

# IS ACCESS ANTI-SPAM PRODUCT OVERVIEW

PROPOSAL FOR IS ACCESS AT ITT CONNECT AND  
INTERNET SOLUTIONS





## Table of Contents

|  |   |
|--|---|
| Table of Contents .....  | 2 |
| Introduction – spam Overview .....   | 3 |
| What is “true spam”?.....  | 3 |
| E-mail address harvesting .....  | 3 |
| False positive rate .....  | 4 |
| IS’ Anti-Spam solution.....  | 4 |
| Solution Features .....  | 5 |
| Solution Benefits .....  | 5 |
| Brightmail anti-Spam software.....   | 5 |
| Brightmail evaluation criteria defined.....                                      | 6 |
| Open Source vs. Commercial anti- spam services, i.e. free vs. paid services..... | 8 |
| Glossary.....  | 9 |

### List of Tables

|  |   |
|--|---|
| Table 1 Anti-spam false positive rate comparison ..... | 4 |
|--|---|

### List of Figures

|   |   |
|---|---|
| Figure 1 Determining whether messages are Spam..... | 6 |
|---|---|

## Introduction – spam overview

Spam is a modern day scourge plaguing millions of e-mail users worldwide, it is the electronic equivalent of junk mail, wasting users' time and money - by virtue of the precious international bandwidth resources used in delivering spam to an end users mailbox, since most spam originates outside the borders of South Africa.

An estimated 30% to 50% of all e-mail messages on the Internet can be considered as spam, indicating that organisations are faced with an ever increasing, constantly evolving spam problem.

### What is “true spam”?

It is extremely difficult to define “true spam” without straying into the realm of opinion and personal choice. For the purposes of the IS' Anti-Spam service, true spam is defined as:

- ?? Pure spam — senders that don't exist, con artists, pornography
- ?? Chain letters, hoaxes, urban legends
- ?? There is another category which falls into a grey area since it comprises legitimate organisations trying to make a living, making use of so called “opt in” mailing lists. E-Mail is only sent to users if their consent has been given. It is up to the personal choice of the recipient whether or not the message is desired. In such cases, it should be up to the recipient, since blurring the definition of spam means an increased risk of deleting e-mail which an end user wishes to receive.

### E-mail address harvesting

Spammers "harvest" addresses in every conceivable place. If a user has ever registered a software product, asked a question on a technical-support bulletin board or participated in an online discussion group, their e-mail address may be pick-up by a spammer harvesting addresses.

Additionally, spammers launch harvesting attacks against companies in an effort to obtain valid e-mail addresses. A harvesting attack is launched by sending messages to the target company addressed to BobA, BobB, BobC, BobD, ... BobZ @company.co.za. The receiving message transfer agent (MTA) will be kept busy dealing with the response to this message, example: "BobA is not at this domain ... nor BobB ... nor BobC. But, I can deliver this one for you!" The spammer learns from the delivery and non-delivery reporting which e-mail addresses are valid and which are not. Anti-spam vendors have reported that nearly 50% of the connections made to their service are attempts to harvest addresses, which they fortunately have technology to rebuff.

*Note: a recent spam message advertises harvesting software for a nominal fee of \$39.95 that will harvest general e-mail lists from mail servers. The software boasts of an ability to harvest 100,000 e-mail addresses directly from e-mail servers in an hour!*

Another useful tool in the fight against e-mail address harvesting is the use of lax naming conventions – easy to guess naming conventions play into the hands of spammers. It is important to re-look at the specific naming convention applied by your organisation to ascertain whether or not this puts your organisation at risk from receiving spam as a result of a harvesting attack. It is important to have defences against harvesting attacks. If Bob.Arnold@company.com is a valid address, then Tom.Smith@company.com is probably also a valid address. Like random telephone dialling, spammers guess at addresses and try all manner of combinations. With a computer, compiling a randomized list of 1,000 user names is not difficult.

The longer a user has an e-mail address, the more likely they are to receive spam. A general rule of thumb is that if an employee has been with a company for more than a year, they will likely

begin receiving spam. One solution is to change addresses, but that is annoying enough as a consumer and nearly impossible for a business user to do. Besides, it's only a temporary tactic. The spammers will "learn" the new address and, within a year, you will be back to the same level of spam. Another solution is to deploy an Anti-spam solution.

