

Buyer's Guide: Anti-spam

Test: Spam in the wild

We throw real traffic at 16 anti-spam products

By Joel Snyder
Network World, 09/15/03

Practically every vendor on the planet claims to be able to solve your spam woes. So we tested 16 products on a live production network to see who could back those claims. For the entire month of June, we threw a live mail stream, spam and all, at the products to see who could survive the spam onslaught, and who would choke.

Estimates of the amount of unwanted e-mail range from 40% to 75%, but we can give you an exact percentage - 69%. That's how much spam we saw during the month of June. And things are getting worse, not better - in a similar Network World test we ran in February (see review), only 50% of the mail stream was spam.

Fortunately, good products prevailed, and can help you significantly reduce your spam problem. With a very broad field, including service-based mail filters, appliances and traditional software on Unix and Windows, network managers should be able to solve their spam problem with a minimum of disruption - to the accolades of their users.

How well do they work?

We tested mail-filtering gateways by feeding them an e-mail stream in real time, as it came into our labs (see "How we did it"). Each product received two scores. The first score, sensitivity, measures how well the filter identified spam. A perfect score would be 100%. The second score is the false-positive rate, the ability of the filter to make sure that non-spam messages do not get tagged as spam. A perfect false-positive rate would be 0%. (For more about the different ways to measure a spam filter, see "Spam and statistics".)

Spam filtering is such that a high sensitivity naturally also would have a high false-positive rate. Similarly, a low false-positive rate might let a lot of spam through. We feel that enterprise network managers would be more concerned with false positives, so we asked the vendors to tune their products for a false-positive rate of about 1%.

Products from seven companies - Cloudmark, Corvigo, MailFrontier, MX Logic, Postini, Trend Micro and Tumbleweed Communications - met our 1% requirement.

To identify the top products in filtering spam, we looked for a sensitivity rate of at least 80%. Products from seven companies also met that level - ActiveState, Cloudmark, Computer Mail Services, MailFrontier, Postini, Singlefin and Tumbleweed. (For complete results, see graphic.)

Combining these lists gives us the top overall performers: Cloudmark's Authority, MailFrontier's Anti-Spam Gateway (ASG), Postini's Perimeter Manager and Tumbleweed's Messaging Management System (MMS).

Of course, your results will vary, depending on your own message-stream characteristics and how well you tune the products. For example, Postini's spam-detection engine is at the heart of Trend Micro's recently released Spam Prevention Service (SPS). However, we got very different results with the two products, largely because Postini officials told us to tune their product using one set of numbers, while the Trend Micro team gave us a different set. This resulted in both a higher false-positive rate and lower spam sensitivity for Trend Micro.

Many vendors predicted that their false-positive rate would be much lower than 1%. Corvigo's CTO said its cus-

Accuracy

Percentage of spam the filters caught. An 80% accuracy rate would be acceptable in an enterprise setting.

Vendor	Accuracy
Postini	94.0%
MailFrontier	89.4%
ActiveState (Probable and Maybe folders)	89.4%
Singlefin	86.2%
Cloudmark (Spam and potential spam folder)	85.1%
Corvigo (Junk + Bulk folder)	84.6%
Computer Mail Services	83.4%
Cloudmark (Spam folder only)	82.0%
Tumbleweed (SpamHi + Spam folder)	81.3%
ActiveState (Probable folder only)	80.5%
Corvigo (Junk folder only)	77.9%
MX Logic	77.0%
SurfControl	76.5%
Tumbleweed (SpamHi folder only)	72.2%
Trend Micro	60.3%
Clearswift	48.5%
EasyLink	23.1%
GFI	3.6%

False positives

The percentage of non-spam messages that are marked as spam. A 1% rate or lower would be deemed acceptable in a company.

Vendor	False positive
Postini	0.4%
MX Logic	0.5%
MailFrontier	0.7%
Corvigo (Junk folder only)	0.7%
Trend Micro	0.8%
Tumbleweed (SpamHi folder + Spam folder)	1.2%
Cloudmark (Spam folder)	1.3%
Tumbleweed (SpamHi folder)	1.4%
Cloudmark (Spam folder + Potential Spam folder)	1.6%
Clearswift	2.3%
ActiveState (Prob folder only)	2.9%
Singlefin	2.9%
SurfControl	3.3%
ActiveState (Prob folder + Maybe folder)	7.2%
Corvigo (Junk folder + Bulk folder)	16.5%
EasyLink	20.5%
Computer Mail Services	23.4%
GFI	56.3%

NOTE: Some companies used different folders/categories for spam filtering as such their results are listed twice.



tomers report a false-positive rate between 10 and 100 times better than our tests showed. That's easy to understand, because most of the false positives we saw fell into the category of "mail that wouldn't be missed by users," such as news stories forwarded by friends, e-mail from online merchants and postings to mailing lists. For example, Postini, which had the lowest false-positive rate, missed 28 messages it marked as spam. Of those, only five were messages we wanted to see. If we hadn't been combing our mail carefully, we wouldn't have noticed those messages as missing.

