

# White Paper

---

## Preventing Spam in a Windows Environment

*By Bill Boswell, CTO Windows Consulting Group*

*Bonus Section: McAfee Case Study; Closing the Door on Spam.*

**Sponsored by:**



## Preventing Spam in a Windows Environment

- Bill Boswell, Windows Consulting Group, Inc.

When the public was first introduced to the Internet back in the early 90s, unsolicited e-mail was only a minor nuisance. It was simply seen as the digital equivalent of paper junk mail, to be discarded unless found by chance to contain something useful. In the last couple of years, though, this nuisance has transformed itself into a pernicious waste of time and system resources as an unrelenting flood of offers for pills, porn and personal enhancement products pours through the Internet and onto your Exchange servers.

Analyst groups have tallied the toll of junk e-mail in the millions, even billions, of dollars in an effort to show decision-makers that it is worth expending money to stop its spread. Government bodies debate ways to punish spammers without harming productive uses of the Internet. The courts have taken a more and more conservative stand against senders of unsolicited e-mail. The backlash against spam has become so overwhelming that no responsible organization would now even think of sending even one e-mail without express consent of the recipient. Still, spam continues to erode brand identity as large numbers of e-mails contain offers for otherwise reputable products, thus linking the innocent manufacturers with slimy marketing tactics.

As IT managers and engineers, we not only have to deal with irate users tired of inboxes filled with worthless information, we need to be concerned with potentially felonious offers that our users might accept, either unwittingly or because they do not know the danger. A single click on the part of a user could open up our organizations to exploits, embarrassments, and possible litigation. For example, unsolicited e-mail advertising software at drastically reduced prices often comes from pirates who, having duped an unwary user into accepting the fraudulent offer, then multiply their profits with credit card scams.

I live in the southern latitudes, and when I think of spam, I can't help but think of fire ants. They propagate voluminously, swarm over any likely landscape, bite viciously when threatened, and make life miserable for anyone who encounters them. When I look for ways to kill fire ants in my backyard, I don't waste time stomping on one or two of them. I look for pesticides that disrupt their life cycle and keep them from spreading. A good antispam strategy does the same thing. You need to first identify how unsolicited e-mail makes its way into your system then block those paths in as many places as possible.

### Avoid Getting Listed as a Spam Target

A spam campaign begins life as a list of e-mail addresses. Your first job, then, is to educate your users not to do things that get them put onto those lists. That can happen in a variety of ways. Users who sign up for Internet access from unscrupulous ISPs often find that their e-mail addresses have become negotiable commodities. Inform your users which ISPs to avoid. Users often leave contact information at web sites that then sell their addresses to spammers. Encourage them to use alternative e-mail addresses rather than their corporate address or to obfuscate their address in a way that only a human reading it would be able to interpret, such as `tony@REMOVE_BEFORE_SENDING.company.com`.

---

Spammers love to harvest addresses from web sites. One method to avoid a great deal of spam is to obfuscate addresses you put on your public web pages by using ASCII equivalents for the characters. For example, instead of putting your clear-text address in a “mailto:” entry, such as

```
<a href="mailto:Sales@Company.com">
```

you would use the ASCII numbers that correspond to the letters, which looks like this

```
<a href="mailto:sales@company.com">
```

This doesn’t necessarily guarantee spam avoidance – some search bots look for translated addresses – but it sure does help reduce the volume of spam. For a quick look at the ASCII equivalent of your e-mail address, visit [www.wbwip.com/wbw/emailencoder.html](http://www.wbwip.com/wbw/emailencoder.html).

Frankly, though, even if you and your users are diligent about avoiding inappropriate uses of e-mail addresses and you obfuscate your addresses wherever they appear on your web sites, you delay the inevitable. I’m convinced that if I were to leave an e-mail address scrawled on a napkin in a restaurant, the account would get unsolicited e-mail within hours.

## Avoid Getting Validated as a Spam Target

In true entrepreneurial fashion, vendors aggregate the addresses they harvest into huge lists that get sold through various licit and illicit channels. Valid addresses have more value, so spammers use a variety of schemes to verify that an e-mail address represents a person.

