Network Working Group
Internet-Draft
Expires: February 2004

C. Kularski
Highland School
of Technology
August 2003

SPAM Reduction Through Creative Addressing

draft-kularski-spam-spamreduce-03.txt

Status of this Memo

Abstract

This document describes a procedure that users can follow to significantly cut down on the amount of SPAM that they receive. SPAM/UCE (Unsolicited Commercial Email) has become a problem for most Internet users, there is currently no complete solution to the problem. Once the procedure described in this document the user can expect to see dramatically reduced SPAM. Some user refinement may be required at first, but this procedure is very low maintenance.

Copyright Notice

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 RFC-2119 [i].

A Virtual Address as used in this document is an email address not directly existing on the server, but it specified by a catch-all.

Table of Contents

1. General Description

The key to making this procedure for SPAM elimination work with currently available server software is having two email boxes available. Each of these boxes MUST meet a different set of criteria described later in this document.

This procedure can be used by corporate administrators, Internet Service Providers (ISP), or users who have the resource of their own email server.

1.1 Proper Implementation

Because the procedure described in this document drastically changes the way user receives email the implementation should either be performed at the user’s request, or with significant prior notification.

2. Mailbox 1

This mailbox can be used with automated systems, and just about any other purpose, except for those purposes noted for Mailbox 2 in Section 3. This is the only mailbox that is valid for use with non-human senders.

For this mailbox the server will need to recognize each user as their own sub-domain (ex. Jane Doe uses janedoe.example.net). The mailbox MUST have a sub-domain or FQDN (Fully Qualified Domain Name) associated with it. An MX (Mail Exchanger) record should point the domain to the email server on which the account resides. The mailbox will be a general collection box for receiving all of the email pointed at a catch-all mailbox. The address of the real email box should remain private, unless Section 4 is utilized. If Jane uses jane@janedoe.example.net to login to her email box, she should never release jane@janedoe.example.net to anyone as her email address. Each entity that is to receive an email address from the user should be given a unique address, so that the user has the ability to terminate an email address that has been SPAMed and possibly sold to a mass mailing list. If Jane were communicating with the Internet Engineering Task Force she could communicate her email address as being IETF@janedoe.example.net.

Any email alias for common services, roles or functions, as defined by RFC 2142 [ii], should be defined as aliases and pointed to those users on the overseeing organization's domain (ISP, Corporation, etc). If the user has his/her own server those roles (especially Postmaster) it is highly recommended that the user point those addresses to the white listed box [Section 3].

2.1 The Email Server Software

The Email server MUST have software capable of handling a catch-all system. The catch all mailbox needs to point to a single mailbox on the server. The mailbox may reside on the same domain or a separate domain, depending upon the user's needs and the server capabilities. The email server SHOULD support the "X-RCPT-TO" email header to allow for identifying mail that may be disguised.

Some email servers will often tell the sending server the destination of any type of forwarded address, including for catch-alls. This MUST NOT be allowed to occur on a server where the procedure described in this document is implemented.

2.2 Alternative to Catch-all

A more user involved, but more reliable alternative to the catch-all method for the first mailbox is having each user to specify their own list of acceptable addresses.

In this method only the accepted addresses will be able to receive mail. This can be done through the use of mail server aliases being added for each approved address, or having the catch-all in place and having everything sorted out by a mail server rule that checks a list of approved addresses. Messages received that are not on the approved recipients list should be moved to a queue.

2.3 When SPAM Occurs

After a short amount of time in circulation one or more of the user's virtual addresses will begin to attract SPAM. As soon as SPAM is received the "X-RCPT-TO" or "TO" lines in the header should be checked to verify the address that the mail was destined for. The virtual address should be immediately discontinued from use.

A few options exist for what to do with the virtual address after it is identified as a SPAM recipient. First, the virtual address can be created as an alias and forwarded to a dead-end mailbox that is automatically cleared after a certain amount of time (or is never permanently recorded). The second option is a little less drastic, the virtual address can be created as an alias and pointed to another actual account residing on the user's domain. For example, Jane can get all of her SPAMed virtual addresses pointed to spam@janedoe.example.net where she can later sort the mail manually, or by a conventional SPAM identification program.

