



Best Practices in Anti-SPAM

Advisory document for the ETIS community

Editor: Folkert Visser, KPN
Date: September 2007
Status: Draft



Participants in the Anti-SPAM Co-Operation Project are:

- Koninklijke KPN N.V., The Netherlands
- Belgacom, Belgium
- Telenor, Norway
- TeliaSonera, Finland
- McAfee, USA
- Ironport, USA
- TNO, The Netherlands

Anti-SPAM Co-Operation Project

Project leader: Richard Kerkdijk, TNO

Project owner: Fred Werner, ETIS

© 2007 Participants in Anti-SPAM Co-Operation Project

Disclaimer

This document contains material which is the copyright of certain ETIS members and may not be reproduced or copied without permission.

The commercial use of any information contained in this document may require a license from the proprietor of that information.



Contents

CONTENTS	3
1 INTRODUCTION	4
1.1 BACKGROUND	4
1.2 PURPOSE OF THIS DOCUMENT	4
1.3 READER'S GUIDE.....	4
2 ESSENTIAL BEST PRACTICES	6
3 SUPPLEMENTARY MEASURES	12

1 Introduction

1.1 Background

The abuse of electronic messaging systems to send unsolicited bulk messages, commonly referred to as spam, poses a severe problem for the users of e-mail as well as for carriers and Internet Service Providers (ISPs). Some sources claim that up to 80% of all current e-mail traffic consists of spam. Moreover, in the past few years spam has evolved from simple annoyance to large scale criminal practice. Consequences include annoyance, loss of productivity, over-provisioning of e-mail storage and, increasingly, fraud at the cost of carriers, ISPs and end-users.

Early 2007, ETIS joined forces with four of its member ISPs (Belgacom, KPN, Telenor and Teliasonera), two renowned vendors (McAfee and Ironport) and Dutch research and consulting firm TNO to look into the possibilities of establishing *carrier grade* anti-spam. Specifically, these partners initiated a pilot project to explore the potential of active anti-spam cooperation. The project was concluded with great success in the summer of 2007.

As a supplementary result from the project, the participating partners decided to produce a specification of best practices in anti-spam as an advisory document for the ETIS community. Here, it was felt that the momentum of the project offered an attractive opportunity to fulfil a widespread demand.

1.2 Purpose of this document

The purpose of this advisory document is twofold:

1. Specify a baseline set of best practice measures in anti-spam that is both effective and feasible within the business context of most ETIS member ISPs
2. Specify a supplementary set of measures that may be adopted by ETIS member ISPs which already have an effective anti-spam baseline in place, but wish to develop their maturity in this field in further.

Note that the project team is aware of the existence of other best practice specifications in the field of anti-spam, for instance provided by MAAWG (the Messaging Anti-Abuse Working Group). In most instances, however, these specifications are rather complex and pose stringent demands on ISPs adopting them. With this advisory document, the project team set out to offer ETIS member ISPs a low threshold yet sufficiently effective alternative.

1.3 Reader's guide

In accordance with the aforementioned objectives, the document consists of three specific elements:

1. **Essential best practices.** This is the baseline referred to above. It consists of ten anti-spam measures each ETIS member ISP should adopt. Chapter 2 provides an elaborate description of each.
2. **Supplementary measures.** This is a set of twelve anti-spam measures all ETIS member ISP should consider, actual adoption depending on their specific configuration as well as current maturity of anti-spam strategy. Chapter 3 provides a brief description of each.



Please note that the contents of this specification are by no means intended to remain static. New insights or specific spam developments will most likely trigger updates in the course of time. Readers are therefore recommended to periodically check for new versions on www.etis.org.

2 Essential best practices

E1: Outbound port 25 block	Impact: Medium
Description Block all direct outbound port 25 traffic from (residential) customers towards the Internet. Outbound e-mail traffic should only be allowed via the ISP mailservers.	
Motivation Most malware uses it's own SMTP-engine that directly sends spam or other abuse towards the Internet. By blocking direct port 25 trafic most malware will be rendered ineffective without reduction of service for customers. The effect will be most noticeable at the Abusedesk where the number of complaints from the Internet about customers will reduce significantly.	
Implementation guidance The technical implementation depends on the technical architecture from the ISP. Most ISPs will have a more or less centralised network. This probably is the best place for implementing the outbound port 25 block. Via firewalls or other filters access towards the ISP e-mail-infrastructure can be assured. Experience has shown that an unannounced implementation of an outbound port 25 block might cause a temporary increase in customers. Even if this type of traffic was banned in the general terms of usage. To decrease the peak a phased (geographically or technically) implementation should be considered. Also a press-release or general message for the customers should be released. Customers not willing to use the ISP e-mail-infrastructure can be advised to make use of the guidelines in RFC 2476.	

