MARID Working Group
Internet Draft
Document: draft-ietf-marid-core-03.txt

J. Lyon
Microsoft Corp
M. Wong
pobox.com

Expires: February 2005

August 2004

Sender ID: Authenticating E-Mail

Status of this Memo

Abstract

Internet mail suffers from the fact that much unwanted mail is sent using spoofed addresses -- "spoofed" in this case means the address is used without the permission of the domain owner. This document describes the following: mechanisms by which a domain owner can publish its set of outgoing MTAs, mechanisms by which SMTP servers can determine what email address is allegedly responsible for most proximately introducing a message into the Internet mail system, and whether that introduction is authorized by the owner of the domain contained in that email address.

The specification is carefully tailored to ensure that the overwhelming majority of legitimate emailers, remailers and mailing list operators are already compliant.

Table of Contents

1. Introduction....3
2. Problem Statement....3
2.1 Positive Problem Statement....3
2.2 Negative Problem Statement....4
3. Decision Model....4
4. Determining the Purported Responsible Address....5
5. Actions Based on the Decision....5
5.1 Neutral or None or PermError....6
5.2 Pass....6
5.3 Fail....6
5.4 SoftFail....6
5.5 TempError....6
6. Security Considerations....6
6.1 DNS Attacks....7
6.2 TCP Attacks....7
6.3 Forged Sender Attacks....7
6.4 Address Space Hijacking....7
7. Implementation Guidance....8
7.1 Simple E-mailers....8
7.2 E-Mail Forwarders....8
7.3 Mailing List Servers....8
7.4 Third-Party Mailers....9
7.5 MUA Implementers....9
8. IANA Considerations....9
9. Acknowledgements....9
10. References....10
10.1 Normative References....10
10.2 Informative References....10
11. Authors' Addresses....11

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1. Introduction

Today, a huge majority of unwanted email contains headers that lie about the origin of the mail. This is true of most spam and substantially all of the virus email that is sent.

This document describes a mechanism such that receiving MTAs, MDAs and/or MUAs can recognize mail in the above category and take appropriate action. For example, an MTA might refuse to accept a message, an MDA might discard a message rather than placing it into a mailbox, and an MUA might render that message in some distinctive fashion.

In order to avoid further fragmentation of the Internet email system, it is necessary that the Internet community as a whole come to a consensus as to what mail senders should do to make their mail appear non-spoofed, and how mail receivers should determine whether mail is spoofed. On the other hand, it is not necessary to reach a consensus regarding the actions that various parties take once a message has been determined to be spoofed. This can be done unilaterally -- one agent might decide to discard a spoofed message while another decides to add a disclaimer.

2. Problem Statement

2.1 Positive Problem Statement

Briefly stated, the mechanisms of this document allow one to answer the following question:

When a message is transferred via SMTP between two UNRELATED parties, does the SMTP client host have permission to send mail on behalf of the mailbox that allegedly caused the most recent introduction of the message into the mail delivery system?

As seen from the question, this mechanism applies to unrelated parties: it is useful at the point where a message passes across the Internet from one organization to another. It is beyond the scope of this document to describe authentication mechanisms that can be deployed within an organization.

The mechanism of this document also seeks to authenticate the mailbox associated with the MOST RECENT introduction of a message into the mail delivery system. In simple cases, this is who the mail is from. However, in the case of a third-party mailer, a forwarder or a mailing list server, the address being authenticated is that of the third party, the forwarder or the mailing list.

This document provides means to authenticate the DOMAIN of the appropriate email address; it is not directed at the local-part. A domain owner gets to determine which SMTP clients speak on behalf of addresses within the domain; a responsible domain owner should not authorize SMTP clients that will lie about local parts.

In the long run, once the domain of the sender is authenticated, it will be possible to use that domain as part of a mechanism to determine the likelihood that a given message is spam, using, for example, reputation and accreditation services. (These services are not the subject of the present mechanism, but it should enable them.)

2.2 Negative Problem Statement

Following are several alternate questions, which this specification makes no attempt to answer:

1. Is the host at a particular IP address authorized to act as an SMTP client?

2. Is an SMTP client authorized to use a particular domain name in its SMTP EHLO command?

3. Is an SMTP client authorized to use a particular email address in an SMTP "MAIL FROM:" command?

4. Was a message really authored by who it claims to be authored by?

3. Decision Model

The essence of this specification is:

Given an email message, and given an IP address from which it has been (or will be) received, is the SMTP client at that IP address authorized to send that email message?