Some false positives were understandable, but regrettable. A message with the subject line "IOS fw guru" looked like spam to many of the filters, but turned out to be a job offer for our test lab.

Tuning to improve performance

Most of the products can be tuned to increase sensitivity and decrease false positives, but how this tuning is accomplished and who is responsible for it makes all the difference. There are two main tools used in tuning mail filters. First is the threshold that determines whether a message is spam. The best products offered a series of levels, often expressed as percentages, based on its guess as to whether a message is spam.

Cloudmark's Authority is an excellent example of this. Each message passing through the system is assigned a number from 1 to 99, indicating Authority's confidence that the message is spam. The higher the number, the more likely a message is spam. The system manager picks actions based on these thresholds. If a message gets a 99 (very likely to be spam), then the message is dropped. At thresholds of 88 or higher, the message is probably spam, and the system manager might select to add "[SPAM]" to the subject line before passing it on to the end user.

In a corporation, using different thresholds and creating different actions for different users are important for a successful implementation. The products with the best control features were Corvigo's MailGate and Postini's Perimeter Manager. With Perimeter Manager, an end user can log on to the spam management interface at any time and adjust his settings up or down as needed. Vircom's modusGate

also supported per-user settings, but only under the control of the system administrator. Many other products included per-domain or per-group settings, all controlled by the network manager. If you can divide your users into domains, this feature might work well for you. These include MX Logic's Email Threat Management Service, ActiveState's PureMessage, SurfControl's E-Mail Filter, EasyLink's MailWatch and Tumbleweed's MMS. This doesn't mean that some of the other products couldn't do per-domain settings, it's just that we found configuring this feature with their interfaces so clumsy that it obviously wasn't part of the product design.

Black and white (lists)

A corollary to spam thresholds is the management of whitelist and blacklist membership. Of these, whitelists are the most important - the list of senders that always should be passed through and never considered spam.

Blacklists are the opposite - everything they send is considered spam. The theory is that an aggressive spam filter will show fewer false positives and higher sensitivity if it has a good whitelist from which to operate. The products we tested treated the whitelist with differing degrees of importance. Several have automatic whitelisting as an integral part of their product - the system adds an address to the user's whitelist or the company's whitelist as soon as the user sends a message.

MailFrontier's ASG does this by monitoring the logs of your outgoing mail server, while Corvigo's MailGate and GFI's MailEssentials would do this if you used their products as your outgoing mail relay.

But automatic whitelisting has its own dangers: If a user turns on a "on vacation" auto-reply message, they'll add every spammer who sends them a message to the whitelist. MailFrontier's team responded to this issue by telling us that people in corporations never use vacation messages anymore. We think that's crazy and strongly disagree.

Many products supported per-user whitelist/blacklist settings under user control, including Corvigo's MailGate, MailFrontier's ASG, MX Logic's Email Threat Management Service, Postini's Perimeter Manager and Singlefin's E-mail Protection Services. For Singlefin, this is a critical feature: Because its spam filter has no threshold tuning available, whitelists and blacklists are the only tools available to improve performance. With Singlefin's very high false-positive rate, only a comprehensive whitelist would make this perform acceptably.

All the other products supported per-system or per-group blacklist and whitelists, although sometimes the facilities for this feature were clumsy.

Not all whitelists are just for "allowed senders." Some products, including MailWatch, MMS, modusGate, Praetor, PureMessage and SurfControl, let you create a whitelist on elements such as message content. For example, a company might want to whitelist messages that have one of their product names in the body of the message, especially if these are sent to the sales team.

Three products failed to provide adequate tuning facilities that might have fixed their poor performances. Clearswift's Mailsweeper has no ability to tune spam thresholds - something either is or is not spam, yet the built-in rules caught less than half the spam it could have, with a false-positive rate eight times higher than the best products. Both GFI's MailEssentials and Computer Mail Services' Praetor used the same technique: If bad words or phrases appear even once in a message, it must be spam.