E2: Acceptable Use Policy (AUP)	Impact: Medium
Description An ISP should make sure that their acceptable use policy (AUP) for end-users makes it possible for the ISP to react correctly to spam incidents. The AUP for corporate customers/business customers with their own network should require the customer to handle spam complaints as fast as possible - preferably with a set time limit - and take appropriate action to stop the problem.	
Motivation Sending or relaying spam should violate the AUP, and the ISP should make sure the AUP allows it to filter, limit or block one or more services until the problem has been solved.	

E3: Multilevel Abuse Handling	Impact: Medium
<p>Description ISPs should implement multiple increased levels of sanctions in case of ongoing abuse-complaints against customers.</p> <p>The first level should be to inform the customer, next a full SMTP-block should be considered followed by a quarantine-like isolation of the customer. The next level can be a temporary shutdown of the access and the ultimate sanction is a permanent termination of service.</p>	
<p>Motivation Most customers are unaware of the abuse they are causing. By progressively increasing the level of sanctions customers will be given enough time and incentive to take appropriate measures.</p>	
<p>Implementation guidance The implementation of the best-practices is very dependant on the organisation and architecture of the ISP. The only useful guide that can be giving is that the approach should be technical and organisational.</p>	

E4: Use Fully Qualified Domain Names (FQDN) in EHLO/HELO	Impact: Medium
<p>Description The ISP's outgoing mailservers should be configured to use the correct information in EHLO or HELO commands, to avoid being blocked or delayed by anti-spam solutions.</p>	
<p>Motivation Rejecting or delaying emails based on invalid or incorrect EHLO/HELO commands will stop spam from certain types of spam software, and takes very little resources. If an ISP does not use the correct information in EHLO/HELO, mail from these servers might be dropped or blocked by other providers.</p>	

E5: Educate your customers	Impact: Medium
<p>Description An ISP should educate its customers, by making information about spam and how to handle spam easily available for the customers.</p>	
<p>Motivation Better educated users would result in fewer customers replying to spam-mail and fewer addresses exposed to spammers. It may also reduce the risk of the customers' computers being infected and used as spambots.</p>	

E6: Sender verification	Impact: Low
<p>Description</p> <p>The domain portion, of the envelope sender address should be DNS verified. That means that the envelope sender domain must resolve, and an A or MX record must exist for that specific domain.</p> <p>Therefore accepting MTAs should never accept any mail, that is being delivered with an envelope sender whose domain is not resolving.</p> <p>This verification should take place during the SMTP conversation.</p> <p>If the DNS request returns "SERVFAIL", the MTA should return a temporary (4xx) SMTP code. In case of "NXDOMAIN", a permanent (5xx) SMTP code should be sent.</p>	
<p>Motivation</p> <p>Spam and unwanted mail, is frequently sent by senders whose domains or IP addresses cannot be resolved by DNS. DNS and/or sender verification means that you can get reliable information about senders and process mail accordingly.</p> <p>Bounce messages are sent to the MAIL FROM address, so in case the incoming message is not deliverable, the ISP is able to send back a bounce message, to an envelope sender domain that exists (or at least has MX/A record).</p>	
<p>Implementation guidance</p> <p>Most MTA software supports the verification of the envelope sender (domain). Check the manual of your specific MTA for more information.</p>	

E7: Set up an active abuse department	Impact: High
<p>Description</p> <p>With an active abuse department your company will have more striking power in the war against spam, this will however require quite some effort from the people responsible to setup the department.</p> <p>The definition of an active abuse department is: "An abuse desk that actively processes the data they receive and continually searches for more data to process, in order to be able to find and react upon spammers in their own network."</p> <p>Each company needs to be able to track which customers sent which e-mails, spammers cannot be tracked without this possibility. One also need to take into consideration local legislation regarding storage of customer data.</p> <p>When buying or creating a system for handling incidents it's important to place emphasis on:</p> <ul style="list-style-type: none">• Automation It is very difficult to handle the large amounts of incident reports that an abuse department usually receives manually. Therefore there should be a high focus on automation as this will reduce response time dramatic. However this must not reduce the quality of the work.• Customer dialog This is an important part of solving incidents for an abuse department,	

with a poor customer dialog you will probably end up with many customers failing to disinfect their systems.

Consider to use some dedicated persons from customer service to handle this part. They are better to talk to customers than technicians are.