**False positive rate**

The accuracy of any anti-spam product is defined by the stated false positive rate, defined as the percentage of legitimate e-mail messages that are incorrectly identified as spam.

Incorrectly filtering small amounts of legitimate e-mail creates the same (or more) productivity loss as spam itself, so it is essential to minimise as far as possible the false positive rate of the anti-spam mechanism deployed. Illustrative examples:

- 1) A journalist sends an announcement of a friend's wedding to 100 of his friends who work for various companies. The message is full of excitement and CAPITAL LETTERS and, for emphasis, many exclamation marks !!!!!, as well as an invitation to give money to a charity in honour of the happy couple. This message may be identified as spam by tools in place at nearly half the recipients companies.
- 2) A bookstore catering to mystery books lovers is corresponding with one of its customers. It has no trouble receiving the customer's messages, but the shop's replies are not getting through to the customer. Finally, the customer calls their Internet Service Provider (ISP), and finds out that the ISP is blocking messages from mysterylovers.com because it has the word "lover" in the domain name.

These are typical problems encountered when using a poor anti-spam tool - one that leaps too quickly to the wrong conclusions. These types of tools usually stop around 20 to 30 percent of the spam, and catch too many false positives. A slightly better class of tool uses "point system," where each spam tactic identified in the message earns one or more points. The message is not declared to be spam, however, until a certain threshold of points is reached. In the above examples, the e-mails would not garner enough points to qualify them as spam, and thus are likely to have been received by the recipients.

IS in choosing an anti-spam software vendor performed a due diligence on the various options available, the table below contrasts the stated false positive rate of various products evaluated (the table reflects a representative sample).

|                                   | <b>Brightmail</b>           | <b>PureMessage</b>           | <b>SpamAssasin</b>        | <b>MailMarshall</b>          |
|-----------------------------------|-----------------------------|------------------------------|---------------------------|------------------------------|
| <b>Stated False Positive Rate</b> | 0.0001% i.e. 1 in 1 million | 0.01%, i.e. 1 in 10 thousand | Depends on implementation | 0.01%, i.e. 1 in 10 thousand |

**Table 1 Anti-spam false positive rate comparison**

By far the best option is to deploy an anti-spam tool that makes use of a multitude of methods in identifying individual messages as spam, thereby minimising the risk of false positives.

**IS' Anti-Spam solution**

Internet Solutions (IS) recognises that e-mail is a central part of doing business today and as such an organisations' messaging platform is a critical system that must be designed with high availability, redundancy and scalability in mind.

IS has partnered with Brightmail a US based company in providing its Anti-Spam service to customers. Brightmail is generally viewed as having the dominant share of the global service provider market and is a market leader in the Anti-spam space.

## Solution Features

- ?? **Outsourced e-mail spam scanning engine:** the Anti-Spam service is a fast, easy-to-use anti-spam solution that enforces your organisation's IT Acceptable Use Policy (AUP) while protecting against spam and loss of confidential data.
- ?? **Shared platform solution:** the Anti-Spam service is available on a shared platform, leveraging the benefits of economies of scale for our customers, delivering cheaper, more sustainable solutions.
- ?? **Mail server protection:** integral to the services' anti-spam capabilities.
- ?? **Mail spool:** in the instance that a client's line or mail server is no longer available, IS Anti-Spam spools mail on behalf of a customer, hence providing enhanced uptime for a client's mail service.
- ?? **Reporting services:** detailed reports are available to clients. These provide excellent management information and illustrate email usage patterns and savings. Reports are made available 24 hours a day via the IS Customer Zone.
- ?? **Fully supported, high availability solution:** support staff are available 24 hours a day to ensure the system is always available. IS is responsible for ensuring that the environment and service components are fully redundant and thus offer maximum availability.

