pre-workgroup
Internet-Draft
Expires: June 22, 2006

J. Fenton
Cisco Systems, Inc.

December 19, 2005

Analysis of Threats Motivating DomainKeys Identified Mail (DKIM)
draft-fenton-dkim-threats-02

Status of this Memo

Copyright Notice

Abstract

This document provides an analysis of some threats against Internet mail that are intended to be addressed by signature-based mail authentication, in particular DomainKeys Identified Mail. It discusses the nature and location of the bad actors, what their capabilities are, and what they intend to accomplish via their attacks.

Table of Contents

1. Introduction

DomainKeys Identified Mail (DKIM) [I-D.allman-dkim-base] defines a simple, low cost, and effective mechanism by which email messages can be cryptographically signed, permitting a signing domain to claim responsibility for the use of a given email address. Message recipients can verify the signature by querying the signer's domain directly to retrieve the appropriate public key, and thereby confirm that the message was attested to by a party in possession of the private key for the signing domain.

Once the attesting party or parties have been established, the recipient may evaluate the message in the context of additional information such as locally-maintained whitelists, shared reputation services, and/or third-party accreditation. The description of these mechanisms is outside the scope of this effort. By applying a signature, a good player will be able to associate a positive reputation with the message, in hopes that it will receive preferential treatment by the recipient.

This effort is not intended to address threats associated with message confidentiality nor does it intend to provide a long-term archival signature.

1.1. Terminology and Model

The following diagram illustrates a typical usage flowchart for DKIM:

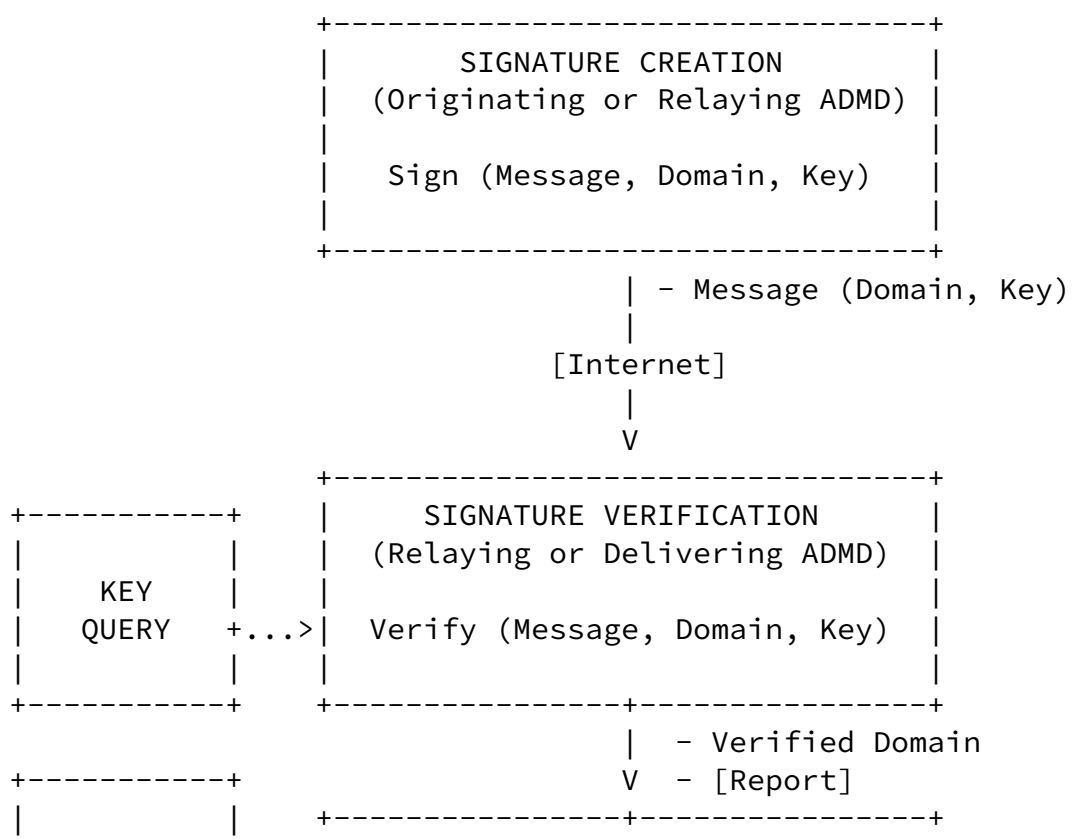
SIGNATURE CREATION
(Originating or Relaying ADMD)
Sign (Message, Domain, Key)
- Message (Domain, Key)
[Internet]
SIGNATURE VERIFICATION
(Relaying or Delivering ADMD)
KEY QUERY
Verify (Message, Domain, Key)
- Verified Domain
- [Report]

