pre-workgroup
Internet-Draft
Expires: April 24, 2006

J. Fenton
Cisco Systems, Inc.

October 21, 2005

Analysis of Threats Motivating DomainKeys Identified Mail (DKIM)
draft-fenton-dkim-threats-01

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.

2. The Bad Actors

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 in addition attacks on Internet infrastructure itself, such as DNS cache-poisoning attacks and IP routing attacks via compromised network routing elements.

3. Capabilities of the Bad Actors

3.1. General 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 key 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

3.2. Advanced capabilities

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.

4. Location of the Bad Actors

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. Bad actors 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").

4.1. 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 wasier 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.

4.2. 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.

4.3. 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.

5. 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.

5.1. 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.

5.2. 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.

5.2.1. 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.

5.2.2. 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 convincing 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.

5.2.3. 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.

Reputation attacks of this sort are sometimes based on the retransmission (often referred to as a "replay") of a legitimately sent message. DKIM provides little protection against such acts, although the key used to sign the original instance of the message can be revoked, which limits the time window available for such attacks. Other reputation attacks, involving the fabrication and transmission of a fictitious message, are addressed by DKIM since the

bad actor would not, without inside assistance, be able to obtain a valid signature for the fabricated message.

6. 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. Some representatives of these categories of bad acts are described below. Additional postulated attacks are described in the Security Considerations section of [I-D.allman-dkim-base].

6.1. Unsigned Messages

Messages without signatures may be sent in an effort to exploit the incremental deployment of message signatures. In many cases, a recipient may not be able to make a determination about unsigned messages from a domain, and therefore will need to accept the message (although perhaps at a lower delivery priority). This situation is mitigated by the use of the DKIM Sender Signing Policy (SSP) [I-D.allman-dkim-ssp] that indicates whether or not a given domain signs all of its messages. Nevertheless, the possibility of signature breakage due to legitimate modification of the message may limit the ability of SSP to dictate harsh treatment of messages without valid signatures.

Messages with invalid signatures may also be introduced by bad actors. The intent may be to make the message appear as though it was legitimately sent, but "broken" in transit, i.e. that the message was modified, rendering the signature invalid. At least until the causes of signature breakage are well understood, messages with invalid signatures need to be evaluated as if the invalid signature isn't present at all.

6.2. Use of Throw-Away Addresses

Bad actors may also introduce messages with valid signatures on behalf of domains they control, perhaps "throw-away" domains registered under false pretenses which are difficult to trace. In other words, the existence of a message signature does not imply that the message is "good". The use of such domains will undoubtedly give rise to domain-based accreditation and reputation systems. Until these are available, local reputation, mostly in the form of whitelists, can be maintained by domains to improve the deliverability of email from domains with which they have business or other relationships.

Accreditation and reputation, or even local whitelists, require a reliable identity on which to base their assertion, and in the case of reputation on which to base any feedback reports. Message signing provides an identity which is intended to be sufficiently reliable for this purpose, and it (or some other reliable mechanism) is necessary for accreditation and reputation systems to operate.

Modification of messages by mailing lists and other legitimate agents requires that a mechanism be created for signing of messages by other than the originating domain. This provides a bad actor with an additional avenue through which it might attempt to circumvent message authentication. A bad actor might attempt to pose as a mailing list which modifies a message and adds its own signature taking responsibility for the message. If this signature is from an untraceable domain, little assertion of the legitimacy of the message is provided by this signature. For this reason, accreditation, reputation, and local reputation in the form of white lists is at least as important for these signatures from third parties as they are for origination address signatures.

6.3. Message Replay

Message replay is a term used to describe the retransmission of already-signed messages. Here the bad actor obtains an account with a domain such as a consumer ISP, and sends an undesirable message to an external address controlled by the bad actor or an accomplice. That message, having now obtained a signature, is forwarded to other recipients without the authorization of the signing domain. It is closely related to one of the reputation attacks described above.

This bad act is basically indistinguishable from a number of acceptable acts, such as the transparent forwarding of messages by a recipient to multiple addresses. For this reason, DKIM is not particularly effective at detecting and eliminating this bad act. Prompt key revocation may mitigate this problem; however, since verification typically occurs as messages are received by recipient domains, time is of the essence.

Other means to mitigate this bad act include the use of content filtering on messages being signed, and business models which enforce more accountability for subscribers whose messages are to be signed by DKIM.

6.4. Control of Key Management

In cases where the Bad Actor is in control of the DNS zone for a domain's keyspace subdomain (_domainkey), or is able to influence the records in that zone, message signing may present a new opportunity

to interfere with the receipt of legitimate messages or to sign messages not illegitimately. This could occur if the DNS provider for the domain is not reliable or if the security measures used by the DNS provider are breached by the bad actor. DKIM is fully dependent on the key information which it is provided by DNS, so independent means such as audits of the key records would be required to mitigate this threat.

7. IANA Considerations

This document defines no items requiring IANA assignment.

8. Security Considerations

This document describes the security threat environment in which DomainKeys Identified Mail (DKIM) is expected to provide some benefit.

9. Informative References

[I-D.allman-dkim-base]
Allman, E., "DomainKeys Identified Mail (DKIM)", draft-allman-dkim-base-00 (work in progress), July 2005.

[I-D.allman-dkim-ssp]
Allman, E., "DKIM Sender Signing Policy", draft-allman-dkim-ssp-00 (work in progress), July 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.

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.

Appendix B. Acknowledgements

The author wishes to thank Phillip Hallam-Baker, Eliot Lear, Tony Finch, Dave Crocker, Barry Leiba, Arvel Hathcock, Eric Allman, and Jon Callas 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.

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.