002.

[San05] S. Santesson. X.509 Certificate Extension for S/MIME Capabilities, May 2005. draft-ietf-smime-certcapa-05.txt.

[WY05] X. Wang and H. Yu. How to Break MD5 and Other Hash Functions. In *Proceedings of Eurocrypt '05*, 2005.

[WYY05] X. Wang, Y. Yin, and H. Yu. Collision Search Attacks on SHA1, 2005. `http://theory.csail.mit.edu/~yiqun/shanote.pdf`.

[Ylo96] Tatu Ylonen. SSH – secure login connections over the Internet. In *Proceedings of the Sixth Usenix Unix Security Symposium*, pages 37–42, July 1996.

[Ylo05] T. Ylonen. SSH protocol architecture, 2005. draft-ietf-secsh-architecture-22.txt.