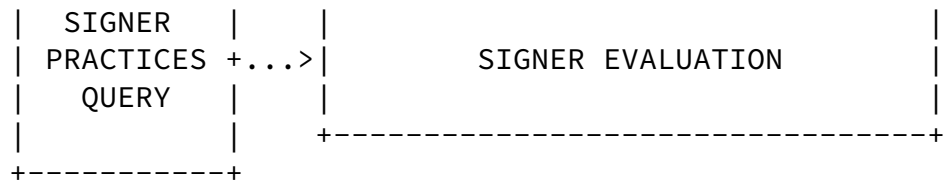

Definitions of some terms used in this document may be found in Appendix A.

Placeholder for some discussion of 2821 vs. 2822 solutions, etc.

2. The Bad Actors

2.1. Characteristics

The problem space being addressed by DKIM is characterized by a wide range of attackers in terms of motivation, sophistication, and capabilities.

At the low end of the spectrum are bad actors who may simply send email, perhaps using one of many commercially available tools, which the recipient does not want to receive. These tools may or may not falsify the origin address of messages, and may, in the future, be capable of generating message signatures as well.

At the next tier are what would be considered "professional" senders of unwanted email. These attackers would deploy specific infrastructure, including Mail Transfer Agents (MTAs), registered domains and possibly networks of compromised computers ("zombies") to send messages, and in some cases to harvest addresses to which to send. These senders often operate as commercial enterprises and send messages on behalf of third parties.

The most sophisticated and financially-motivated senders of messages are those who stand to receive substantial financial benefit, such as from an email-based fraud scheme. These attackers can be expected to employ all of the above mechanisms and additionally may attack the Internet infrastructure itself, e.g., DNS cache-poisoning attacks; IP routing attacks via compromised network routing elements.

2.2. Capabilities

In general, the bad actors described above should be expected to have access to the following:

1. An extensive corpus of messages from domains they might wish to impersonate

2. Knowledge of the business aims and model for domains they might wish to impersonate

3. Access to public keys and associated authorization records published by the domain

and the ability to do at least some of the following:

1. Submit messages to MTAs at multiple locations in the Internet

2. Construct arbitrary message headers, including those claiming to be mailing lists, resenders, and other mail agents

3. Sign messages on behalf of potentially-untraceable domains under their control

4. Generate substantial numbers of either unsigned or apparently-signed messages which might be used to attempt a denial of service attack

5. Resend messages which may have been previously signed by the domain

6. Transmit messages using any envelope information desired

As noted above, certain classes of bad actors may have substantial financial motivation for their activities, and therefore should be expected to have more capabilities at their disposal. These include:

1. Manipulation of IP routing. This could be used to submit messages from specific IP addresses or difficult-to-trace addresses, or to cause diversion of messages to a specific domain.

2. Limited influence over portions of DNS using mechanisms such as cache poisoning. This might be used to influence message routing, or to cause falsification of DNS-based key or policy advertisements.

3. Access to significant computing resources, perhaps through the conscription of worm-infected "zombie" computers. This could allow the bad actor to perform various types of brute-force attacks.

4. Ability to "wiretap" some existing traffic, perhaps from a wireless network.

Either of the first two of these mechanisms could be used to allow the bad actor to function as a man-in-the-middle between sender and recipient, if that attack is useful.

2.3. Location

In the following discussion, the term "administrative unit", taken from [I-D.crocker-email-arch], is used to refer to a portion of the email path that is under common administration. The originator and recipient typically develop trust relationships with the administrative units that send and receive their email, respectively, to perform the signing and verification of their messages.

