



**TREND**  
M I C R O™

Douglas Otis  
Trend Micro, Inc. NSSG  
[Doug\\_Otis@trendmicro.com](mailto:Doug_Otis@trendmicro.com)

## Can Trust in Email Be Restored?

Without identifying by name who can be held accountable, a common practice protects email's utility as an efficient form of communication. This practice blocks abusive sources based upon the IP address. By using anonymous addresses, this stifles development of trust and accountability by name. Abuse is often discerned when much of the email from an IP address has been sent to recipients that did not expressed a desire for its receipt. In the email vernacular, this practice is called "block-listing" when the source has not adhered to an "opt-in" criteria for the distribution of bulk email. Although an "opt-in" criteria may impose differing sending limits, this criteria is essential as a practical means to ascertain which sources cause the greatest harm to email's utility. A better method has not yet been devised, even though when based upon the IP address, the source of the message remains unnamed.

## CAN-SPAM Demonstrates DMA's Influence

Some email marketers advocate a wholly unworkable and unsafe "opt-out" method, codified in the US Federal CAN-SPAM law, which depends upon the source of the message being trustworthy. Fortunately, a practical "opt-in" criteria may still be used as a matter of policy instead. Opting-out greatly increases both the recipient's burden and risks. Opting-out requires responding to the flood of undesired messages. Merely responding increases the trading value of the now "verified" email-address, which may invite more messages from a series of unscrupulous senders. When the "opt-out" method is offered as a link, just clicking on the link may compromise the recipient's system. The "opt-out" method is too onerous and dangerous to be considered a means for protecting email's utility. Applying "opt-in" criteria, as a matter of the recipient's policy, is preferable over the destruction of email's utility.

## Source Identifiers

There are only a few verifiable source identifiers within email where trust can be established. The sender's IP address or the host-name that is provided in the HELO announcement are two verifiable source identifiers. A new source identification method based upon public-key cryptography is being tested called DomainKeys (DK), along with an ongoing IETF effort called DomainKeys Identified Mail (DKIM). Although DK or DKIM provide a source identifier by name that can transverse several stages of delivery, both methods exclude the message envelope from the signature, which limits the scope of named source's accountability.

## Arbiter of Abuse for DK/DKIM Signed Messages

Normally the message envelope, which lists the recipients, provides a basis for assessing abuse. This process typically uses the "opt-in" criteria, but

**Trend Micro Incorporated**

the envelope's information unfortunately can not be safely ascribed to a signing-domain. Bad actors can modify the envelope of a message without affecting the validity of the message's signature. Therefore, ascribing abuse to the signing-domain, based upon messages being sent to unwilling recipients, may corrupt behavioral information accrued for signing-domains.

Since the DK/DKIM signature excludes the message envelope, no assurance related to the number of messages, intended recipients, or the return-path is provided. Just the domain introducing the signed portion of the message can be verified. Regardless of the number of unwilling recipients, it may not be possible to ascribe all abusive messages to the signing-domain. When limited to assessing just signed content, this process will be resource intensive. The bad acts of a signing-domain will be confined only to those of a criminal nature, such as offering malware or dangerously misleading information and links.

## Vetted Signing Domains

Signatures by themselves do not provide a safe means for white-listing due to a potential for message replay abuse. Signatures are also expensive and perilous when used for block-listing. Although a DK/DKIM signature verifies the initial message source by name, assurances conveyed to the recipient should preclude messages from unknown signing-domains which might be controlled by bad actors. Rather than presuming a domain is trustworthy until proven otherwise, a vetted list of trusted signing-domains should qualify any assurance being conveyed. Due to the nature of content assessment, vetting may be unable to respond fast enough to deter abuse of such assurances. This is unlike the simpler methods used with block-listing, which presumes the source does not abuse until proven otherwise.

## Trust-Marks Safely Assure Recipients

A new trust-mark can safely convey that a message is from a "trusted" signing-domain. Indications that an email-address is associated with a signing-domain is something easily accomplished by any bad actor controlling a signing-domain. Conveying just the existence of an email-address association with an email identifier dangerously presumes the recipient is capable of a rather perilous examination of the email-address.

