BMW AG

Comparing antispam appliances and services

SPAM SLAM

Spam filters can help smooth the waves in your inbox, as long as they are reliable and don't have too many side effects. We'll show you what we found when we tested five antispam appliances and two service providers. **BY TOBIAS EGGENDORFER**

Filtering spam as it enters the network keeps control in the hands of IT professionals, increasing the effectiveness of the filter and improving user productivity by reducing the glut of mail. In previous issues of Linux Magazine, we have outlined various software-based scenarios for integrating spam filtering with a mail server. In this article, we'll look at some alternatives to conventional software-based filters such as SpamAssassin. In particular, we'll report on some hardware-based spam filtering appliances, and we'll examine a representative pair of Internet spam filtering service providers.

The Test Environment

The typical approach to testing spam products is to take known spam from large archives [2] and run the filters against it (see the "Filtering Techniques" box). At best, this will tell you whether the vendors have done their homework, and whether their filters have been tested and optimized to combat known spam. However, spammers constantly develop new obfuscation methods with the aim of tricking existing filters, and they continually test their junk mail against known filtering techniques.

Spammers often set up email accounts with major mail providers to see if junk mail delivered to the account junk will make it past the filter. Some bulk mail programs even integrate SpamAssassin [3] in order to test mass mail before sending it on its way. The spammer will modify the text until the message finally slips past popular filters.

Thus, the detection rate for known spam does not say much about the quality of a filter's heuristics. It is hard to say how a filter will react to the daily flood of spam based on historic data. This is

Mail Exchange in DNS

A Mail Transfer Agent (MTA) first discovers the IP address to which it is permitted to deliver an email message [4]. To do so, it extracts the domain name from the mail address and sends an MX Query (Mail Exchange) to a DNS server. The server responds with the IP address of the authorized mail server. For redundancy, domains are allowed to have multiple MX entries with different priorities, where a lower value denotes a higher priority.

The sending MTA establishes a connection to the top priority MX and attempts to deliver the message to the MX. If the attempt fails, the MTA contacts the MX with the next highest priority. The secondary mail servers then attempt to pass the messages on to the internal mail server with the highest priority.

In productive environments, this typically means that the highest priority MX is located on the company premises and managed by the local IT department. Strict spam filtering typically applies here. The carrier often has additional backup MX servers. These servers do not usually apply filters, so this target is far more lucrative for a spammer. If the spammer succeeds in delivering their junk mail, they have reached their target, or at least managed to evade a number of filters, such as IP-based blacklists. The blacklists will view the provider's legitimate, low-priority MX as the source, rather than as the spammer.

Figure 1: Five appliances took part in our lab. From bottom to top: Symantec Mail Security 8260, McAfee Secure Content Management Appliance 3200, Ironport C10 Email Security Appliance, Canit Anti-Spam, and IKU Sponts-Box.

what prompted us to let the filters tackle unknown spam fresh off the Internet. Simply redirecting an existing mail account would not be sufficient, because the spam battle starts with the SMTP dialog. Thus, each test candidate needed its own domain and its own DNS MX entry (see the box titled "Mail Exchange in DNS").

Four Months of Fresh Spam

The first thing we needed to do was to guarantee a continuous supply of spam to multiple addresses. To do so, we registered ten domains and published a website with four email addresses in each domain. The domains went online six months before the test started, and many links from other pages pointed to them. We added a selection of keywords to whet the harvesters' appetites, and we sat back to wait for the spammers to take the bait.

About six months later, the domains and the two webmail accounts were perfectly prepared; they each received enough mail and similar amounts of spam. During the main test phase, each victim address received 50 to 100 fresh spam mails.

The mail server for the domains was a machine with ten IP addresses assigned to it. Each of the 10 IPs was entered as the MX for one registered domain. We used Sendmail to pick up the messages. Sendmail simply passed the messages to a small Perl script, quoting the envelope to address; the script separated the body from the header and entered the results in a database.