Bad actors or their proxies can be located anywhere in the Internet. Certain attacks are possible primarily within the administrative unit of the claimed originator and/or recipient domain have capabilities beyond those elsewhere, as described in the below sections. Bad actors can also collude by acting in multiple locations simultaneously (a "distributed bad actor").

2.3.1. Externally-located Bad Actors

DKIM focuses primarily on bad actors located outside of the administrative units of the claimed originator and the recipient. These administrative units frequently correspond to the protected portions of the network adjacent to the originator and recipient. It

is in this area that the trust relationships required for

authenticated message submission do not exist and do not scale adequately to be practical. Conversely, within these administrative units, there are other mechanisms such as authenticated message submission that are easier to deploy and more likely to be used than DKIM.

External bad actors are usually attempting to exploit the "any to any" nature of email which motivates most recipient MTAs to accept messages from anywhere for delivery to their local domain. They may generate messages without signatures, with incorrect signatures, or with correct signatures from domains with little traceability. They may also pose as mailing lists, greeting cards, or other agents which legitimately send or re-send messages on behalf of others.

2.3.2. Within Claimed Originator's Administrative Unit

Bad actors in the form of rogue or unauthorized users or malware-infected computers can exist within the administrative unit corresponding to a message's origin address. Since the submission of messages in this area generally occurs prior to the application of a message signature, DKIM is not directly effective against these bad actors. Defense against these bad actors is dependent upon other means, such as proper use of firewalls, and mail submission agents that are configured to authenticate the sender.

In the special case where the administrative unit is non-contiguous (e.g., a company that communicates between branches over the external Internet), DKIM signatures can be used to distinguish between legitimate externally-originated messages and attempts to spoof addresses in the local domain.

2.3.3. Within Recipient's Administrative Unit

Bad actors may also exist within the administrative unit of the message recipient. These bad actors may attempt to exploit the trust relationships which exist within the unit. Since messages will typically only have undergone DKIM verification at the administrative unit boundary, DKIM is not effective against messages submitted in this area.

For example, the bad actor may attempt to apply a header such as Authentication-Results [I-D.kucherawy-sender-auth-header] which would normally be added (and spoofing of which would be detected) at the boundary of the administrative unit. This could be used to falsely indicate that the message was authenticated successfully.

As in the originator case, these bad actors are best dealt with by

controlling the submission of messages within the administrative unit. Depending on the characteristics of the administrative unit, cryptographic methods may or may not be needed to accomplish this.

3. Representative Bad Acts

One of the most fundamental bad acts being attempted is the delivery of messages which are not authorized by the alleged originating domain. As described above, these messages might merely be unwanted by the recipient, or might be part of a confidence scheme or a delivery vector for malware.

3.1. Use of Arbitrary Identities

This class of bad acts includes the sending of messages which aim to obscure the identity of the actual sender. In some cases the actual sender might be the bad actor, or in other cases might be a third-party under the control of the bad actor (e.g., a compromised computer).

DKIM is effective in mitigating against the use of addresses not controlled by bad actors, but is not effective against the use of addresses they control. In other words, the presence of a valid DKIM signature does not guarantee that the signer is not a bad actor. It also does not guarantee the accountability of the signer, since that is limited by the extent to which domain registration requires accountability for its registrants. However, accreditation and reputation systems can be used to enhance the accountability of DKIM-verified addresses and/or the likelihood that signed messages are desirable.

3.2. Use of Specific Identities

A second major class of bad acts involves the assertion of specific identities in email.

Note that some bad acts involving specific identities can sometimes be accomplished, although perhaps less effectively, with similar looking identities that mislead some recipients. For example, if the bad actor is able to control the domain "examp1e.com" (note the "one" between the p and e), they might be able to convince some recipients that a message from admin@examp1e.com is really admin@example.com. Similar types of attacks using internationalized domain names have been hypothesized where it could be very difficult to see character differences in popular typefaces. Similarly, if example2.com was

controlled by a bad actor, the bad actor could sign messages from bigbank.example2.com which might also mislead some recipients. To

the extent that these domains are controlled by bad actors, DKIM is not effective against these attacks, although it could support the ability of reputation and/or accreditation systems to aid the user in identifying them.