This simplistic approach to filtering spam just won't work, as demonstrated by the very high number of false positives MailEssentials and Praetor had. Both products allowed for a network manager to edit the rules to filter spam, but the inherent limitations of this technique is that the manager would be tuning forever. With so many other products that

Anti-spam performance

Vendors	Delivery rate (messages/second)
MailFrontier	20*
Cloudmark	20*
Trend Micro	20*
Tumbleweed	10
Corvigo	7.25
Clearswift	6.7
Postini	6
EasyLink	3.8
MX Logic	3.6
ActiveState	3.2
SurfControl	3
Vircom	2.6
GFI	2.4
Singlefin	1.25
Computer Mail Services	0.5

* Test setup was for 20 messages/sec.; products theoretically could deliver at higher speeds.



work better out of the box, why would you want to tune forever? GFI acknowledges the limitations of its current approach. The company is adding Bayesian filtering to Version 9 of its product, to be released later this year, which should improve overall performance. One product submitted for this test, Gordano Messaging Server, was dropped because it not only uses the same poor algorithm as MailEssentials and Praetor, but it also comes with an empty rule set, and the vendor doesn't even provide an initial guess at a word list.

What happens to spam?

The next area where products quickly differentiated themselves was in their ability to manage the spam once it was identified. All the products, except EasyLink's MailWatch, can tag spam, typically by adding a string to the subject line (such as "[SPAM]") or adding a header to the message (such as "X-Spam: yes") or both. This method is the lowest level of spam catching because it means the message still hits the corporate mail server and has to be identified and managed by the end user. Although all modern mail clients can catch and segregate messages into folders (or directly delete) when tagged, we think large enterprise networks will want to keep spam as far from their e-mail servers as possible.

Blocking the message at the anti-spam gateway means having a quarantine facility. As the filter identifies spam, it goes into the quarantine instead of the user's mailbox.

Corvigo, MailFrontier, MX Logic, Postini and Singlefin all offered a per-user quarantine, which the user can manage via a Web portal. With a few clicks, a user can see his quarantined messages, release false positives and immediately add the sender to his whitelist. This is a very scalable way to handle spam and gives the user maximum control and flexibility.

Three other products, ActiveState's PureMessage, Cloudmark's Authority and Vircom's modusGate, used a mail-based quarantine: Users get periodic notification of their spam in an e-mail, and they can click or reply to the message to release messages for delivery. This isn't as clean or desirable an approach. For example, we found things such as license keys often are marked as spam because they are short messages with lots of nonsense words in them. If you had to wait a day for the quarantine notification to show up so you could release that one urgent license key you've been waiting for, it wouldn't be a happy day.

Of course, not every company needs a per-user quarantine. If you don't care

whether messages get onto your mail servers, you can tag the spam and make users responsible for receiving and managing their own spam. Other products have per-system or per-group quarantine settings. The theory is that this lets a full-time e-mail administrator paw through the quarantine and take the burden away from end users. The problem with this is that it doesn't work - our own experiment in picking through every message showed us just how difficult it can be to decide whether a message is spam. But in some scenarios (such as small offices or environments where e-mail is tightly controlled), these approaches might work.

Different types of anti-spam

We also discovered that there are two major kinds of products in the anti-spam game today. Some are custom-built and aimed at spam filtering, such as MailFrontier's ASG, Cloudmark's Authority, Trend Micro's SPS and Corvigo's MailGate. Others come at spam filtering from the general policy enforcement, content management and mail firewall side of the house (GFI's MailEssentials, Clearswift's Mailsweeper, and Tumbleweed's MMS fit into this category).

One of the many differences in these products is in the available set of actions for messages. For example, Trend Micro's SPS has only two possible actions: Add something to the subject line, or redirect the message to some other mailbox (or both). However, if you bundle SPS with Trend Micro's older InterScan Messaging Security Suite, a popular mail content-management system, you get five more things you can do with e-mail, including a systemwide quarantine or simple deletion.

Mail content management tools are obvious places to perform spam filtering, along with other tasks. Virus scanning is the most popular, and all but five products included virus scanners or made them available as an integrated option. (Cloudmark, Computer Mail Services, Corvigo, MailFrontier and Trend Micro had no integrated virus scanning capability.)

In a couple of cases, though, adding spam protection to an existing content management tool wasn't well executed. Trend Micro's SPS is a perfect example: It's a stand-alone, simple application that is completely un-integrated into its other products. We also had a bad experience with Clearswift's add-on to Mailsweeper, its venerable policy manage-

ment system. The integration in that case was exceptionally poor. GFI's MailEssentials had similar issues - a badly thought-out anti-spam product bolted onto an older and more stable e-mail management tool.