As any experienced administrator knows, the fastest way to validate your address is to click the “Do Not Send Mail To Me Again” hyperlink in a message. Other hyperlinks can also be buried in the message, so you should teach your users to avoid clicking on anything in an unsolicited message.

Another validation scheme exploits e-mail clients that accept messages with HTML content. Embedded in the message is a hyperlink to a graphics file on the spammer’s web site. Here’s an example:

```

```

When the user opens the message or views it in the preview pane, the e-mail client connects to the target specified in the hyperlink to load the image. When used in conjunction with unsolicited e-mail, these links act as beacons that notify the spammer when a live recipient has gotten the message. Outlook 2003 and Exchange 2003 Outlook Web Access (OWA) block beacons unless the user explicitly chooses to load the image or to trust the sender.

This beacon blocking has its disadvantages. Some popular message tracking systems operate by embedding a hyperlink in your outbound messages. This hyperlink points at a graphic on the message tracking system’s web site. When the recipient opens the message, the e-mail client reads the hyperlink and downloads the graphic. This automatically sends a notification to the sender that the message has been opened. These tracking systems will not work if the client blocks access to the hyperlink target.

## Avoid open SMTP Relays and Proxies

Once a spammer has loaded a server with hundreds of thousands of e-mail addresses, the server now needs a way to send those messages without interference. This requires either a spam-friendly ISP (some countries have national service providers notorious for turning a blind eye to spam) or a launching point that has not been previously identified as a source of bulk e-mail. Such a launching point would be willing to accept incoming messages in the form of anonymous Simple Mail Transfer Protocol (SMTP) transactions and to forward those messages to their ultimate destination. This can happen if you or a colleague inadvertently set up a public-facing SMTP server configured as an *open relay*. Figure 1 shows how an open relay works.

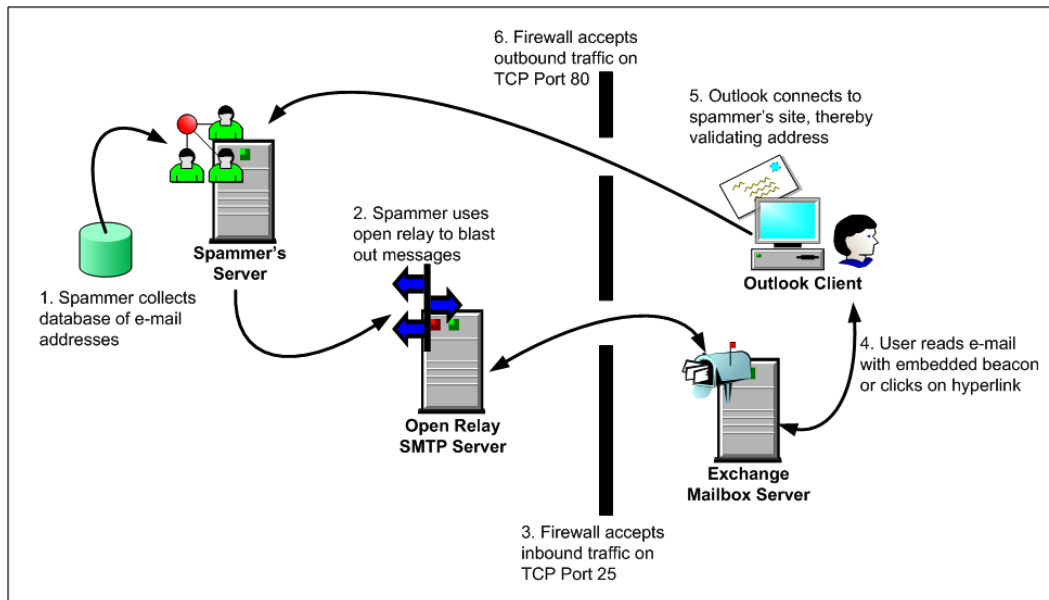


Figure 1 – Example message flow through open SMTP relay

An open relay does not require the Exchange service to be loaded. The Internet Information Services (IIS) suite on all modern Windows platforms, including desktops running Windows 2000 Professional and XP, has an SMTP service. To prevent relaying, the default restrictions of the SMTP service only permits relaying for computers in the same organization, as shown in Figure 2. In spite of the default setting, all it takes is one administrator, intent on troubleshooting a problem, to click the radio button title “All Except The List Below” and you have yourself a problem.