3.2.1. Exploitation of Social Relationships

One reason for asserting a specific origin address is to encourage a recipient to read and act on particular email messages by appearing to be an acquaintance or previous correspondent that the recipient might trust. This tactic has been used by email-propagated malware which mail themselves to addresses in the infected host's address book. In this case, however, the sender's address may not be falsified, so DKIM would not be effective in defending against this act.

It is also possible for address books to be harvested and used by an attacker to send messages from elsewhere. DKIM would be effective in mitigating these acts by limiting the scope of origin addresses for which a valid signature can be obtained when sending the messages from other locations.

3.2.2. Identity-Related Fraud

Bad acts related to email-based fraud often, but not always, involve the transmission of messages using specific origin addresses of other entities as part of the fraud scheme. The use of a specific address of origin sometimes contributes to the success of the fraud by helping convince the recipient that the message was actually sent by the alleged sender.

To the extent that the success of the fraud depends on or is enhanced by the use of a specific origin address, the bad actor may have significant financial motivation and resources to circumvent any measures taken to protect specific addresses from unauthorized use.

3.2.3. Reputation Attacks

Another motivation for using a specific origin address in a message

is to harm the reputation of another, commonly referred to as a "joe-job". For example, a commercial entity might wish to harm the reputation of a competitor, perhaps by sending unsolicited bulk email on behalf of that competitor. It is for this reason that reputation systems must be based on an identity that is, in practice, fairly reliable.

4. Attacks on Message Signing

Bad actors can be expected to exploit all of the limitations of message authentication systems. They are also likely to be motivated to degrade the usefulness of message authentication systems in order to hinder their deployment. Both the signature mechanism itself and declarations made regarding use of message signatures (often referred to as Sender Signing Policy, Sender Signing Practices or SSP) can be expected to be the target of attacks.

The sections below begin with a table summarizing the postulated attacks in each category along with their expected impact and likelihood. The following criteria were used in scoring the attacks against these criteria:

Impact:

High: Affects the verification of messages by an entire domain or multiple domains

Medium: Affects the verification of messages by specific users, MTAs, and/or bounded time periods

Low: Affects the verification of isolated individual messages only

Likelihood:

High: All users of DKIM should expect this attack on a frequent basis

Medium: Users of DKIM should expect this attack occasionally; frequently for a few users

Low: Attack is expected to be rare and/or very infrequent

4.1. Attacks Against Message Signatures

Summary of postulated attacks against DKIM signatures:

Attack Name	Impact	Likelihood
Theft of private key for domain	High	Low
Theft of delegated private key	Medium	Medium
Private key recovery via timing attack	High	Low
Chosen message replay	Low	M/H
Signed message replay	Low	High
Denial-of-service attack against verifier	High	Medium
Denial-of-service attack against key service	High	Medium
Canonicalization abuse	Low	Medium
Body length limit abuse	Medium	Medium
Use of revoked key	Medium	Low
Compromise of key server	High	Low
Falsification of key service replies	Medium	Medium
Publication of malformed key records and/or signatures	High	Low
Cryptographic weaknesses in signature generation	High	Low
Display name abuse	Medium	High

Compromised system within originator's network	Medium	Medium

4.1.1. Theft of Private Key for Domain

Message signing technologies such as DKIM are vulnerable to theft of the private keys used to sign messages. This includes "out-of-band" means for this theft, including burglary, bribery, extortion, and the like, as well as electronic means for such theft, such as a compromise of network and host security around the place where a private key is stored.

Keys which are valid for all addresses in a domain typically reside in MTAs which should be located in well-protected sites, such as data centers. Various means should be employed for minimizing access to private keys, such as non-existence of commands for displaying their value, although ultimately memory dumps and the like will probably contain the keys. Due to the unattended nature of MTAs, some countermeasures, such as the use of a pass phrase to "unlock" a key, are not practical to use.

4.1.2. Theft of Delegated Private Key

There are several circumstances where a domain owner will want to delegate the ability to sign messages for the domain to an individual user or a third-party associated with an outsourced activity such as a corporate benefits administrator or a marketing campaign. Since these keys may exist on less well-protected devices than the domain's own MTAs, they will in many cases be more susceptible to compromise.