When the appliances started to arrive – some accompanied by a service engineer and others by snailmail – each device was assigned a domain and the MX IP address (Figure 2). This avoided the need to change the DNS data, which might have then affected the test results.

SMTP Proxy

Of course, we needed to change the MX entry to accommodate the service providers, Expurgate and Spam Stops Here. To use these services, you need to enter the provider's IP as the MX for your own domain. Expurgate and Spam Stops Here analyze and filter the incoming messages and forward them to the target mail server. This approach outsources spam filter management to the service provider, however, relying on an external provider to manage your mail environment does mean having a lot of trust.

The appliances we tested work like incoming SMTP proxies: the external mail server uses SMTP to contact the appliance and attempt to deliver mail. If the mail passes the filter, the appliance then uses SMTP again to route it to the local mail server. It's business as usual for the local mail server, and for the users on the internal network who remain blissfully unaware of all this. They can continue to either fetch mail from the server or read their mailboxes on the server itself.

We deliberately avoided testing the virus filtering features of the appliances. Virus checking was beyond the scope of our test, and for that matter, our test would not have revealed much about the appliance, since detection rate is defined by the virus scanner the appliance deploys. However, a virus scanner can also reduce the volume of unsolicited mail by removing any worms it detects from the mailbox. Some products will even let you deploy multiple scanners at the same time to compensate for inadequacies of individual scanners, a capability that is reflected in the licensing fee.

Spam Quality

Not all of the vendors we contacted actually provided test equipment, and this left three test domains unused – we had originally planned for one domain to provide comparative results for filtered versus unfiltered domains. What we originally considered to be a drawback turned out to be a bonus, giving us the ability to ascertain a more stable mean value for the spam volume.

As all the email addresses were known only to spammers, the mail servers received only spam at first. This is useful for ascertaining a filter's detection rate, however, it does not tell you anything about false positives; and false positives are a major criterion for defining spam filter quality. A single legitimate mail message incorrectly classified as spam by the filter can cause more damage than letting ten spam messages through.

To ascertain the false positive rate, we bombarded the candidates with legitimate messages toward the end of the

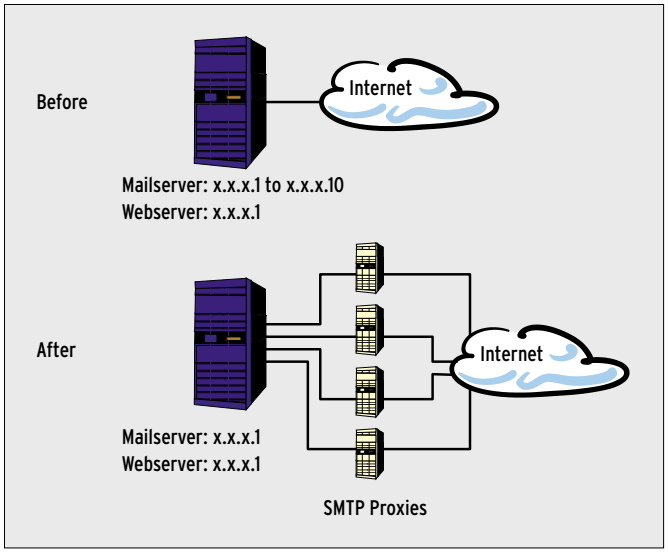

Figure 2: During the harvester phase (top), a mail server collected spam for 10 IP addresses in 10 separated domains. During the test phase, the appliances were introduced as SMTP proxies for the addresses.

