



CONTENT

Introduction.....1

Definition of Spam.....2

Definition of E-mail Relationship.....2

Definition of Relationship Enforcement.....2

Challenge Response Systems.....2

Network Box Relationships and Management.....4

The Network Box Relationship System.....4

Network Box Spam Score Adjustments.....5

Conclusion.....8

AUGUST 2009

No part of this publication including text, examples, or illustrations may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical or otherwise, for any purpose, without prior written permission of Network Box Corporation Limited.

Network Box Corporation Limited,
 16th Floor, Metro Loft,
 38 Kwai Hei Street, Kwai Chung,
 Kowloon, Hong Kong
 Telephone: +852 2736-2083
 Fax: +852 2736-2778
www.network-box.com

Introduction

Today, Spam email is ubiquitous and aggravatingly persistent. Since Network Box started work on anti-spam and anti-virus systems, anti-spam signatures have shot passed the 3 million mark. Anti-virus has jumped from 30,000 known viruses to more than a million. In 2009, Network Box is adding support for over 50,000 new virus variants each month.

While viruses can be blocked almost 100% of the time, traditional spam filters are able to block up to only 95 – 98%. Unfortunately this still allows a goodly number of unwanted messages through considering the volume of spam in circulation nowadays. While spam filtering is effective for the more obvious types of Spam, spammers have also become more sophisticated and are outsmarting the filters and their creators, who are themselves evolving in an effort to combat the spammers.

At the same time, resources dedicated to anti-spam are increasingly dear: hundreds of thousands of signatures, real-time DNS checks, multiple engines, thousands of regexes, OCR, etc.. The end result is that benchmarking shows anti-spam takes 7 – 10 times the resources of anti-virus. The biggest culprit is the signature engines, with RBL checks following a close second.

But, in a never-ending battle of wits, intelligence and technological advancement, how does one effectively and efficiently eradicate, or minimise, Spam from one's network / system?

The most effective and efficient method is to use the email relationships method. This is not content filtering; it is, instead, a technological advancement that can help eliminate Spam.

Definition of Spam

To effectively eliminate Spam, however, requires an understanding of it:

The solution, therefore, lies in E-mail Relationship and Relationship Enforcement.

Unsolicited e-mail, often of a commercial nature, sent indiscriminately to multiple mailing lists, individuals, or newsgroups; junk e-mail.

Source: www.thefreedictionary.com

Spammers most often harvest email addresses through trawling web sites, monitoring newsgroups, Trojan'd machines, buying lists from each other, etc.

As already indicated, blocking and filtering such emails is not enough. A closer analytical look reveals the most common factor in Spam is the unrecognised sender — unless we have opted into a mailing list, we generally block, blacklist or delete email from people we don't recognise. Considering this type of email behaviour, the solution, therefore, lies in E-mail Relationship and Relationship Enforcement.

Definition of E-Mail Relationship

An e-mail relationship is an established connection between a sender and a recipient, based on pre-programmed or learned behaviour.

Relationships are presented as a tuple (sender email address, recipient email address and some attribute to validate the sender).

Definition of Relationship Enforcement

Enforcement, on the other hand, is actively blocking, delaying or forcing a manual response.

For example, in SMTP, we can drop / reject email, and we can temporarily refuse to accept email. We can also quarantine then challenge the sender to manually confirm their humanity (challenge-response).

Challenge-Response Systems

Challenge-response is an anti-spam technique used to ensure that the person sending you an email message is a human being.

Currently, the number of systems that offer or perform Challenge / Response (C/R) to an effective standard are rare. Many systems claim high efficiency rates, but actual tests show a high variance. On the other hand, more security

vendors are beginning to look into email behaviour to combat increasing Spam sophistication.

The system is a basic and limited form of relationship enforcement. It works by monitoring the sending email address of incoming messages. If the sender is previously unknown, a challenge-response system will quarantine the incoming message and challenge the sender to authenticate themselves. If the sender authenticates successfully, the email is released from quarantine and the sender's email address is added to a whitelist to prevent future challenges.

Normally, the challenge is either a straightforward email reply or a web-based hyperlink to be clicked, but, it can take various forms. Both the issuing of the challenge and the handling of the response are automated by the system.

Historically, users have complained about challenge-response systems for a number of reasons. These reasons stem from many problems, but as at writing, the top three problems are:

- *Susceptibility to Forging* . In general, developers of C/R systems agree that this is the main weakness still prevalent despite today's technology: The sender address of an email is simple to forge. This leads to problems with innocent users (who have had their email address spoofed) receiving challenges for spam they did not send.