In order to mitigate this exposure, keys used to sign such messages can be restricted by the domain owner to be valid for signing messages only on behalf of specific addresses in the domain. This maintains protection for the majority of addresses in the domain.

4.1.3. Private Key Recovery via Timing Attack

Timing attacks are a technique whereby the private key is recovered by observing the time required to sign a series of messages. It

requires both the ability to submit messages for signing as well as the ability to accurately measure the time required to compute the signature.

In most cases, an MTA has are enough variables (system load, clock resolution, queuing delays, etc.) to prevent the signing time from being measured accurately enough to be useful for a timing attack. Furthermore, while some domains, e.g., consumer ISPs, would allow an attacker to submit messages for signature, with many other domains this is difficult. Other mechanisms, such as mailing lists hosted by the domain, might be paths by which an attacker might submit messages for signature, and should also be considered as possible vectors for timing attacks.

4.1.4. Chosen Message Replay

Chosen Message Replay (CMR) refers to the scenario where the attacker creates a message and obtains a signature for it by sending it through an MTA authorized by the originating domain to him/herself or an accomplice. They then "replay" the signed message by sending it, using different envelope addresses, to a (typically large) number of other recipients.

Due to the requirement to get an attacker-generated message signed, Chosen Message Replay would most commonly be experienced by consumer ISPs or others offering email accounts to clients, particularly where there is little or no accountability to the account holder (the attacker in this case). One approach to this problem is for the domain to only sign email for clients that have passed a vetting process to provide traceability to the message originator in the

event of abuse. At present, the low cost of email accounts (zero) does not make it practical for any vetting to occur. It remains to be seen whether this will be the model with signed mail as well, or whether a higher level of trust will be required to obtain an email signature.

Revocation of the signature is a potential countermeasure. However, the rapid pace at which the message might be replayed (especially with an army of "zombie" computers), compared with the time required to detect the attack and implement the revocation, is likely to be problematic. A related problem is the likelihood that domains will

use a small number of signing keys for a large number of customers, which is beneficial from a caching standpoint but presents a problem revoking some signatures and not others. To this end, "revocation identifiers" have been proposed which would permit more fine-grained revocation, perhaps on a per-account basis. Messages containing these identifiers would result in a query to a revocation database, which might be represented in DNS. Further study is needed to determine if the benefits from revocation (given the potential speed of a replay attack) outweigh the transactional cost of querying the revocation database.

4.1.5. Signed Message Replay

Signed Message Replay (SMR) refers to the retransmission of already-signed messages to additional recipients beyond those intended by the sender. The attacker arranges to receive a message from the victim, and then retransmits it intact but with different envelope addresses. This might be done, for example, to make it look like a legitimate sender of messages is sending a large amount of spam. When reputation services are deployed, this could damage the originator's reputation.

A larger number of domains are potential victims of SMR than of CMR, because the former does not require the ability for the attacker to send messages from the victim domain. However, the capabilities of the attacker are lower. Unless coupled with another attack such as body length limit abuse, it isn't possible for the attacker to use this, for example, for advertising.

Many mailing lists, especially those which do not modify the content of the message and signed headers and hence do not invalidate the signature, engage in a form of SMR. The only things that distinguish this case from undesirable forms of SMR is the intent of the replayer, which cannot be determined by the network.

4.1.6. Denial-of-Service Attack Against Verifier

While it takes some compute resources to sign and verify a signature, it takes negligible compute resources to generate an invalid

signature. An attacker could therefore construct a "make work" attack against a verifier, by sending a large number of incorrectly-signed messages to a given verifier, perhaps with multiple signatures each. The motivation might be to make it too expensive to verify messages.

While this attack is feasible, it can be greatly mitigated by the manner in which the verifier operates. For example, it might decide to accept only a certain number of signatures per message, limit the maximum key size it will accept (to prevent outrageously large signatures from causing unneeded work), and verify signatures in a particular order.