When diving into the full-fledged content management arena, we found some products where integration of anti-spam functionality gave the network manager an incredibly rich feature set. Choosing among these products can be difficult. The first thing a network manager has to decide is whether he needs anti-spam and anti-virus capabilities (because these are so common to every company), or whether the additional tools of content management and policy enforcement would be valuable. If so, products such as Mailsweeper, MMS, Praetor and PureMessage had the greatest power and flexibility in both managing e-mail content and offering a wide variety of possible actions. For example, you could use MMS to identify incoming messages from an airline's "Internet-only specials" mailing lists and defer delivery until after business hours - unless they were being sent to your travel department.

Managing configurations

Defining your e-mail policy, whether it's anti-spam, anti-virus or content/policy means driving some sort of GUI. Most products picked either a local management application, such as a Win32 or Microsoft Management Console tool, or offered a Web GUI; SurfControl gave us both. For some products, including Cloudmark's Authority and Trend Micro's SPS on Unix, you also could simply edit the local configuration files. We were also impressed with Postini's batch command language, which is useful if you needed to update settings on hundreds or thousands of users or define many policies. Of course, the simpler the product, the simpler the GUI. This makes comparing products with many functions, such as SurfControl, to spam-filter-only products, such as Authority, a little unfair.

The highest level of flexibility came in two Sieve-based products (Vircom's modusGate and ActiveState's PureMessage) and in Computer Mail Services' Praetor, which uses Visual Basic as its scripting language. All invited the network manager to get down-and-dirty, writing in Basic or Sieve language, defining their own policies and rules for searching, classifying and identifying messages of interest. This can be important, because we found that not every vendor has a



good idea of how to abstract out the idea of e-mail policy configuration in its GUI. One of the best was Tumbleweed's MMS, which we found easy and intuitive, letting us construct rules using any of the dozens of criteria, actions and notifications MMS supported - without having to dive into a programming language. MMS had some rough edges, though: If you wanted to add a string to the end of a subject line (such as "[SPAM]"), you used one kind of rule, and if you wanted to put the same string at the beginning, you used another. Whoever thought that up wasn't thinking clearly.

Other systems brought a lot of power, but with a serious lack of flexibility. Praetor is a good example. The authors of Praetor came up with a bunch of very interesting and typical scenarios that could be used, wizard-style, to build a mail policy. But like many wizards, they also are completely inflexible. Go down the path given to you, and it's easy; but decide you want to whitelist traffic for different domains in different ways, and you're kicking and screaming all the way. It's not impossible in Praetor, but a lot harder than it needs to be.

Getting updates is another management problem worth solving. Not every product needed spam "signature" updates because of the algorithms being used, but some that did weren't designed all that well. For example, Computer Mail Services' Praetor and GFI's MailEssentials both depend on word lists to help identify spam, and neither one has automated updates. Possibly this is because the word lists shipped are such poor spam filters that any network manager would have customized them so much that an update wouldn't do any good.

Corvigo's MailGate also requires manual updates, and while Clearswift, Cloudmark and ActiveState all support automatic updates, they don't build them into the product. That's dangerous because without simple things such as cryptographic hashes applied to a virus or anti-spam update, it's easy to propagate bad or corrupted updates and start bouncing or filtering mail inappropriately. Such a lackluster and amateur design to such a critical part of the system can be a costly problem to fix.

Speed can be a problem

We ran some simple performance tests on the anti-spam gateways to see if speed was going to be a problem. In our tests, we used a traffic generator to throw a stream of 10,000 e-mail messages (at approximately 20 messages per second) at each product. We were

happy to see that some products could keep up well.

In the lab, Computer Mail Services' Praetor, which slowed down to one message every two seconds, turned in the worst performance. This might be caused by the internal architecture of Praetor, which converts your anti-spam rules into Visual Basic for actual filtering. That number gave us pause because that's only about 40,000 messages per day, well below what even a midsize company would see. (For more on our performance tests, see "How we did it".)

We also saw fairly low numbers out of ActiveState, modusGate, MailEssentials and SurfControl. In the case of SurfControl, we tested the software on slightly different hardware than everyone else: The CPU speed was faster, but the I/O subsystem was Advanced Technology Attachment-based rather than SCSI. Because a typical mail relay will see peak loads of five to 10 times the average load, it wouldn't take too much of a burst to make any of those products fall behind in the middle of the day.