- Response types

Consider to apply some access lists with some restricted access to the customers internet connection (this can also be a walled garden) when the first reports of an incident is received. This will drastic reduce the numbers of incidents sent to the abuse department and the number of spam sent from your customers.

Consider to join one or more anti abuse forums, in these forums you will be able to share ideas and problems with colleagues facing the same challenges as your abuse department. The methods used by spammers will continue to develop and it is a continuous battle to keep track of all of them alone.

Some of the anti abuse forums are:

- ETIS (<http://www.etis.org>)
- MAAWG (<http://www.maawg.org>)
- FIRST (<http://www.first.org>)

According to RFC2142 there are a number of mailboxes that are required to exist in an network environment. Among these are abuse, noc, security, postmaster and hostmaster.

These mailboxes must be monitored, you will receive reports of spamming (or other inappropriate behaviour) to one or more of these mailboxes.

Motivation

Setting up an active abuse department will with a high probability get your company's e-mail servers less blacklisted in different RBLs. This will just get more important as the use of IP reputation keeps to grow.

Experience shows that a lot of customers spam each other, so your customers will also have a reduction in the amount of spam received.

Helping customers to disinfect their computers will also reduce the number of bots in your network. These bots can potentially be used for other types of attacks as well. Which will be of great risk for both your company and others connected to the internet.

All of these actions in combination will highly increase the quality of service delivered to your customers.

E8: Customer SPAM and virus filter	Impact: Medium
<p>Description Protect your users by providing an antivirus and antispam solution on the endpoint.</p>	
<p>Motivation In order to avoid false-positives and due to issues of privacy, elimination of spam and malware within the mail infrastructure can only remove a lesser percentage of malicious email.</p> <p>Documents, of which content is split over more than one email, can only be investigated at the point of assembly. If such document contain malware an engine on the desktop is required to detect the exploit.</p> <p>Protection of the user also means that less infection and exploitation occurs and therefore less spambots on the DSL network.</p>	
<p>Implementation guidance Form a strategic partnership with a vendor of anti-virus and anti-spam consumer services in order to provide the appropriate software to customers.</p> <p>This service does not need to be provided free-of-charge and may indeed be used as an additional revenue stream.</p>	

E9: Hybrid Anti-SPAM solution	Impact: High
<p>Description Employ hybrid technology for SPAM filtering</p>	
<p>Motivation No single anti-spam technology is enough to combat the threat of spam. IP-reputation on its own only removes connections from a subset of spam sources. It is not possible to block all IP-addresses sending spam and it might block legitimate email from the same IP addresses.</p> <p>Content-analysis can take care of detecting a lot of spam, without invading privacy. However as content-analysis only happens within the DATA phase, this require more resources.</p> <p>Techniques such as SPF and DKIM on their own are very ineffective.</p>	
<p>Implementation guidance The only solution is to use a collection of tools. At the very least this should include:</p> <ul style="list-style-type: none"> - Highly dynamic IP-based sender reputation - Content analysis of an email during the DATA phase <p>In addition the use of SPF and DKIM validation at the gateway point.</p> <p>A strategic partnership with a provider of such technology is recommended as the threat landscape is forever changing.</p>	

E10: Use feedback loops between ETIS partners	Impact: Medium
<p>Description</p> <p>Monitor the IP range(s) of any partner that wants to cooperate, for fighting mail abuse. That means, that for the IP range(s) of each partner, a daily summary mail should be generated, regarding the mail abuse that was seen. This report must contain the following information, about each offending IP address individually:</p> <ul style="list-style-type: none">• How much spam was detected• How much messages in total were monitored• Timestamp, from address, subject of the last monitored spam message <p>Moreover, each partner should also commit itself, to actively process the feedback reports, that it receives from the other partners.</p>	
<p>Motivation</p> <p>The pilot project that was conducted by an ETIS working group during the first half of 2007, clearly showed the great value of a feedback loop between partners. Not only does the participating party receive feedback from its own IP space, to be able to pinpoint mail abuse sources; but also incoming mail abuse will visibly reduce, by informing parties about their IP ranges, so that the problem can be eliminated at the source.</p>	
<p>Implementation guidance</p> <p>The required information that should be contained within the feedback reports, was carefully selected, so that it should be available in general mail log files. The main idea, is that the party that wants to join up, creates a script that parses the mail log files on a daily base, and sends a report about the monitored IP ranges to the other partners. Depending on the software being used, the existing partners may be able to share knowledge and information, for setting up the feedback loop with the concerned mail server software.</p>	

3 Supplementary measures

Authenticated SMTP

Description

Offer SMTP authentication for residential customers (RFC2554) and split mail streams (auth/non-auth)

Motivation

Allow your customers to authenticate to your mail relays. Proof of identity using SMTP-auth allows your business to be more trustworthy with the connection and be sure that it is a legitimate message. Spyware infected PC's are not able (yet) to learn the SMTP-auth details for users yet.