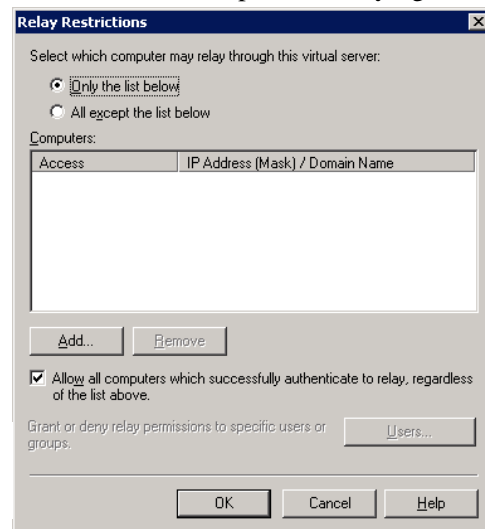


Figure 2: Default Relay Restrictions settings for Windows SMTP

Thankfully, open SMTP relays have gotten rare thanks to the diligence of administrators who configure SMTP servers. It is now proxy servers that constitute the majority of illicit SMTP message sources. This is partly due to the many small businesses and individual users who install an inexpensive (or free) proxy service application to simplify connecting several PCs to a broadband Internet line. Proxy applications are simple to configure when compared to firewalls, which often leads to problems when the person doing the installation doesn't know their power.

For example, consider a medical office where the doctor's 15-year-old nephew installs a proxy server on a Windows 98 desktop connected to a new DSL box because the proxy vendor's web site claims that the product "speeds up the Internet." The nephew misconfigures the software, puts the proxy on the public interface, and a spammer discovers it. By burying SMTP transactions inside of HTTP commands, the spammer connects through the proxy to an SMTP server and begins funneling messages to the Internet.

Open proxies have also become extraordinarily common thanks to the proliferation of viruses that either include a proxy service in the payload; or, like the SoBig virus, download a proxy service following the initial infection.

## Are You An Unknowing Open Relay or Proxy?

Spammers find open relays and open proxies by scanning for machines that respond to common ports then testing those ports using SMTP relay requests and proxy HTTP CONNECT requests. You can do these same kinds of tests yourself on the public interface of your servers. *Note:* Always get written permission from management before doing security probes. You might find yourself afoul of a corporate policy and land in serious trouble.

From the public side of your network, connect to your public Exchange server using telnet with the following syntax:

```
telnet <exchange_server_name> 25
```

The server should respond with an SMTP banner, something like this:

```
220 EX1 Microsoft ESMTPL MAIL Service, Version: 6.0.3790.0 ready at Wed, 30 Oct 2003 14:32:31 -0700
```

Now get a list of SMTP capabilities by entering [EHLO](#) (Extended Hello). When the server replies, use three SMTP commands to send a message -- MAIL FROM, RCPT TO, and DATA -- with the following syntax:

```
mail from: totallybogus@fabricatedaddress.biz
250 2.1.0 totallybogus@fabricatedaddress.biz...Sender OK
rcpt to: someusername@yahoo.com
250 2.1.5 someusername@yahoo.com
data
354 Start mail input; end with <CRLF>.<CRLF>
Subject: Your assistance most graciously and desperately needed
Let me introduce myself. I am the grandson of the Duke of Earl and I need your help.
.
250 2.6.0 <EX-S1HM3S01bpH71y00000008@ex1.actualsmtpdomain.com> Queued mail for delivery
quit
221 2.0.0 ex1.actualsmtpdomain.com Service closing transmission channel
Connection to host lost.
```

Notice that the MAIL FROM address does not need to exist. This permits a spammer to hide his real identity. In practice, spammers load lots and lots of fake information into an SMTP message.