This question will usually be asked by an SMTP server as part of deciding whether to accept an incoming mail message. However, this question could also be asked later by a different party. An MUA, for example, could use the result of this question to determine how to file or present a message.

There are four steps to answering this question:

(1) From the headers of the email message, extract the "purported responsible address". This is the mailbox that the message claims is responsible for the most recent introduction of the message into the delivery system. This step is described in detail in section 4 below. A separate specification, [Submitter], describes an SMTP extension that allows an SMTP server to perform this check at the time of the SMTP MAIL command instead of the SMTP DATA command.

(2) Extract the domain part of the purported responsible address. Call this the "purported responsible domain".

(3) Call the check_host function defined in [Protocol], passing the following parameters:
a. The IP address (either IPv4 or IPv6) from which the message is being or has been received.
b. The purported responsible domain from step (2) above.
c. The purported responsible address from step (1) above.

The result of the check_host function is one of the values "Neutral", "Pass", "Fail", "SoftFail", "None", "TempError" or "PermError". Section 5 describes how these results are used by MTAs receiving messages. This specification imposes no requirements on parties performing this test in other environments.

4. Determining the Purported Responsible Address

The purported responsible address (PRA) of a message MUST be determined using the algorithm described in [PRA].

If the Sender ID check is being performed by an MTA as part of receiving an e-mail message, and the PRA algorithm cannot determine a PRA, then the message SHOULD be rejected with error "550 5.1.7 Missing purported Responsible Address".

5. Actions Based on the Decision

When the Sender ID test is used by an SMTP server as part of receiving a message, the server should take the actions described by this section.

The check_host function returns one of the following results. See [Protocol] for the meaning of these results.

5.1 Neutral or None or PermError

An SMTP server receiving one of these results SHOULD NOT reject the message for this reason alone, but MAY subject the message to heightened scrutiny by other anti-spam measures, and MAY reject the message as a result of this heightened scrutiny.

5.2 Pass

An SMTP server receiving this result SHOULD treat the message as authentic. It may accept or reject the message depending on other policies.

5.3 Fail

An SMTP server receiving this result SHOULD reject the message with a "550 5.7.1 Sender ID xxx - yyy" SMTP error, where "xxx" is replaced with the additional reason returned by the check_host function and "yyy" is replaced with the explanation string returned by the check_host function.

5.4 SoftFail

An SMTP server receiving this result SHOULD NOT reject the message for this reason alone, but MAY subject the message to heightened scrutiny by other anti-spam measures, and MAY reject the message as a result of this heightened scrutiny. A message for which the result is "SoftFail" is less likely to be authentic than a message for which the result is "Neutral".

5.5 TempError

An SMTP server receiving this result MAY reject the message with a "450 4.4.3 Sender ID check is temporarily unavailable" error code. Alternatively, an SMTP server receiving this result MAY accept a message and optionally subject it to heightened scrutiny by other anti-spam measures.

6. Security Considerations

This entire document describes a new mechanism for mitigating spoofed email, which is today a pervasive security problem in the Internet.

Assuming that this mechanism is widely deployed, the following sections describe counter-attacks that could be used to defeat this mechanism.

6.1 DNS Attacks

The new mechanism is entirely dependent on DNS lookups, and is therefore only as secure as DNS. An attacker bent on spoofing messages could attempt to get his messages accepted by sending forged answers to DNS queries.

An MTA could largely defeat such an attack by using a properly paranoid DNS resolver. DNSSEC may ultimately provide a way to completely neutralize this class of attacks.

6.2 TCP Attacks

This mechanism is designed to be used in conjunction with SMTP over TCP. A sufficiently resourceful attacker might be able to send TCP packets with forged from-addresses, and thus execute an entire SMTP session that appears to come from somewhere other than its true origin.

Such an attack requires guessing what TCP sequence numbers an SMTP server will use. It also requires transmitting completely in the blind - the attack will be unable hear any of the server's side of the conversation.

Attacks of this sort can be ameliorated if IP gateways refuse to forward packets when the source address is clearly bogus.

6.3 Forged Sender Attacks

This mechanism chooses a purported responsible address from one of a number of message headers, and then uses that address for validation. A message with a true Resent-From header (for example), but a forged From header will be accepted. Since many MUAs do not display all of the headers of received messages, the message will appear to be forged when displayed.