Recipients often only see the "display-name" and do not know which Unicode character repertoires are being used. ASCII Compatible Encoding of international domain names (ACE labels) are potentially exploited regardless of the form shown to the recipient. In addition, many recipients have little knowledge of the domain-name hierarchy. Therefore, increased use of sub-domains would be counter productive, as few recipients understand the difference between a hyphen and a period in the domain name.

## Retaining Trust with Key Tags

Retaining trust by the signing-domain demands stringent control of all messages that are to receive an assurance of trust, since a signed errant message may be replicated any number of times. Using the signature to bestow a trust-mark may conflict with a policy of signing all messages, because perhaps not all users and MTAs sending messages have been properly vetted or secured to safely ensure a retention of trust.

### **Trend Micro Incorporated**

Although not currently a feature of DK or DKIM, tags added to the signature verification keys can indicate whether the signing-domain considers the source of the message trustworthy. The signer's trust assertion provides a method for the domain to retain trust in their vetted and secured sources, while still signing all messages. Messages signed by trustworthy domains, but where the key is not also tagged trustworthy, should not convey any assurances of "trusted" to the recipient, unless overridden by a local database maintained by the recipient.

For example, a signing-domain may consider delegated keys too risky to tag as trustworthy. Until discovered, and for as long as the TTL of a key allows retention within a DNS cache, a bad actor could continue to send deceptive signed messages damaging trust in the domain, especially when there is reliance upon trust-marks. To protect the trust of the signing-domain, messages being sent from unvetted users or insecure systems should not be signed with keys tagged as trustworthy. Tagged keys could be independent of the email-address, as in the case of insecure systems.

Providers offering low cost or free services will not be able to adequately vet their many users, who often number in the millions. The provider may wish to send special messages from the same domain and have these messages receive a trustworthy evaluation. Perhaps these could be messages from a system administrator indicating that the user's system appears compromised, or that their payment information is no longer valid. The use of trusted/non-trusted key tags would permit a signing-domain to both retain their trustworthy status and to prevent other unvetted users within the same domain from spoofing their own customers with "trusted" messages.

## Most Email is Not Trustworthy

With the high cost of evaluating a violation of trust by a signing-domain, constrained use of "trusted" keys should significantly reduce the burdens placed upon a signing-domain evaluation process. Assessing message content is resource intensive, when compared to the typical envelope-based "opt-in" criteria, often used to determine bad acts. Assessing verifiable actions of a signing-domain is also prolonged when it involves human interaction across hundreds of languages. To help alleviate this situation, key tags can significantly minimize evaluation costs and delays.

## Message Acceptance Based Upon the Signature

Once a bad actor obtains an account, they can send themselves signed abusive messages. These signed messages can then be re-sent in bulk elsewhere, while still retaining valid original signatures. Most domains will be exposed to this simple replay exploit. When the recipients are able to associate a message identifier with the signing-domain, this qualifies the message as not being replayed. A verified HELO could be one such associated identifier that does not restrict the email-addresses being signed.

For security reasons, each MTA should be assigned a unique message verification key. Setting up a DNS resource record to verify the HELO could be a simple SRV resource record as defined by the CSV draft, see <http://mipassoc.org/csv>. The SRV resource record provides for the lowest overhead when making a signature association. In contrast, compiling addresses from SPF records may require hundreds of DNS queries, and these records may limit flexibility in which email-addresses can be signed. A

### **Trend Micro Incorporated**

verifiable HELO is much easier to implement as it does not require the MTA administrator to list all the IP addresses currently used to carry a domain's messages.

## Delayed Acceptance

Using SPF records, referenced from some email-address within the signing-domain, may establish an association with the signing-domain and also indicate the message is not being replayed. However, there is risk these SPF records may disrupt the delivery of some emails, while also incurring higher overhead. There is no perfect solution, as any technique to associate an email identifier with the signing-domain may fail when valid messages transverse a mediator. When a signing-domain association is not found, a protective strategy can be to delay acceptance of the message.