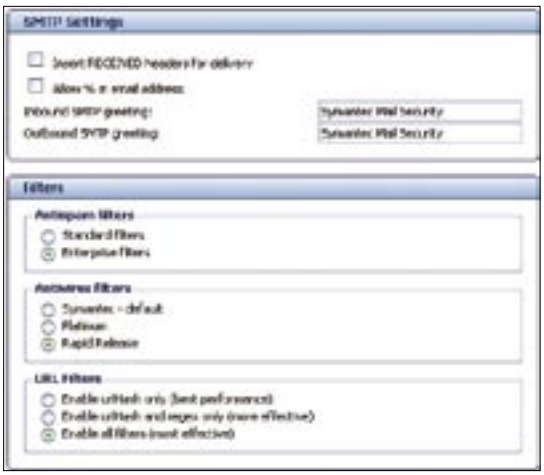

Figure 3: This hidden menu appears on the Symantec appliance when the admin presses [Shift]+[A] on the web interface. The Rapid Release option was set manually by the Symantic engineer. The SMTP Greeting settings are also quite interesting.

test. To allow this to happen, a large number of users let us use their inboxes, thus ensuring a good representative sample of genuine *ham* (non-spam) messages. Users either forwarded messages from their own mailboxes, including newsletters, or wrote their own text emails, some in foreign languages such as Bulgarian and Turkish. Of course, this is torture for filters that attempt to perform language analysis.

At the end of the test phase, we put the systems through stress testing. For a couple of days, we had the MX records of two domains, which are subject to a volume of spam a couple of orders greater than in our test environment, redirect their junk mail to one appliance, thus forcing it to handle somewhere in the region of 35,000 emails per day.

Symantec Mail Security 8260

All vendors were asked to supply an off-the-shelf product suitable for a small to medium-sized business. Despite this, Symantec sent us its top-notch device, the 8260. This machine is designed for enterprises with upward of 1000 mail accounts, and it is said to be able to handle up to 10,000 accounts.

The system supports clustering for environments that require higher performance. The hardware comprised a Dell Poweredge Server with a 19'' rack case and a front plate that was customized by Symantec. The flagship product sports two Xeon CPUs with a clock speed of 3 GHz. 2 GB RAM and two Raid 1 mirrored 73 GB hard disks provide sufficient storage capacity. The doubly redundant power supplies underline the fact that the box is clearly targeted at the professional market. The operating system is Red Hat Enterprise Linux 3.0. Root login is not envisaged. According to Symantec, the system uses a hardened version of Postfix.

The configuration interface is web-based via HTTPS on port 41433. Initial configuration can either use a serial console or the local keyboard and display. The initial configuration simply sets up the hostname and network environment. The license key, and filter direction setup (incoming or outgoing), both use the web interface. The Symantec engineer dug into his bag of tricks and conjured up a secret menu by pressing a keyboard shortcut [Shift] + [A] in the *Settings* menu (Figure 3).

Secret Menu

The engineer enabled the *Rapid Release* setting for the antivirus filter; according to the engineer, this setting is standard for other customers. As we were not interested in the virus filtering functionality, we left the setting as is. The secret menu gave us access to more interesting settings. For example, you can modify the SMTP greeting, which gives admins the ability to hide the appliance from simple probes launched by attackers.

The spam filter is based on Brightmail's Antispam; Symantec acquired Brightmail in June 2004. Among other things, it uses SPF, various blacklists and whitelists, URL filtering, and Sender ID for detection purposes. Additionally, Symantec uses signature detection and a hash-based comparative detection algorithm. Each individual filter returns a value used to score the spam probability. The values all contribute to a fixed, and non-customizable, total spam score.

There is no way for us to disable the different filtering mechanisms on an individual basis. However, it is possible for us to set individual values for the spam score for user groups (defined on an LDAP server).

In our lab, the Symantec system achieved a perfect score of no false positives, but it identified less than 90 percent of all spam, thus letting over 10 ercent of all spam through (Figure 4).

McAfee Secure Content Management Appliance 3200

McAfee also shops with Dell. The McAfee 3200 appliance supplied the Symantec appliance's little brother: a 2.8 GHz Xeon and 1GB RAM are all McAfee needs for up to 1000 users. The hard disk subsystem is a SCSI Raid 1 array. McAfee also uses Red Hat Linux, however, the appliance has a lot more software on board. A Secure Web Gateway was pre-installed on the test device. (This product is available individually, and was not taken into consideration in our lab.)

