Spam tagging explained

The Janet Mailer Shield service can attempt to identify spam (also known as UBE or Unsolicited Bulk Email) and add a tag with a score indicating the possibility of an e-mail being spam. The higher the score, the more likely it is that it will be spam. The Janet Mailer Shield achieves this by passing the e-mail through a program that uses several criteria to try and discover the nature of the e-mail being sent. A commonly used resource are the global DNSBLs (Domain Name Server Blackhole Lists), which are DNS (Domain Name Server) lists set up by several groups to indicate which servers on the Internet are currently sending spam ? due either to deliberate spamming by the owners or to poor configuration or security flaws being exploited by a third party.

More information about DNSBLs is available at: http://www.dnsbl.info/dnsbl-list.php [1]

The program also performs its own filtering in addition to using several DNSBLs, and is updated as required to try to minimise false positives (e-mail incorrectly seen as spam) whilst maximising correctly identified spam.

Description of tags

If an organisation requests to have the tagging option enabled, the e-mail passed to it will have several extra headers added.

X-JMS-Connection-From: 10.1.2.3

> This header should always be present, and the IP address (the series of digits at the end) will be the IP address that passed the e-mail to the Janet Mailer Shield. As the Janet service relays your mail to you (remote mail servers no longer connect to you directly), e-mails cannot be blocked based purely on the IP address of the server connecting to you. However, some software may allow you to use the information given in this header line to enforce your own local blocklist for servers you find particularly annoying.

X-Spam-Bar: ########## Score 10 (3,10) angel1:DATA/E42F1.f0*

This header should always be present. The end-users should use it to direct probable spam into a separate spam folder, which they should check at least once a day to see if there are any legitimate e-mails (false positives) before deleting the rest.

If the Janet Mailer Shield system believes something is spam, it will score 8 or more; if not, the score will be 7 or less. The ?X-Spam-Bar:? will contain a row of #?s equal to the score, so even fairly crude mail filters should be able to identify it. End-users should set up a rule equivalent to:

?If header "X-Spam-Bar:" contains "########" save mail in spam-folder?

This would also catch spam scoring more than 8. The details of how you achieve this will depend on what software you run locally. The IT support at organisations using the Janet Mailer Shield tagging will need to instruct end-users on how best to configure their local rules.

Users should not filter on any scores below 8. Tuning revolves around 8+ equalling spam, so a score of 7 should not be filtered. Filtering on scores greater than 8 is fine: however, the higher values will result in less spam being detected.

X-Spam-Bar: WinExe (Probably a virus)

This second ?X-JMS:? will only turn up if the Janet Mailer Shield identifies a Windows® binary executable in the e-mail. The ?WinExe? tag will also be in the first ?X-Spam-Bar:? but some simple filters are unable to find fixed text in a variable header string, so this second header will be added with the important tag at the start. Please note that the Janet Mailer Shield is not checking for viruses, only Windows® executables. Organisations should still run their own virus checker on the local mail hub and on desktops.

WinExe e-mails will also be given a spam score of at least 10.

Interacting with the filters

All of the above are designed to make it easier for end-users to handle spam. With one exception (mentioned later), the Janet Mailer Shield will not block e-mails to an organisation. It is recommended that the organisation mail hub passes the e-mail on to the end-users because false positives cannot be totally eliminated, and any other course of action may result in important e-mail not being delivered. Even if someone receives hundreds of spam each day, if the spam is suitably filtered into a separate folder then most of the problem goes away and makes the handling of spam trivial. This approach also means the end-users can provide useful feedback to the filtering system.

For the purposes of interacting with the filter, spam is any unsolicited bulk or commercial e-mail. E-mails sent by sponsors or owners of lists or groups you have subscribed to are not spam. If you report them as being spam then you will probably end up with all e-mails from the list or group being tagged.

There are three addresses by which end-users may interact with the tagging system.

false@jms.ja.net [2]

> Any false positives (non-spam messages scoring 8 or more) should be sent here with full headers. If the content of the message is personal or in any way confidential then it should not be forwarded with the headers. Only one report should be sent in each e-mail. Please do not send batches in digests or multiple attachments.
>
> A full set of headers will contain header lines starting with ? Received: ?. If you want to see an example of what e-mail headers from your organisation look like, send an e-mail from your normal address to echo@jms.ja.net [3]. The returned e-mail will include the full set of headers.
>
> False positives will be treated as a priority, and you will generally receive a response to these ? although as more organisations join the Janet Mailer Shield and use the tagging, there may come a point at which you will only get a response if more information is required or if you need to do something.

spam@jms.ja.net [4]

> Any spam which scores 7 or less should be sent here with full headers and body text. Only one report should be sent in each e-mail. You will not get a reply to these (except on rare occasions where some advice or other information needs to be given). They will simply be scanned through to look for ways of modifying the filters.

help@jms.ja.net [5]

> For most e-mail related issues and problems, end-users should contact their local postmaster, but if you need to interact with the administrators of the spam tagging system and it is not simply to report spam or false positives, you may use this address. Organisation administrators who are requesting help or advice on behalf of their end-users may also use this address for quick queries concerning the tagging, but should generally use the normal Janet Mailer Shield contacts as given elsewhere.