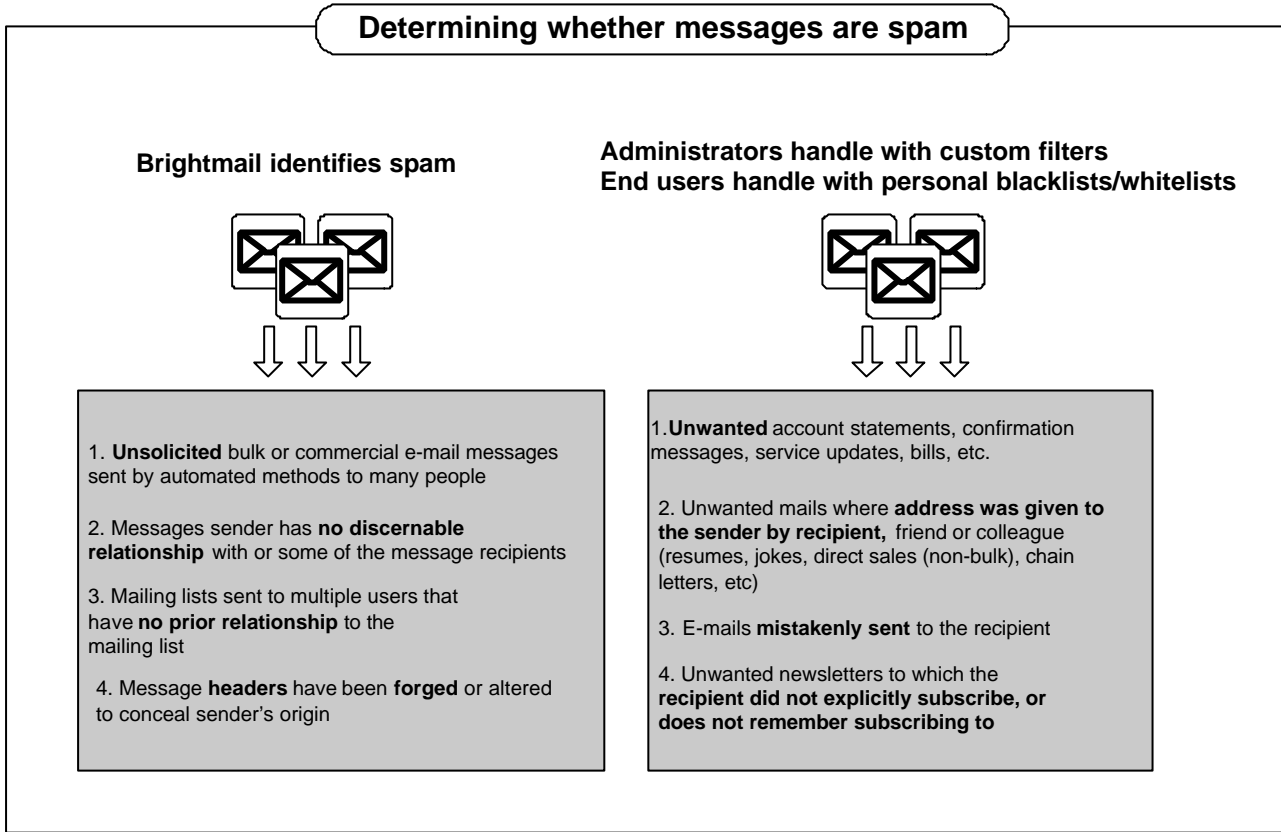
## Solution Benefits

- ?? **Enhance network protection and enforce your organisation's AUP:** as e-mail becomes the most important channel for business communications, scanning of e-mail is increasingly recognised as a necessary component of network protection for businesses.
- ?? **Increase leased line efficiency:** Anti-Spam has been shown to block up to 60% of incoming mail, translating to a direct cost saving on bandwidth and leading to much more efficient use of the Internet line.
- ?? **Outsourced and cost effective solution:** the solution is entirely outsourced; hence costly software and hardware purchases are avoided, as well as costs to hire internal staff to provide a similar solution are avoided. Software licenses are negotiated in bulk, and hardware is shared amongst clients to reduce costs thus offering a solution that benefits from these scale economies.

## Brightmail anti-Spam software

Brightmail has a strict definition of spam, taking great care to distinguish messages which are legitimate from those which are not, without impinging on the end users personal choice. Additionally, Brightmail can identify unsolicited commercial e-mail (legitimate companies that have not done double-opt-in list registration) and warns the sender to pay greater attention or risk being labelled a "spammer."

It is imperative that all parties within an organisation agree on the definition of spam. Brightmail uses the following guidelines to distinguish spam from legitimate email communication and recommends that companies formulate an e-mail usage guideline based on the principles outlined in the diagram below.