In order to neutralize this attack, MUAs will need to start displaying at least the header that was verified.

6.4 Address Space Hijacking

This mechanism assumes the integrity of IP address space for determining whether a given client is authorized to send messages from a given PRA. In addition to the TCP attack given in section 6.2, a sufficiently resourceful attacker might be able to alter the IP routing structure to permit two-way communication using a

specified IP address. It would then be possible to execute an SMTP session that appears to come from an authorized address, without the need to guess TCP sequence numbers or transmit in the blind.

Such an attack might occur if the attacker obtained access to a router which participates in external BGP routing. Such a router could advertise a more specific route to a rogue SMTP client, temporarily overriding the legitimate owner of the address.

7. Implementation Guidance

This section describes the actions that certain members of the Internet email ecosystem must take to be compliant with this specification.

7.1 Simple E-mailers

A domain that injects original email into the Internet, using its own name in From headers, need do nothing to be compliant. However, such domains SHOULD publish e-mail policy records in DNS.

7.2 E-Mail Forwarders

A program that forwards received mail to other addresses MUST add an appropriate header that contains an email address that it is authorized to use. Such programs SHOULD use the Resent-From header for this purpose.

Additionally, e-mail forwarders SHOULD publish Sender ID records for their domains, and SHOULD use MTAs for which the Sender ID check yields a "pass" result.

Some of today's forwarders already add an appropriate header (although many of them use Sender rather than Resent-From.)

7.3 Mailing List Servers

A mailing list server MUST add an appropriate header that contains an email address that it is authorized to use. Such programs SHOULD use the Resent-From header for this purpose.

Additionally, mailing list servers SHOULD publish Sender ID records for their domains, and SHOULD use MTAs for which the Sender ID check yields a "pass" result.

Most of today's mailing list software already adds an appropriate header (although most of them use Sender rather than Resent-From).

7.4 Third-Party Mailers

A program that sends mail on behalf of another user MUST add an appropriate header that contains an email address that it is authorized to use. Such programs SHOULD use the Sender header for this purpose.

Additionally, third-part mailers servers SHOULD publish Sender ID records for their domains, and SHOULD use MTAs for which the Sender ID check yields a "pass" result.

Many, but not all, of today's third-party mailers are already compliant.

7.5 MUA Implementers

When displaying a received message, an MUA SHOULD display the purported responsible address as defined by this document whenever that address differs from the RFC 2822 From address. This display SHOULD be in addition to the RFC 2822 From address.

When a received message contains multiple headers that might be used for the purported responsible address determination, an MUA should consider displaying all of them. That is, if a message contains several Resent-From's, a Sender and a From, an MUA should consider displaying all of them.

8. IANA Considerations

This document contains no actions for IANA.

9. Acknowledgements

Variations on the idea of using a DNS record to check the legitimacy of an email address have occurred multiple times. The earliest known work is [Vixie]; others include [RMX], [SPF] and [CallerID].

The current document borrows heavily from each of the above, and incorporates ideas proposed by many members of the MARID working group. The contributions of each of the above are gratefully acknowledged.

10. References

10.1 Normative References

[PRA] J. Lyon, "Purported Responsible Address in E-Mail Messages", draft-ietf-marid-pra-00. Work in progress.

[Protocol] M. Wong and M. Lentczner, "The SPF Record Format and Test Protocol", draft-ietf-marid-protocol-01. Work in progress.

[RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119.

10.2 Informative References

[CallerID] Microsoft Corporation, Caller ID for E-Mail Technical Specification, http://www.microsoft.com/mscorp/twc/privacy/spam_callerid.mspx.

[RMX] H. Danisch, "The RMX DNS RR and method for lightweight SMTP sender authorization", draft-danisch-dns-rr-smtp-04. Work in progress.

[SPF] M. Lentczner and M. Wong, "Sender Policy Framework (SPF): A Convention to Describe Hosts Authorized to Send SMTP Traffic", draft-mengwong-spf-01. Work in progress.

[Submitter] E. Allman and H. Katz, "SMTP Service Extension for Indicating the Responsible Submitter of an E-mail Message", draft-ietf-marid-submitter-03. Work in progress.

[Vixie] Paul Vixie, "Repudiating Mail-From", http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg00658.html

11. Authors' Addresses

Jim Lyon
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA
jimlyon@microsoft.com

Meng Weng Wong
Singapore
mengwong@dumbo.pobox.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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