Testing your filters

It is useful to be able to test your filters when you set them up. To make this easier you can

force the system to add a user-defined tag. To do this send e-mail to the echo server:

To: echo@jms.ja.net [3]

Subject: set-x-spam: ##########

The echo server will echo back a response, but it will use the bar you have given, including any additional text. You may add text up to a limit of about 80 characters. The bar must start with at least 4 ?#?s. The following request line would also be acceptable:

To: echo@jms.ja.net [3]

Subject: set-x-spam: #### (Testing 4 #?s)

Additional information

Bad and good source IP addresses

Over a period of time, the Janet Mailer Shield may find that nothing but spam or some other taggable e-mail comes from certain IP addresses. If this continues and not even one non-taggable message is seen, the IP address sending the e-mail will get marked as bad and every subsequent message for that IP address will be tagged with the same values without examining it in detail. Every so often, the server will be rechecked, and if that check is satisfactory the cache will be cleared, but if the e-mail is still then bad further mails will continue to use the cached values up until the next cache-check period.

End-users may notice this when someone they regularly communicate with gets a virus on their PC and sends a lot of Windows® executable viruses to other Janet Mailer Shield customers. The gateway they come from will probably be marked as bad, and if they then send a normal e-mail shortly afterwards then that may well be marked as a Windows® executable too. Once they clean their PC then things should return to normal, although that may take several hours after the last bad e-mail was seen.

The converse is also true. If nothing but good e-mail is seen from a gateway then the Janet Mailer Shield will also note that. If extremely little spam comes from a particular gateway which is cached as being good, the spam will be passed through using the cached result, but if even one bad e-mail is seen during a check then the cache will be cleared.

For the most part, only very busy servers or spam hosts tend to send enough to end up marked as good or bad, and the markings are mostly long-term.

Virus checking

Whilst the Janet Mailer Shield will try to spot the most obvious binary Windows® executables, it does not perform any virus checking. All organisations should run a competent virus checker on their desktops: it is also recommended to run a virus checker on their mail hubs.

Dropping binary Windows® executables at the gateway

As mentioned above, the Janet Mailer Shield normally expects e-mail to be passed on to the end-users for them to take the actions they feel are required. The one possible exception is Windows® executables, which are usually viruses. During particularly heavy virus outbreaks

the numbers of viruses hitting an organisation can cause problems for local mail hubs. Most organisations can cope, albeit with delays, but some servers cannot. If your server is unable to handle the kind of load that a new virus outbreak causes, and is also unable to use the tag headers to drop the e-mail itself, then the Janet Mailer Shield operators may agree simply to drop all identified messages with binary Windows® executables which are destined for your organisation. Please note that they will not be quarantined and no notifications will be sent to the sender or the recipients. They will simply disappear.

Any users legitimately sending self-extracting archives or other executables to colleagues will need to encode them in some fashion. Using pkzip/WinZip may work, although attempts are made to try and identify those too. At present, if the first file in the archive is text then it should get through. You may find it more reliable to install some third party software such as gzip in order to exchange binary Windows® executables.

If your organisation already has a policy of barring executables then there should be no problem with the above, However it is important to note that the Janet Mailer Shield will not generally unpack archives or encoded messages, so you need to have software at your hub which can fully enforce your organisation?s policy.

DNS Blackhole Lists

DNS Blackhole Lists (DNSBLs) are lists of IP addresses from which you might choose not to accept mail connections. Each list consists of a number of entries in the DNS.

DNSBLs each have their own criteria and operating practices, but almost all have the same way of interrogating the list. The Spamhaus Blackhole List (SBL) maintained by Spamhaus.org [6] illustrates the process. In this example it is applied to ns0.ja.net, a system very unlikely ever to be present in the SBL:

1. The list uses IP addresses rather than domain names, found using the usual DNS tools (e.g. host). For example, host ns0.ja.net gives 128.86.1.20.
2. Reverse the order of the four address components: so, 128.86.1.20 becomes 20.1.86.128.
3. Append the zone name for the list concerned. For the SBL the zone is sbl.spamhaus.org, giving in this example 20.1.86.128.sbl.spamhaus.org.
4. Look up this contrived domain name in the DNS:
   host 20.1.86.128.sbl.spamhaus.org
   This gives a response indicating that there is no entry for this address in this list:
   Host 20.1.86.128.sbl.spamhaus.org not found: 3(NXDOMAIN).

For many of the DNSBLs the artificial IP address 127.0.0.2 is added to the list for test purposes. Instead of failing as above, looking up this address will give a positive result.

In normal use a mail system will perform this test for the IP address of each incoming connection. For listed addresses it will take some action such as refusing the connection, or accepting it and marking the transferred message in some way. For unlisted IP addresses, the mail system will accept the connection and transfer the message in the usual way.

Links

[1] http://www.dnsbl.info/dnsbl-list.php
[2] mailto:false@jms.ja.net
[3] mailto:echo@jms.ja.net
[4] mailto:spam@jms.ja.net
[5] mailto:help@jms.ja.net
[6] http://www.spamhaus.org/