4.1.7. Denial-of-Service Attack Against Key Service

An attacker might also attempt to degrade the availability of an originator's key service, in order to cause that originator's messages to be unverifiable. One way to do this might be to quickly send a large number of messages with signatures which reference a particular key, thereby creating a heavy load on the key server. Other types of DoS attacks on the key server or the network infrastructure serving it are also possible.

The best defense against this attack is to provide redundant key servers, preferably on geographically-separate parts of the Internet. Caching also helps a great deal, by decreasing the load on authoritative key servers when there are many simultaneous key requests. The use of a key service protocol which minimizes the transactional cost of key lookups is also beneficial. It is noted that the Domain Name System has all these characteristics.

4.1.8. Canonicalization Abuse

Canonicalization algorithms represent a tradeoff between the survival of the validity of a message signature and the desire not to allow the message to be altered inappropriately. In the past, canonicalization algorithms have been proposed which would have permitted attackers, in some cases, to alter the meaning of a message.

Message signatures which support multiple canonicalization algorithms give the signer the ability to decide the relative importance of signature survivability and immutability of the signed content. If an unexpected vulnerability appears in a canonicalization algorithm

in general use, new algorithms can be deployed, although it will be a slow process because the signer can never be sure which algorithm(s) the verifier supports. For this reason, canonicalization algorithms, like cryptographic algorithms, should undergo a wide and careful review process.

4.1.9. Body Length Limit Abuse

A body length limit is an optional indication from the signer how much content has been signed. The verifier can either ignore the limit, verify the specified portion of the message, or truncate the message to the specified portion and verify it. The motivation for this feature is the behavior of many mailing lists which add a trailer, perhaps identifying the list, at the end of messages.

When body length limits are used, there is the potential for an attacker to add content to the message. It has been shown that this content, although at the end, can cover desirable content, especially in the case of HTML messages.

If the body length isn't specified, or if the verifier decides to ignore the limit, body length limits are moot. If the verifier or recipient truncates the message at the signed content, there is no opportunity for the attacker to add anything.

If the verifier observes body length limits when present, there is the potential that an attacker can make undesired content visible to the recipient. The size of the appended content makes little difference, because it can simply be a URL reference pointing to the actual content. Recipients need to use means to, at a minimum, identify the unsigned content in the message.

4.1.10. Use of Revoked Key

The benefits obtained by caching of key records opens the possibility that keys which have been revoked may be used for some period of time after their revocation. The best examples of this occur when a holder of a key delegated by the domain administrator must be unexpectedly deauthorized from sending mail on behalf of one or more addresses in the domain.

The caching of key records is normally short-lived, on the order of hours to days. In many cases, this threat can be mitigated simply by setting a short time-to-live for keys not under the domain administrator's direct control (assuming, of course, that control of the time-to-live value may be specified for each record, as it can with DNS). In some cases, such as the recovery following a stolen

private key belonging to one of the domain's MTAs, the possibility of

theft and the time required to revoke the key authorization must be considered when choosing a TTL. The chosen TTL must be long enough to mitigate denial-of-service attacks and provide reasonable transaction efficiency, and no longer.

4.1.11. Compromise of Key Server

Rather than by attempting to obtain a private key, an attacker might instead focus efforts on the server used to publish public keys for a domain. As in the key theft case, the motive might be to allow the attacker to sign messages on behalf of the domain. This attack provides the attacker with the additional capability to remove legitimate keys from publication, thereby denying the domain the ability for the signatures on its mail to verify correctly.

The host which is the primary key server, such as a DNS master server for the domain, might be compromised. Another approach might be to change the delegation of key servers at the next higher domain level.

This attack can be mitigated somewhat by independent monitoring to audit the key service. However, it may be difficult to detect the publication of additional keys by such means until the selector(s) added by the attackers are known.

4.1.12. Falsification of Key Service Replies

Replies from the key service may also be spoofed by a suitably positioned attacker. For DNS, one such way to do this is "cache poisoning", in which the attacker provides unnecessary (and incorrect) additional information in DNS replies, which is cached.

DNSSEC [RFC4033] is the preferred means of mitigating this threat, but the current uptake rate for DNSSEC is slow enough that one would not like to create a dependency on its deployment. Fortunately, the vulnerabilities created by this attack are both localized and of limited duration, although records with relatively long TTL may be created with cache poisoning.