This also includes the problem of backscatter and polluting the Internet. As the sender is forged, he receives a challenge for a message he never sent. For example, if Recipient A receives a message sent by a forger who is imitating Sender B. Recipient A's system then sends a C/R email to Sender B asking him to verify his identity and that he sent the email. However, Sender B never sent that email.

- *Loops*. If two users with challenge-response systems talk to each other, the result can be a deadlock, where both systems challenge each other.
- *Response Rate*. For challenge-response systems, the response rate may be as low as 40%, even for real, human senders. According to these studies, people are scared to respond, meaning that once you turn on such a system, you start losing mail due to people wanting to contact you but are educated against, or too scared of or too lazy to go through the hurdles of the response process.

It is important not to blind cc yourself, and also not to whitelist yourself.

It is important not to blind cc yourself, and also not to whitelist yourself. It is relatively common for spammers to forge a message with the same 'from' address as the 'to' address.

Unfortunately, despite the obvious benefits to users, C/R systems are still getting a bad reputation, mostly pointing to the increased stress load on a

The Network Box relationship enforcement system has the ability to operate without challenge-response

network operating a C/R system, on the server side. Other shortfalls highlight the client side: potentially, senders and recipients can get stuck in a loop of challenge / authenticate with neither party receiving the email message in question.

Fortunately, there are workarounds and many systems already have this built in. Conforming to RFC3834 also helps ensure the system does not fall into an endless loop of challenge / authenticate.

So, for all its teething problems, the relationship system is an ideal foundation for a smart challenge-response system based on a full knowledge database.

Network Box Relationships and Management

Using the above information, the Network Box relationship enforcement system is able to tackle the listed (and non-listed) issues differently. By having a 'learning' period before 'enforcement' is enabled, we have the ability to be more selective about whom we challenge.

The Network Box Relationship system provides a database that tracks who is sending what to whom. It tracks the history (such as Spam and virus rates) and establishes a trust relationship between the tuple of sender, sender attribute (e.g., IP address) and recipient.

The Network Box Relationship System

At the same time, the Network Box relationship enforcement system has the ability to operate without challenge-response, but can benefit from it for the last 1 – 2% of cases. The system works thus:

1. A central relationship database is maintained. This database stores relationships for all email accounts managed by Network Box.
2. The relationship is defined as sender + recipient + type + score.
 - *Sender* is the sender email address + attributes. Attributes includes IP address, network address, country, reverse-IP domain and others.
 - *Recipient* is the recipient email address.
 - *Type* describes the relationship detail and identifies how the relationship was established.
 - *Score* indicates the trust and strength of the relationship.
3. The sum of all scores for all detail records make up the total relationship score for that sender and recipient.
4. The system always learns:
 - If you send an email outbound from A to B, it will create (or strengthen) the score of the reverse relationship (sender B, recipient A, detail type 'outbound email').

The system uses the relationship database to automatically adjust anti-spam scores based on relationship strength.

- If you go through a challenge successfully, that relationship will be strengthened.
 - If you send your colleague an email from the same IP or country, that relationship will be strengthened, and so on.
6. When the system is in one of its 'enforcement' modes, it has the ability to enforce relationships. It makes this decision based on its configuration, the score and detail records. The enforcement options include:
- Challenge-response
 - Spam quarantine
 - Policy quarantine
 - Temporary deferment
7. The relationship database becomes a valuable database of historical information. For example, if an EXE file comes in, the system can review your previous relationship with the sender (email address, domain, country or IP address) and decide the likelihood of a virus.
8. The relationship database can be queried.

Implementing the above into the system and ensuring it works smoothly requires spam score calculations and adjustments. To effectively do so and block Spam while permitting genuine emails through, the system uses the relationship database to automatically adjust anti-spam scores based on relationship strength.

Network Box Spam Score Adjustments

Relationship Spam Score Adjustments are usually enabled for inbound SMTP email not coming via a backup MX server (i.e., direct from the sender). The system will also not run if the email has already been specifically whitelisted or blacklisted.

This system can be used to be:

- More aggressive towards known Spam sources; and
- Kinder to known good sources of non-Spam mail.

It is also exceptionally effective with messages in the 'border line' zone (typically scoring between 5.0 and 9.0 before adjustment), and can 'tip the scale' to mark the messages as Spam / not Spam. The system accounts for the previous relationship history of that sender and recipient pair. It also has a confidence mechanism to attempt to authenticate the sender, even if the SMTP protocol itself does not have sender authentication capability.