**Figure 1 Determining whether messages are Spam**

Diagram Source: Brightmail Anti Spam Enterprise Edition 5.5 Reviewer's Guide

**Brightmail evaluation criteria defined**

The Brightmail rules module provides primary anti-spam protection. In order to keep up to date with the latest spam attacks, Brightmail employs automated rule creation and delivery technologies, delivering updated Brightmail rule sets to IS' Brightmail servers approximately every 5-10 minutes.

Within the rules module, there are five types of filters, each designed to combat different types of spam, necessary in order to combat the complex spam attacks that spammers now resort to. The filters have the following functionality:

| Filter Type | Methodology  | Purpose  |
|-------------|--|--|
| Body Hash   | Reduces the body of an email message to an essential fingerprint   | Traps spam characterised by a common message body and complex, highly randomised headers.  |
| BrightSig2  | Strips Spam messages of random HTML code and incorporates fuzzy analysis to identify the underlying "DNA" of an evolving Spam attack | Defuses HTML-based spam attacks that evade most filtering techniques. Allows Brightmail to group seemingly random spam messages into a common attack that can be efficiently filtered. |
| Header      | Users expression filters that target the headers and subject   | Creates tight, targeted filters that identify telltale spam  |

|           |  |  |
|-----------|--|--|
|           | lines of Spam messages   | characteristics with almost no false positives.  |
| Heuristic | Scores messages against a large set of heuristic filters. If a message achieves more than a specified score, it is considered to be Spam             | Enables the proactive identification of spam based on inherent characteristics. The application of weights to each rule, guards against false positives. |
| URL       | Matches embedded URLs that often appear in Spam messages with a list of Spam URL's compiled by the Brightmail Logistics and Operations Centre (BLOC) | Identifies and filters spammer's untended URL, which is often the sole purpose of prevalent "all to action" spam messages.                               |

Source: Brightmail Anti Spam Enterprise Edition 5.5 Reviewer's Guide



## Open Source vs. Commercial anti- spam services, i.e. free vs. paid services

It is not the intention of this document to open the Open Source vs. Commercial services debate, however, the issue of free vs. paid for anti-spam solutions requires urgent redress.

IS in deploying commercially available anti-spam software and expensive hardware within the IS network are forced to recoup this cost by billing our customers for the Anti-Spam service.

This is in fact one of the large differentiators between IS' Anti-Spam service and that of other service providers offering free anti-spam services. IS leverages Brightmail's BLOC which employs teams of people to constantly monitor and identify new forms of spam presenting themselves to users. New rules are created based on their findings which are sent to Brightmail's customers' servers (example IS) around the world every 10 -15 minutes.

Service providers running free anti-spam services generally do so by deploying Open Source anti-spam software (example SpamAssasin) which places the onus on the service provider to constantly update the spam filter rules, this in turn requires teams of people do accomplish. One has to ask the question: "can this be done effectively by a single organisation, without incurring large resourcing and infrastructure costs – thereby jeopardising the free model?"

**Bottom Line:** Low-cost and inadequate spam-fighting tools will not stop spam, and will require a great deal of staff time to manage the product and associated complaints. Every enterprise should implement a competent product or service to guard the boundary of the enterprise. It is, unfortunately, part of the cost of doing business on the Internet.



## Glossary

**“Blacklists and White Lists”:** The use of blacklists as the only anti-spam tactic is entirely unsatisfactory. Used as one data point in a point system, blacklists can be helpful. An enterprise can also create a white list of domains that are always allowed to receive e-mail, no matter what their content is;

**“Content Analysis”:** Includes one or more of the following capabilities:

- ?? A set of rules to search for known spammer tactics;
- ?? A set of rules to search for known chain letters, hoaxes and urban legends;
- ?? The ability to look for words and phrases in a targeted "words list" (for example, porn, financial services);
- ?? The ability to do contextual analysis;

**“Heuristics”:** Algorithms, whose intention it is to predict messages which may be tomorrow's spam based on characteristics in the header or body. The term heuristics is most-often applied to viruses, where algorithms try to find mutants of previously seen viruses, or new initiatives containing similar malicious behaviours. The term is less precise when applied to spam;