4.1.13. Publication of Malformed Key Records and/or Signatures

In this attack, the attacker publishes suitably crafted key records or sends mail with intentionally malformed signatures, in an attempt to confuse the verifier and perhaps disable verification altogether. This attack is really a characteristic of an implementation vulnerability, a buffer overflow or lack of bounds checking, for example, rather than a vulnerability of the signature mechanism itself. This threat is best mitigated by careful implementation and creation of test suites that challenge the verification process.

4.1.14. Cryptographic Weaknesses in Signature Generation

The cryptographic algorithms used to generate mail signatures, specifically the hash algorithm and the public-key encryption/ decryption operations, may over time be subject to mathematical techniques that degrade their security. At this writing, the SHA-1 hash algorithm is the subject of extensive mathematical analysis which has considerably lowered the time required to create two messages with the same hash value. This trend can be expected to continue.

The message signature system must be designed to support multiple signature and hash algorithms, and the signing domain must be able to specify which algorithms it uses to sign messages. The choice of algorithms must be published in key records, rather than in the signature itself, to ensure that an attacker is not able to create signatures using algorithms weaker than the domain wishes to permit.

Due to the fact that the signer and verifier of email do not, in general, communicate directly, negotiation of the algorithms used for signing cannot occur. In other words, a signer has no way of knowing which algorithm(s) a verifier supports, nor (due to mail forwarding) where the verifier is. For this reason, it is expected that once message signing is widely deployed, algorithm change will occur slowly, and legacy algorithms will need to be supported for a considerable period. Algorithms used for message signatures therefore need to be secure against expected cryptographic developments several years into the future.

4.1.15. Display Name Abuse

Message signatures only relate to the address-specification portion of an email address, which some MUAs only display (or some recipients

only pay attention to) the display name portion of the address. This inconsistency leads to an attack where the attacker uses an From header field such as:

From: "Dudley DoRight" <whiplash@example.org>

In this example, the attacker, whiplash@example.org, can sign the message and still convince some recipients that the message is from Dudley DoRight, who is presumably a trusted individual. Coupled with the use of a throw-away domain or email address, it may be difficult to bring the attacker to account for the use of another's display name.

This is an attack which must be dealt with in the recipient's MUA. One approach is to require that the signer's address specification

(and not just the display name) be visible to the recipient.

4.1.16. Compromised System Within Originator's Network

In many cases, MTAs may be configured to accept, and sign, messages which originate within the topological boundaries of the originator's network (i.e., within a firewall). The increasing use of compromised systems to send email presents a problem for such policies, because the attacker, using a compromised system as a proxy, can generate signed mail at will.

Several approaches exist for mitigating this attack. The use of authenticated submission, even within the network boundaries, can be used to limit the addresses for which the attacker may obtain a signature. It may also help locate the compromised system that is the source of the messages more quickly. Content analysis of outbound mail to identify undesirable and malicious content, as well as monitoring of the volume of messages being sent by users, may also prevent arbitrary messages from being signed and sent.

4.2. Attacks Against Message Signing Policy

Summary of postulated attacks against signing policy:

Attack Name	Impact	Likelihood

Look-alike domain names	High	High
Internationalized domain name abuse	High	Medium
Denial-of-service attack against signing policy	Medium	Medium
Use of multiple From addresses	Low	Medium

4.2.1. Look-Alike Domain Names

Attackers may attempt to circumvent signing policy of a domain by using a domain name which is close to, but not the same as the domain with a signing policy. For instance, "example.com" might be replaced by "examp1e.com". If the message is not to be signed, DKIM does not require that the domain used actually exist (although other mechanisms may make this a requirement). Services exist to monitor domain registrations to identify potential domain name abuse, but naturally do not identify the use of unregistered domain names.

4.2.2. Internationalized Domain Name Abuse

Internationalized domain names present a special case of the look-

alike domain name attack described above. Due to similarities in the appearance of many Unicode characters, domains (particularly those drawing characters from different groups) may be created which are visually indistinguishable from other, possibly high-value domains. This is discussed in detail in Unicode TR 36 [UTR36]. Surveillance of domain registration records may point out some of these, but there are many such similarities. As in the look-alike domain attack above, this technique may also be used to circumvent sender signing policy of other domains.