Providing this rule helps an ISP differentiate trustworthy & not so trustworthy SMTP connections and even consider splitting mail streams so that authenticated mail senders earn a better reputation (e.g. Spamcop & Spamhaus) than unauthenticated mail senders.

E-mail submission

Description

Provide email submission services on port 587 (RFC2476)

Motivation

Educate and encourage your users/customers to submit messages using user authentication techniques provided. By users authenticating, ISP mail relay systems can trust the senders as malware infected PC's do not have the ability (yet) to use SMTP authentication techniques. This allows ISP's to reduce security checks & increase delivery speed of mail for users, otherwise non-authenticated senders face stricter security & slower delivery speeds to ensure messages sent are not spam or viruses.

Inbound port 25 block

Description

Block incoming port 25 to your network

Motivation

By routing all inbound e-mail-traffic via the ISP e-mail platform the general e-mail-filters will be able to filter all e-mails from spam. The use of multiple DNS MX records can solve inbound e-mail routing issue's that originally would directly be sent towards customer mailservers.

ISP SPAM filter

Description

Offer inbound/outbound SPAM filter at the gateway/ISP level

Motivation

Many ISP's cannot afford to offer spam filtering due to performance, accuracy or cost reasons. ISP's should deploy solutions that allow them to offer services that they can offer a small monthly charge for to obtain incremental revenue streams. Many vendors offer deals to partner with ISP's to licence features to their customer base.

Published mail error codes

Description

Inbound: incorporate more information in soft/hard bounce errors/documentated mail error codes a la Yahoo/AOL

Motivation

Put more information into MTA soft and hard bounces. Cryptic error messages lead to more helpdesk calls. Include simple language & URL's in the soft & hard bounces to direct users to the reason(s) why the message was not accepted. Good examples are provided by AOL on this (<http://postmaster.aol.com>).

Reject spam

Description

Inbound: drop all positive spam when possible, don't pass on (policy issue) plus "reject is better than accept-and-drop"

Motivation

Drop all positive spam where possible – assuming you have invested in a recognized anti-spam technology. By delivering all positive spam, this act only serves to fill up mailboxes, waste network bandwidth & generate user requests as to why it was not detected. Users whom want to see all their spam can opt out of this policy.

External accreditation

Description

Consider 3rd party reputation accreditation (e.g. Return Path, Habeas)

Motivation

ISP send a lot of email to the Internet. By using 3rd party accredited parties such as Return Path or Habeas, ISPs can get better mail delivery rates using the whitelists established by vendors. Note that this does entail some supplementary obligations on the ISP side.



Correct public records

Description

DNS RIPE/whois record data correct

Motivation

Make sure the DNS records for your organization's registered domains are up to date & correct. Admin's regularly check whois records in DNS to quickly find out who to contact with problems or questions.

Bounce handling

Description

Monitor bounces

Motivation

Many ISP's don't have time to check to see where all the non-deliverable mail is generated from or to where it goes to. By examining the bounce messages details of users with infections can be identified.

Monitor mail traffic

Description

Monitor IMAP/POP3 activity

Motivation

Commercial solutions are available to scan messages not only via SMTP but also within IMAP & POP protocols. Spam signatures are often reactive however messages are not always downloaded and read immediately after receiving. Anti-spam scanning during mail collection is another way to capture spam. Many vendors offer deals to partner with ISP's to license such features to their customer base.

Inbound port 53 block

Description

Block incoming port 53 to residential customers

Motivation

ISP's should rarely see incoming DNS queries. There are two reasons why you would see this:

- an individual hosting a DNS server (if so – they would have a static address)
- a individual's PC is compromised and being used for 'fast-flux' dns queries

ISP's should prevent the DNS ports from entering to prevent exacerbating the problem.

Use sender validation

Description

Use DKIM / Sender-id/ SPF on inbound e-mail

Motivation

Studies show that all organization sending email should use Domain Keys and / or publish SPF records to prove the identity of their SMTP mail servers. They work in different ways (SPF – publish via DNS where you end mail from) (Domain Keys – messages signed with private key, receivers validate message via public key published in DNS).

Large mail hosting organizations use either/both to stop spam & phishing messages reaching their users (Hotmail, Yahoo, Gmail, AOL etc.). If you don't use these techniques your customers messages are more likely to get moved to spam filters and not the recipients inbox.