The RCPT TO command defines the recipient and, of course, must be a valid e-mail address. If the server accepts a RCPT TO address outside of its home SMTP domain from an anonymous connection, you have found an open relay. If you examine the header of the message when you receive it, you will only see the relay server's identity and IP address. You have not revealed your name, your workstation's name, or your IP address. You can see why open relays are highly desirable objects to spammers.

Testing for an open proxy is a little more difficult and requires special tools. One such tool is Proxy Analyzer from G-Lock Software, [www.glocksoft.com](http://www.glocksoft.com). Point Proxy Analyzer at a machine and tell it which port or ports might be compromised. Then, let it see if it can connect to the port as a proxy. The tool has a scoring system to determine how ripe a machine is for exploits.

## Report Abuse by Spammers

If you or your users receive particularly egregious spam or barrages of pornographic messages, you can complain to the system administrator in the originating SMTP domain on the assumption that the administrator didn't know that the system was being abused. Don't ever render your complaint directly. The system administrator might be in league with the spammer and your complaint could invite a flood of additional spam or other, more sinister abuse.

Instead, forward your complaints through Abuse.net at [www.abuse.net](http://www.abuse.net). They take your complaint, remove any identifying marks, and forward it to the administrator of the originating domain if the name is listed in their database. You must register for this free service. If you abuse the service, or if you make vague or unsubstantiated claims or threats in the messages you want them to forward, they can cancel your registration.

## Use Real-time Block Lists

You can go a long way towards stopping the receipt of mail from open relays, proxies and SMTP servers that sit behind dialup connections by enlisting the aid of one or more service bureaus dedicated to keeping tabs on the invalid use of SMTP servers. These service bureaus go by the name of RBLs. The acronym expansion varies. Microsoft uses Real-time Block List. Other expansions include Real-time Blackhole List and Real-time Boycott List. Despite the differences in names, all RBL providers have a similar intent when it comes to inappropriate SMTP use: search it out, identify the source, and inform the public. They do not "block spam" in the traditional sense of providing an active filter. This would open them up to litigation. Instead, they compile passive lists then offer the content of the lists to you for you to do the filtering. Neither you nor the RBL provider take overt action against the spammer's servers.

RBL providers generally fall into two categories: fee and free. The Mail Abuse Prevention System (MAPS) at [www.mail-abuse.org](http://www.mail-abuse.org) is probably the most widely know of the fee-based providers. Examples of free providers include the Distributed Server Boycott List (DSBL) at [www.dsbl.org](http://www.dsbl.org), the Open Relay Database (ORDB) at [www.ordb.org](http://www.ordb.org), SpamCop at [www.spamcop.net](http://www.spamcop.net), and Not Just Another Bogus List (NJABL) at [www.njabl.org](http://www.njabl.org). Some service providers also include domains that have been reported to originate spam. One example is Spamhaus, [www.spamhaus.org](http://www.spamhaus.org). For a complete list of RBL providers, visit the Declude website at [www.decluce.com/junkmail/support/ip4r.htm](http://www.decluce.com/junkmail/support/ip4r.htm). The list includes a brief assessment of each RBL's strengths and weaknesses.

Several providers give their clients a block list that can be loaded into a filter at the client's own e-mail server or into a personal e-mail client. Filtering locally helps speed up traffic through your SMTP server, but the lists get very long and you must keep them updated regularly. In general, the RBL provider wants you to sign a fair-use agreement stating that the list is for your exclusive use and you will not advertise the contents. This is for the RBL providers' protection and yours. Spammers love to litigate and non-spammers who find themselves on an RBL might decide to sue, as well.

Spammers often operate from "friendly" Internet providers and countries that have little or no incentive to stop spammers from operating. You can get a list of these SMTP domains and countries from [www.spamsites.org/live\\_sites.html](http://www.spamsites.org/live_sites.html). Figure 3 shows an example.