3. Mailbox 2

This mailbox can be used for personal communication, public newsgroups, web page contact or a situation where the address will only be used by humans.

For this mailbox the server must support intelligent white listing. Intelligent white listing involves the email box not only receiving email from senders listed on the white list, but also sending an email to those who are not on the white list to give them a chance to verify that they are human by accepting an email at a special address, once that mail is received and the sender is confirmed the sender is automatically added to the white list, and the mail is released from the queue and delivered to the user.

White listing by itself is effective in eliminating SPAM, but is horribly inconvenient, so it MUST be used in conjunction with the catch-all mailbox in Section 2.

If SPAM is found in the white listed mailbox the sender's email address should be removed from the white list and added to the blacklist.

It is preferable to place existing email addresses as the white list protected address once automated systems that must contact the user have been notified of their assigned address on the catch-all system. Doing so will prevent an interruption in email, or the transition period often associated with changing email systems.

4. Combining Both Mailboxes

Maintaining two independent email boxes is not user friendly, nor does it maintain a low amount of network traffic. Maintaining two separate mailboxes is quite resource heavy for both the server and client. The two mailboxes can be combined on most servers that support both catch-all and white list functions.

The proper way to configure both systems as a single mailbox is to set up the catch-all system as specified, and then configure an alias to use white listing. If mail to the white listed alias passes the white list it can be delivered to the user's main mailbox that they keep secret.

5. An Oops Queue

Where possible the email server SHOULD provide access to a queue where rejected mail from the whitelist or mail to an address not specified by the user (if using option in Section 2.2) is stored. One possible way of implementing the queue is to use a web-based interface that connects to a non-user mailbox, such as "queue" or "spam".

The queue should be cleared of mail older than a set time limit such as 30 or 45 days. An alternative to this would be a size based queue. Once the queue reaches a certain size begin deleting old mail on a first-in, first-out method. Consideration SHOULD also be given to a removal method that will remove abnormally large email from the queue without regard for the first-in, first-out method.

6. SPAM Elimination Process

There is a specific process that SHOULD occur for the user to be able to be as SPAM-free as possible. The process uses the procedure from

this document as well as other SPAM-prevention techniques. Each level is dependant upon server capabilities, but as many levels as are available should be utilized.

(1) Verify that the recipient address is valid locally
Recipient address should either directly exist on the server, or be a valid alias that has been user specified, etc. This step requires that the server be used only for incoming mail, and relayed mail is handled by another server.

(2) Verify open-relay status of sending server
If the sending server is listed as an open-relay with an open-relay database the message is most likely SPAM, but you can not be certain, recommendation in this situation is to move to the queue.

(3) Check the message for viruses
If the message contains any viruses it should be dropped, or moved to a quarantine area.

(4) Check mail using weight-based SPAM detection software
Use a SPAM detection software that assigns messages a point value based on keywords, invalid headers, and other information. Use a moderate cut-off weight to prevent valid mail from being flagged as SPAM.

7. Future Considerations

In the future the developers of email server software may want to write the software with the ability to assign each user to their own sub-domain and not have to specify the sub-domain as an independent domain within the sever software configuration.

8. Results of Experiments Performed

Several experiments of the procedure described in this document were performed. In each of the experiments there was no loss of legitimate email, and only about 2% of the mail was identified as SPAM. The experiments were performed with live email accounts and actual users using the mailboxes for a period of 6 months.

The experimental users had an average of about 25 aliases for avoiding SPAM on the catch-all system, and an average of 3.2 addresses on their blacklists to avoid mail going to the whitelist only system.

Security Considerations

There are no security concerns associated with this document, other than those that are already present in current electronic mail protocols.

References

i Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

ii Crocker, D. “Mailbox Names for Common Services, Roles and Functions”, RFC 2142, May 1997

Author's Addresses

Curtis M. Kularski
219 Best St
Bessemer City, NC 28016-9330
United States
Phone: +1 (704) 629-2498
Email: curtis@kularski.net

Full Copyright Statement

Acknowledgement

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