On the service-based spam filters, we couldn't do the same kind of testing, but we did some testing with some statistics showing an alarming performance problem in Singlefin's spam gateway. Although Singlefin received and accepted messages at a good clip, it only could return them to us at 39% of the speed it accepted them, indicating that Singlefin has some performance bottlenecks to work around. Compare this to EasyLink, MX Logic and Postini, all of which could return messages to us as fast as we could send them.

Our performance problems with Singlefin were not just visible in message-processing speed. When logging on to manage user quarantines, we had delays of 15 to 45 seconds between screens. We also found the administrative interface to be tediously slow.

Services vs. software or appliances

Our experience during the month of testing was ambiguous as to which approach was better. In fact, we had disturbing failures on the appliance, service and software fronts during our months of testing.

The most innocuous problem was on our Corvigo appliance, which didn't recover properly from a power failure. It continued to accept messages but wouldn't deliver them. An e-mail to technical support provided an easy solution - reboot the system using the front panel. Sure enough, after a clean reboot, our messages started to appear.

Software-based gateways also had their fair share of failures. Late in the test, the Clearswift Mailsweeper software suddenly started refusing incoming messages, bouncing them all over the Internet, probably because of a time change. Rebooting didn't fix the problem, and Clearswift's 40-hour-per-week technical support staff didn't get back to us for several days, by which time the problem mysteriously went away (perhaps because the clock had caught up with itself).

On the service front, MX Logic turned up a problem during our performance testing that we hadn't noticed before: It was refusing certain kinds of messages. MX Logic fixed the problem within minutes once we convinced the help desk that the problem was on the company's side and had excellent response time for a Saturday morning. Our experience is that all of the approaches have some chance of failure.

Reporting features

Like all enterprise applications, reporting is an important check mark for an anti-spam gateway. Unfortunately, what reports are needed is something on which no two vendors agreed. Across the board, it's hard to say whether one set of reports was strong or weak. Our favorite, though, was clear: Corvigo's MailGate can provide a combination of reporting and log data through its Web interface so you easily can track any message through the system. While other products had this hidden in their mostly undocumented log files, Corvigo was the only one to make it easy to do. Clearswift's Mailsweeper came in a close second.

We also were interested in products that provided a good dashboard function: something that could instantly display system load, queue lengths, message counts and a current status of what's happening on the gateway. The best dashboard was SurfControl's GUI, which gave us incoming mail load, queue lengths, local statistics and status reporting. Corvigo's MailGate, Clearswift Mailsweeper, Cloudmark's Authority and Tumbleweed's MMS all did a good job at giving us instantaneous snapshots of how their servers were operating.

Vircom tried something unusual in its modusGate: It exported its performance and status information via Windows "perfmon" counters. This has a distinct advantage in a Windows shop: You simply can add the statistics you care about to your existing perfmon screens or run a separate copy.



What's best for me?

Although it's easy to rank products based on their ability to identify and filter out spam, your own requirements will determine which product is best for you. Your first decision has to be whether you consider individual user quarantine control and settings important.

This feature will reduce mail server load and give users control over their own spam settings and whitelists. Top-rated products that include this feature include Postini's Perimeter Manager managed service offering, Corvigo's MailGate appliance and MailFrontier's Anti-Spam Gateway

Windows-based software.

If per-user controls and quarantines are not as important as other policy-based and content-based mail filtering, Tumbleweed's MMS appliance did an excellent job of filtering spam and gave very flexible control over mail flows.

Net Results

The breakdown	Filtering performance 40%	Per-user facilities 25%	Advanced functions 20%	Spam filtering control 10%	Speed 5%	TOTAL SCORE
Postini Perimeter Manager v3.3	5	5	3	3	5	4.4
Tumbleweed Communications Messaging Management System	4.5	3	5	4	4	4.15
MailFrontier Anti-Spam Gateway	5	4	2	4	5	4.05
MX Logic Email Threat Management Service	4	4	3	3	5	3.75
Corvigo MailGate v1.5-10	4	5	2	3	3	3.7
ActiveState PureMessage v4.0	3.5	3	5	4	2	3.65
Vircom modusGate v2.15	3	4	5	3	2	3.6
Cloudmark Authority v2.0	4.5	2	3	4	5	3.55
SurfControl E-Mail Filter v4.6	3.5	2	5	4	2	3.4
Singlefin E-mail Protection Service	3.5	4	3	2	3	3.35
Computer Mail Services Praetor v1.5	2.5	2	4	4	1	2.75
Trend Micro Spam Prevention Services v1.0	3.5	1	2	3	5	2.6
Clearswift CS Mailsweeper Anti-spam Edition	2	1	4	2	3	2.2
GFI Software MailEssentials v8.0	1	1	4	4	2	1.95