The McAfee appliance again used the local console for initial configuration, but it also supports a web interface-based approach that uses a cross-over cable to connect the client to the appliance. HTTPS is simply used as the transport

Appliance Security

All the appliances we tested included the word *security* in their names, however, some devices worried us in this respect by using open HTTP for client server communications, or by providing a jumble of tempting-looking Javascripts as a web interface. Didn't there used to be a rule for secure systems that said something about reducing the services to a minimum and keeping the software as simple as possible? Spam filters are very complex bits of software. Adding a web interface to keep users happy with gadgets and gimmicks is very likely to increase their vulnerability. Every single line of code increases the probability of a bug, and thus increases the likelihood of a security hole.

One of our wishes that has remained unfulfilled thus far was for the configuration to support local console-based administration, such as the ability to disable the web server and other remote services. This feature would give the customer the ability to opt for less security risk with (ostensibly) less convenience. The fact that some vendors don't even let users customize the spam filtering rules shows how highly they value their customers' skills.

protocol; client-side, you either need a Java client program, or you need to run a Java applet in your browser. The applet takes slightly longer than rendering a normal HTML page. All in all, the clear-cut Java interface was fairly sluggish and seemed to respond more slowly to user interaction than any other user interface. Unfortunately, the client program did very little to improve performance.

To run the McAfee spam killer, you first need an activation CD. To create the CD, you first need to download the image file (this is just a couple of MB) from the McAfee company homepage, then burn the image onto a CD and insert the CD into the appliance's drive to read the CD. Instead of this convoluted, and environmentally-unfriendly approach, it might be easier and more understandable to customers to simply upload the ISO image directly using the web interface. And it might be even easier to allow the appliance to contact the McAfee site directly.

Like almost all the vendors in our test, McAfee uses a mix of several filter systems. Besides SpamAssassin, they include a Bayesian filter, blacklists and whitelists, and sender authentication. In contrast to Symantec, McAfee supports very granular filter configuration. Although we kept the default settings for our test, the system achieved a laudable result of 97 percent. Unfortunately, it did so at the expense of the false positive rate of no less than 7 percent. This said, the device drops spam into a quarantine folder, which the recipient can access if needed.

Ironport C10 Email Security

Ironport's case makes a nice change from the typical gray masses: the silver 19 inch, 1 HE unit looks elegant, and the unit is custom-designed as an appliance. The vendor even fits a blind cap to the VGA output to keep inquisitive administrators' fingers firmly away from the box (Figure 5).

The interior design is just as unique: the operating system the appliance runs on is called Async OS. This is a mail filter-optimized Linux that works far more efficiently than normal Linux according to Ironport. Apart from this, the vendor states that the system has two 40 GB disks in a RAID 1 array. Instead of publishing the CPU speed, or memory, Ironport simply states that the C10 can protect up to 1000 mail users against spam. This puts it in the same league as the McAfee appliance.

Ironport offers a choice of two spam filters, Brightmail, which is also Symantec's choice, and a proprietary product. We used the second variant in our lab. The filter had a convincing range of granular settings, which admins can configure in a web-based menu. The fact that the web interface uses HTTP rather than HTTPS slightly tarnishes the good impression.

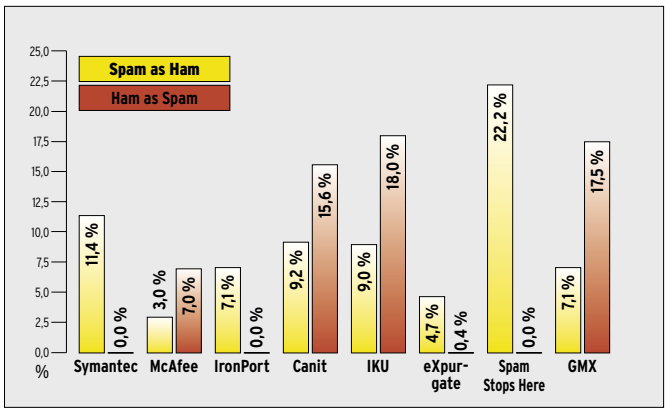

Figure 4: Separate error rates are shown for false negatives (spam incorrectly identified as legitimate) and false positives (legitimate mail identified as spam). The false negative value for Sponts is extrapolated (estimated error rate of between -3 and +5 percent).