4.2.3. Denial-of-Service Attack Against Signing Policy

Just as the publication of public keys by a domain can be impacted by an attacker, so can the publication of Sender Signing Policy (SSP) by a domain. In the case of SSP, the transmission of large amounts of unsigned mail purporting to come from the domain can result in a heavy transaction load requesting the SSP record. More general DoS attacks against the servers providing the SSP records are possible as well. This is of particular concern since the default signing policy

is "we don't sign everything", which means that SSP, in effect, fails open.

As with defense against DoS attacks for key servers, the best defense against this attack is to provide redundant servers, preferably on geographically-separate parts of the Internet. Caching again helps a great deal, and signing policy should rarely change, so TTL values can be relatively large.

4.2.4. Use of Multiple From Addresses

Although this usage is rare, RFC 2822 [RFC2822] permits the From address to contain multiple address specifications. The lookup of Sender Signing Policy is based on the From address, so if addresses from multiple domains are in the From address, the question arises which signing policy to use. A rule (say, "use the first address") could be specified, but then an attacker could put a throwaway address prior to that of a high-value domain. It is also possible for SSP to look at all addresses, and choose the most restrictive rule. This is an area in need of further study.

5. Derived Requirements

This section, as yet incomplete, is an attempt to capture a set of requirements for DKIM from the above discussion. These requirements include:

The store for key and SSP records must be capable of utilizing multiple geographically-dispersed servers.

Key and SSP records must be cacheable, either by the verifier requesting them or by other infrastructure.

The cache time-to-live for key records must be specifiable on a per-record basis.

The algorithm(s) used by the signing domain associated with a given key must be specified independently of the signature itself.

6. IANA Considerations

This document defines no items requiring IANA assignment.

7. Security Considerations

This document describes the security threat environment in which DomainKeys Identified Mail (DKIM) is expected to provide some benefit, and presents a number of attacks relevant to its deployment.

8. Informative References

[I-D.allman-dkim-base]
Allman, E., "DomainKeys Identified Mail (DKIM)", draft-allman-dkim-base-01 (work in progress), October 2005.

[I-D.allman-dkim-ssp]
Allman, E., "DKIM Sender Signing Policy", draft-allman-dkim-ssp-01 (work in progress), October 2005.

[I-D.crocker-email-arch]
Crocker, D., "Internet Mail Architecture", draft-crocker-email-arch-04 (work in progress), March 2005.

[I-D.kucherawy-sender-auth-header]
Kucherawy, M., "Message Header for Indicating Sender Authentication Status", draft-kucherawy-sender-auth-header-02 (work in progress), May 2005.

[RFC2822] Resnick, P., "Internet Message Format", RFC 2822, April 2001.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[UTR36] Davis, M. and M. Suignard, "Unicode Security

Considerations", UTR 36, July 2005.

Appendix A. Glossary

Origin address - The address on an email message, typically the RFC 2822 From: address, which is associated with the alleged author of the message and is displayed by the recipient's MUA as the source of the message.

More definitions to be added.

Appendix B. Acknowledgements

The author wishes to thank Phillip Hallam-Baker, Eliot Lear, Tony Finch, Dave Crocker, Barry Leiba, Arvel Hathcock, Eric Allman, Jon Callas, and Stephen Farrell for valuable suggestions and constructive criticism of earlier versions of this draft.

Appendix C. Edit History

Changes since -00 draft:

- Changed beginning of introduction to make it consistent with -base draft.
- Clarified reasons for focus on externally-located bad actors.
- Elaborated on reasons for effectiveness of address book attacks.
- Described attack time windows with respect to replay attacks.
- Added discussion of attacks using look-alike domains.
- Added section on key management attacks.

Changes since -01 draft:

- Reorganized description of bad actors.

- Greatly expanded description of attacks against DKIM and SSP.
- Added "derived requirements" section.

Author's Address

Jim Fenton
Cisco Systems, Inc.
MS SJ-24/2
170 W. Tasman Drive
San Jose, CA 95134-1706
USA

Phone: +1 408 526 5914
Email: fenton@cisco.com
URI:

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.