Delay in acceptance can be accomplished with a Transient Negative Completion, in conjunction with "Requested action aborted: error in processing" SMTP response; see RFC2821. Although message envelope assessments can be made rapidly, distributing these assessments still requires time. A delay in acceptance allows time for abusive message replays to be discovered and the message source identifiers to be block-listed. A receiving domain may eliminate the delayed acceptance, once the source is considered trustworthy. This strategy would provide protection from threats related to abusive message replay, and yet still allow messages to transverse mediators.

## Email-address Assessment is Problematic

Bad actors can not be identified based upon the email-address domain alone. Only when the email-address domain can be directly associated with the signing-domain, has the domain identifier been verified. However, this verification would then represent the same name as that of the signing-domain. Therefore, when there is a problem, the signing-domain should be held accountable, regardless of any email-address domain associations.

Oddly, the current DKIM proposal directs reports to the email-address domain owner instead of the signing-domain. When the email-address domain is not within a signing-domain, DKIM reporting still assumes only the email-address domain owner is interested in receiving problem reports. This conclusion is questionable, as the hapless email-address domain owner may disregard or be frustrated by reports when they have no means to either trace or to correct the underlying causes. Only the signing-domain is able to take corrective action and curtail abuse.

As with restoring trust, fair accounting must be based upon verified source identifiers. An email-address domain owner's defense from unfair accounting might be to publish close-ended policies that restrict the use of their email-address. This response to unfair accounting erodes the freedoms many enjoy today where their email-address is allowed to transverse the diverse email infrastructure. This freedom allows individuals to communicate using list-servers, e-invites, social aliases, etc. The alternative of publishing open-ended policies offers no protection from unfair accounting, and may even be seen erroneously as actually confirming the email-address domain owner's culpability. The email-address domain owner should therefore be cautious about publishing any email-address policies which might be viewed as unfairly confirming their culpability or causing a disruption in the delivery of their messages.

### **Trend Micro Incorporated**

Assessing signing-domain related behaviors is expensive. Resolving behaviors for an unlimited number of individual email-addresses within the domain only further increases this expense. Therefore, accountability will shift to email-address domain when messages obtain higher ratings based upon the discovery of an email-address related policy. Increased ratings for discovered policies may be viewed as promoting the publishing of email-address policy records. However, the increased rating will also likely be unfairly compensated by an assessment of the email-address domain, rather than the signing-domain. Ironically, promoting the publication of policy records through the use of message ratings also invites abuse that targets the email-address domains that publish open-ended policies.

## Cryptography Requires Auxiliary Defenses

The effort to restore trust does not end with just the DK/DKIM cryptographic signature. The process itself must be defended. Cryptographic techniques represent a moderate overhead where messages must be fully received before the validity of a signature can be verified. This means the cryptographic process is somewhat more vulnerable to Denial of Service attacks than alternative defensive schemes that can be used in conjunction with DK/DKIM. These defensive schemes identify sources based upon the readily available IP address or host-name. However, depending upon the IP address may cause collateral blocking when servers are being shared, as they often are. Fortunately, email already offers a solution for the Denial of Service attack, collateral blocking, and detecting possible message replay. At the beginning of an email exchange session, the host-name of the sending system is provided in the HELO. A new paradigm should be established that ensures the host-name can be verified and associated with the signing-domain. Verifying the HELO permits the same name-based reputations that are already used to vet the message sources, to also defend the cryptographic process.

## In Conclusion

To restore trust in email sources and to minimize the false detection of spoofing attempts, consider implementing DomainKeys. Implement DKIM once it is accepted, where both DomainKeys and DKIM verification will need to be supported during the transition. Establish conventions for tagging keys and verifying HELOs, which are still needed items. Remain cautious about publishing email-address policy records, which can easily disrupt normal uses of email.

Email-address polices can not be fully effective at preventing criminal spoofing attempts, unlike the use of trust-marks that can be safely established with the use of DomainKeys and then DKIM. Reliance upon email-address recognition rather than trust-marks, while not as effective, also invites additional costs when acquiring look-alike domains. Lastly, DK/DKIM related issues should be reported to the signing-domain and not to the email-address domain owner. Slight adjustments to DK/DKIM will ensure email becomes both a trusted and efficient form of communication.