Again, we used the default settings for Ironport. In addition to the typical spam filters, Ironport uses an image filter and an Ironport-specific system known as Senderbase. According to the vendor, Senderbase logs much of the global mail traffic, and is thus capable of quickly detecting new waves of spam and malware outbreaks. In our lab, ham detection proved reliable – no false positives. Ironport had the best spam detection rate of all system with a perfect false positive rate of zeros (just 7 percent of all spam got through).

CanIt Anti-Spam

The CanIt Anti-Spam appliance, which is developed by Roaring Penguin, is supplied in a 1 HE 19", half-depth case. Under the hood, the system has a 3 GHz Pentium 4 CPU, 1 GB RAM, and a 80 GB IDE hard disk.

The manufacturer describes the Debian 3.1-based system as the "leading anti-spam solution." We would question this based on the measured values: The CanIt appliance had the second-highest false positive rate, but it failed to compensate by catching more spam. In all, CanIt had a spam detection rate that puts it firmly in the bottom half of the test field.

On the other hand, the system surprised us with some useful features: users are allowed to define their own filter rules, which the administrator can pre-configure. The box even includes a user manual. Although the manual may be difficult reading for those without some prior knowledge, users who prefer to set up individual spam protection at least have the controls to do so.

Unfortunately, the user interface is unintuitive, but once you have mastered the quirky interface, CanIt gives a wide range of options that might let you improve the poor detection rate by tweaking the settings.

Immature

At various places, the system creates an impression of being immature. For example, setup mode, which requires the administrator to attach a display and a keyboard, has an option to change the root password and the password for the *setup* user. You can use the menu to change the root password, but to change the password for setup, you need to log in to the console as root and enter *passwd setup* manually at the command line.

It is also annoying that the web interface for the box only uses HTTP by default. Referring to HTTPS, the manual says "Setting this up is beyond the scope of this manual, but CanIT-Pro should operate with no changes over HTTPS." The second part of the sentence turned out to be true.

IKU Sponts

Linux Magazine tested the IKU Sponts Appliance a couple of years ago. According to the description on the website, the case should have been either a Mini ITX or a 19'' rack system, but what we got was a Mini ATX system. Apart from this,

the other components seemed to match the Sponts specifications. The big advantage this system offers is that it has no wear and tear parts, apart from the 40 GB IDE disk, and this should mean a longer product life. IKU quotes a limit of 550,000 messages a day for the Debian-based system, which puts it more or less on par with the 1000 users quoted by McAfee and Ironport.

We plugged in a keyboard and a display for the initial setup. The system expects you to enter the network settings at the console. We removed the keyboard for the next boot, and the system hung. The BIOS was set to interrupt the boot process if the keyboard was missing. Fortunately, we soon identified the source of the bug.

Neat Interface

The system's user interface is well designed, and the internal help function really does help. The ability to store messages temporarily on the Sponts Box, and to relay them to the internal mail server via the *Replay* is also useful. The box also runs a POP3 server to keep email service available in case of an internal mail server outage. This is a useful feature that other appliances should think about introducing.

The Sponts Box uses two special approaches to spam protection: the Sponts Effect and SMTP transmission timing analysis. In addition to this, the system has a number of standard filter techniques in place, all of which support individual weighting and customization.