For each message, the system first determines (using envelope analysis and geographic IP location) some basic information on the sender, including:

- E-Mail address

- IP address
- IP / 16 address block
- Country hosting IP

The system then queries the relationship database for matches of the following tuples, and produces an average total across matching relationship records:

- Sender email address, sender country, recipient (75% confidence for a specific recipient address match, 50% for a domain match).
- Sender email address, sender/16 network block, recipient (100% confidence for a specific recipient address match, 75% confidence for a domain match).

These confidence figures are the weightings given to the match. For example, if you have relationship records for previous emails from a specific sender from a specific country to a specific recipient, then the system would treat this as 75% confidence the email is from the same sender.

Scores stored range from -100 to +100 (with -100 being 100% malicious and +100% non-malicious). The database stores these for each trust, Spam, malware and policy value.

In the final stages of anti-spam — scanning — the Email Relationship Spam Score Adjustment System:

1. Takes the average relationship Spam score and multiplies by the confidence (expressed as a percentage) to determine the adjusted Spam weight (expressed as a percentage). For example, if relationship records gave +100 (i.e., 100% not spam) with a 50% confidence, the adjusted Spam rate would be 25% (i.e., on the line 0=ham to 100=Spam, we are half way below the mid-way point of 50=unknown).
2. Maps the adjusted Spam weight onto a configurable sliding Spam score scale to determine and adjustment Spam score. For example, if the sliding Spam scale was -7 to +7, then an email whose relationship history indicated an adjusted spam weight of 25% would result in an adjusted Spam score of -3.5 (i.e., 25% of the way between -7 to +7).
3. Raises a Spam test result to adjust the total Spam score up or down, depending on the Spam weight. This Spam test is known as NB_RELATIONSHIP_SSA.

The system also has the ability to take into account the age and volume (messages processed) of the relationship to further adjust the weightings.

The overall approach is to look at the history and come up with weighted scores (for trust, Spam, malware and policy) based on previous averages weighted by how confident you are of the sender. You can then use that to adjust your Spam scores.

The system is extremely configurable and tunable. It can be configured to only adjust scores up or down by configuration of the adjustment range, and can be individually configured and tuned on each Network Box.

That said, the Network Box implementation of challenge-response is built on the Relationship system foundation. You can selectively enable either per-recipient or per-domain. It can also be configured to only challenge messages not already identified as Spam. It is most suited to those mailboxes experiencing a very high volume of Spam.

Key points of the system are:

- *An initial learning period.* During this time, relationships (e.g., who you are emailing) are learned before enforcement is enabled. This means that the majority of people you communicate with are already in the system, and don't need to be challenged when the system is enabled.
- *Full conformity to RFC3834.* The system avoids email loops and problems with mailing lists and bulk emails.
- *Full conformity to industry standard 'best practices' for challenge-response systems.* The system minimizes the impact on innocent third parties.
- *The system runs after conventional anti-spam and anti-virus have run.* The system runs in addition to the existing protection. It will not challenge the 95 – 98% of Spam and almost 100% of viruses that are blocked by our conventional systems. It is purely targeting the small portion of non-ham that would otherwise get through. However, anti-spam aggressiveness can be tuned down once this system is enabled.
- *Per-recipient basis, but supports 'introductions'.* An outbound email to internal as well as external recipients is treated as an introduction by the database: it records the establishment of a relationship between the external and other internal recipients. The system is also optionally configurable to allow relationships to be shared on a per-domain or per-box basis.
- *Full integration with Mail Portal.* The web-based UI allows recent challenges to be seen and manually released. The periodic email report shows challenges made and responses received, and allows for click-and-release handling as necessary.

You can selectively enable either per-recipient or per-domain. It is most suited to those mailboxes experiencing a high volume of spam.

The primary issue still seen with the system is the large number of automatic responders not conforming to the RFC3834 standard. In such cases, a manual release may still be required. But, the integration to the Network Box Mail Portal makes this relatively simple.

Conclusion

So, while no anti-spam system can ever be 100% effective, the Relationships method is taking us in that direction. It will also lay the foundation for inter-box relationship sharing.

Since implementing the system, we have found the relationship enforcement system to be close to 100% effective against current Spam techniques and resilient to changes in approach by spammers.

For the Network Box, relationship enforcement adds minimal overhead (and reduces it in some cases). For spammers, they will have to redesign their entire database system.

To see the system in action, or for further information, please get in touch with your local NOC, or contact Network Box Corporation Limited at nbhq@network-box.com.

Since implementing the system, we have found the relationship enforcement system to be close to 100% effective.