In our lab, the Sponts Box was configured to respond with a *User unkown* on receiving spam. Because of this, the mail server downstream did not get to see any spam mails, apart from the ones that got past the filter, of course. We were unable to ascertain the detection rate due to this. Although the Sponts Box's logfiles record connection attempts, this value is no help, because spammers sometimes open connections to test mail addresses. Figure 4 shows a detection rate that we extrapolated by reference to the comparative values from the other domains. The ham tests revealed that the Sponts Box is far too aggressive by default. A false positive rate of 18 percent put the box firmly at the bottom of the pack.

Expurgate

Instead of buying an appliance, updating it regularly, and managing the system yourself, you can outsource the job to a service provider. The Expurgate service provides external spam filtering. Expurgate is operated by a German company called Eleven. Customers simply set the MX record for their domains to the Expurgate mail server. To use Expurgate, you need to set up four MX records with the same priority. Each of these hostnames resolves to multiple IP addresses in various subnets, thus implementing

Table 1: The Candidates

Company	Symantec	McAfee	Ironport	Roaring Penguin	IKU	Eleven	Greenview Data
Product	Mail Security 8260	Secure Content Management Appliance 3200	C10 Email Security Appliance	Canit Anti-Spam Appliance	Sponts-Box	Expurgate	Spam Stops Here
Price	EUR 6,310 plus licensing charge. Rates for virus protection and antispam: EUR 35 per license for 100 licenses or EUR 23 per license for 500	EUR 17,680 perpetual license for any number of protected systems, no time limits. Includes one year's support.	EUR 2,000 for 50 users and one year	EUR 2,690 for 300 users, EUR 4 for each additional user	Starts at EUR 664	10 users cost EUR 25 per month, EUR 42 with virus protection. For 1,000 users, EUR 970 without, and EUR 1,700 with virus protection. Discount for 3 year advance payment.	EUR 14 per mail box and year (for 10 to 19 mail boxes). 10,000 mailboxes or more, EUR 6 each plus domain. Maximum transfer per mailbox and month.
Type	Appliance	Appliance	Appliance	Appliance	Appliance	Service	Service
Accounts	More than 1,000	Up to 210,000 messages per hour	Up to 1,000	Not disclosed	550,000 Mails per day	–	–
POP and/ or IMAP	No	No	No	No	Yes	No	No
SMTP Relay	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Administration via	Web interface	Web interface and Java applet	Web interface	Web interface	Web interface	Web page	Web page
Licensing model	Account driven	No limits	Time and account driven	Account driven	By appliance	Time and account driven	Account and domin driven
Filters configurable per user	LDAP users only	No	No	For every user	No	Exceptions for single account	Yes
Hardware	2 x 3 GHz Xeon, 2 GB RAM, 73 GB SCSI Raid 1, redundant power supply	2.8 GHz Xeon, 1 GB RAM, SCSI Raid 1	40 GB Raid 1, not disclosed otherwise	3 GHz Pentium 4, 1 GB RAM, 80 GB IDE	40 GB IDE, not disclosed otherwise	–	–
OS	RHE 3.0	RHE 3.0	Async OS	Debian 3.1	Debian	–	–

Figure 5: Ironport fits a cap to the VGA port; attaching a display is not an option.

redundancy via DNS round robin load balancing. Redundancy is an important advantage of a service provider in comparison to an appliance-based solution.

The Eleven spam filters benefit from the fact that they protect multiple mail servers, which means they come in contact with many spam messages. The service calculates checksums for all incoming messages and uses the checksums to compare the messages. If a checksum occurs more frequently than the mean value, the message can be assumed to be spam. This approach is similar to Ironport's Senderbase approach, but it does away with update cycles.

Endangered Species

Expurgate uses various criteria to prevent the filter from falsely identifying legitimate newsletters as spam. One criterion is the delivery path: spam typically originates from multiple machines on a bot network, whereas newsletters originate from a single source. In addition to this, the company uses what they refer to as spamtrap addresses. These are email addresses only published on websites – much like the addresses used for this test. When the spamtraps receive a message, it is very likely to be spam.

An additional feedback channel gives the ability to tag mail individually as spam or ham. To avoid misuse and errors, Eleven states that a staff member checks messages for legitimacy because users have been known to return phishing mails tagged as spam with notes to the effect that the mails were important messages from their banks.

Unfortunately, administrators have more or less no options for configuring the Expurgate filters. Administrators can except email addresses from filtering and define actions to execute when a message is classified as belonging to a specific spam category. More granularity would be preferable here. Despite this, Expurgate still achieved the second highest spam detection in our lab, at the expense of 0.4 percent false positives.

Spam Stops Here

Canada's Greenview Data follows an approach similar to Expurgate with their "Spam Stops Here" product, except that the configuration options are more granular. For example, administrators can enable and disable filter modules or stop filtering for specific mail boxes. Surprisingly, the spam detection rate for the default settings is very poor: 77.8 percent puts Spam Stops Here right down at the bottom of the list. The product did not produce any false positives.

Compensation Business

The results in Figure 4 are not surprising: the higher the spam detection rate, the higher the false positive tends to be. A better detection rate (less spam getting through) means using more aggressive filters, which may sacrifice ham mails.

It is interesting to see how filters that apply comparative techniques (such as Senderbase and Expurgate) can improve detection rates without affecting false positive performance. Although both values are approximations due to the way the two systems work, they do allow us to evaluate the filter quality. The improved results may also have something to do with a new breed of spam that started to emerge during the test period: image spam hides the message in arbitrarily mangled images. The filters had to become accustomed to the new spam, and rule-based approaches tend to be at a disadvantage in comparison to comparative techniques at first.

There is still a latent danger of comparative filters classifying mailing lists with large numbers of targets as spam. Whitelisting legitimate mailing lists can help to solve this problem, although the sheer bulk of mailing lists makes this difficult to do. To compensate, vendors use additional filtering techniques to further reduce the risk of false positives.

Where vendors combine aggressive filters with automatic error message generation in the SMTP dialog, and this is what Sponts does, a fatal combination can occur. The idea that the sender will open, read, and understand *Mailer-Daemon* messages is one that would only occur to an engineer: users are more likely to complain of being spammed by *Mailer-Daemon*. If you actually read the message, you discover cryptic wording that is more or less impossible to understand. Spam filters should tag, but not remove – they are just not accurate enough to make final choices.

Three Winners

If you take the false positive rate as a important criterion, and you expect a detection rate of more than 80 percent, the contenders left standing are Expurgate, Symantec, and Ironport. Due to its ASP approach (Application Service Provider), Expurgate is hard to compare with the two appliances, and the configuration options were more than spartan.

With respect to filter quality and usability, the Ironport C10 wins against the Symantec solution. We found no big differences in filter quality (Symantec failed to detect 11 percent, and Ironport 7 percent of all spam), and Symantec could easily overtake Ironport following the next update. One positive aspect for Linux fans is that all of the filters we tested are based on Linux systems. ■

Support

The complex test setup was only possible thanks to the Bundeswehr University in Munich, Germany. The Institute for Information Technology Systems provided the mail server, the network connection, and the rackspace for the appliances, in addition to registering the test domains. Lieutentant Carsten Schulz set up the test systems, and performed the measurements, as part of his thesis. Thanks also to Daniel Rehbein, who provided access to two heavily spammed domains for the test series.

INFO

[1] Spam-o-Meter: *http://www.spam-o-meter.com*

[2] Spam archive: *http://spamlinks.net/filter-archives.htm*

[3] SpamAssassin: *http://spamassassin.apache.org*

[4] RFC 0974, "Mail Routing and the Domain System": *http://www.ietf.org/rfc/rfc0974.txt*

[5] "Bot Posse: An insidious botnet attacks Charly" by Charly Kühnast, *Linux Magazine,* August 2006, pg. 68