Many RBL providers also host online lookup services that use reverse DNS queries so you do not need to open any special ports on your firewalls. This is called an RDNS service.

To use an RDNS service, send the RBL provider a DNS request that contains the reverse IP address of the suspected server followed by the RDNS domain name of the RBL service provider. For example, the RDNS domain name for ORDB is `relays.ordb.org`.

An RDNS service identifies spammers by returning an A record with a 127.0.x IP address, where the final x indicates the banned behavior type.

Each provider uses slightly different codes. Here are the NJABL RDNS codes as an example:

- 127.0.0.2 - open relays
- 127.0.0.3 - dial-up/dynamic IP ranges
- 127.0.0.4 - Spam Sources
- 127.0.0.5 - Multi-stage open relays
- 127.0.0.8 - Systems with insecure CGI scripts that turn them into open relays
- 127.0.0.9 - Open proxy servers

You do not need special software for sending an RDNS query to an RBL provider. Just submit the query using Nslookup and view the result. For example, if you get a message you suspect to be spam from a server with the IP address 1.2.3.4, here's the syntax to send an RDNS query to the NJABL provider:

```
C:\>nslookup 4.3.2.1.dnsbl.njabl.org
```

```
Non-authoritative answer:  
Name:      4.3.2.1.dnsbl.njabl.org
```

IP address	Days listed	Trace	Web site	DNS/MX
68.80.30.85				
209.228.4.161	New!	Trace	<a href="http://www.bulk-e-mail-marketing.com">www.bulk-e-mail-marketing.com</a>	REGISTER.COM
69.56.129.222	New!	Trace	<a href="http://www.bulkersclub.com">www.bulkersclub.com</a>	DNS4FREE.ORG
219.93.225.195	New!	Trace	<a href="http://www.a1-bulkemailhost.com">www.a1-bulkemailhost.com</a>	SPAMHAUS.US
208.187.162.4	New!	Trace	<a href="http://www.ravenswoodinc.com/isp.htm">www.ravenswoodinc.com/isp.htm</a>	LANSSET.COM
208.187.165.51	New!	Trace		
216.180.254.177	New!	Trace	<a href="http://www.netpropaganda.com">www.netpropaganda.com</a>	NETPROPAGANDA.COM
64.202.162.37 (plus 2 others)	New!	Trace	<a href="http://www.covat.us">www.covat.us</a>	SECURESERVER.NET
66.227.39.161	New!	Trace	<a href="http://www.abiajwood.com">www.abiajwood.com</a>	CHEAPWEB4U.COM
202.54.194.209	New!	Trace	<a href="http://www.send-safe.com">www.send-safe.com</a>	INTRAGROUP.NET, OGBUS.COM
217.107.162.241	New!	Trace		
66.179.174.44	New!	Trace	<a href="http://www.kash.ca">www.kash.ca</a>	KASH.CA
69.56.134.174	New!	Trace	<a href="http://www.webextractor.com">www.webextractor.com</a>	CHANGEIP.COM

Figure 3: Example spammer IP address and domain listing from SpamHaus

# White Paper

Address: 127.0.0.9

If you don't want to memorize return codes, use Nslookup in interactive mode to get a TXT record corresponding to the A record, if one is available. Here's the syntax:

```
C:\>nslookup
> set type=txt
> 4.3.2.1.dnsbl.njabl.org
```

```
Non-authoritative answer:
4.3.2.1.dnsbl.njabl.org text =

      "open proxy -- 1058170803"
```

If you want a web-based solution to do RBL lookups, visit the Open Relay Database lookup page at [www.ordb.org/lookup](http://www.ordb.org/lookup). This site allows you to enter an IP address to see if it has been tagged as an open relay, and also has the capability of querying other major SMTP abuse sites to see if the IP address shows up as a known spammer, an open proxy, or some other obnoxious use. Figure 4 shows an example report.

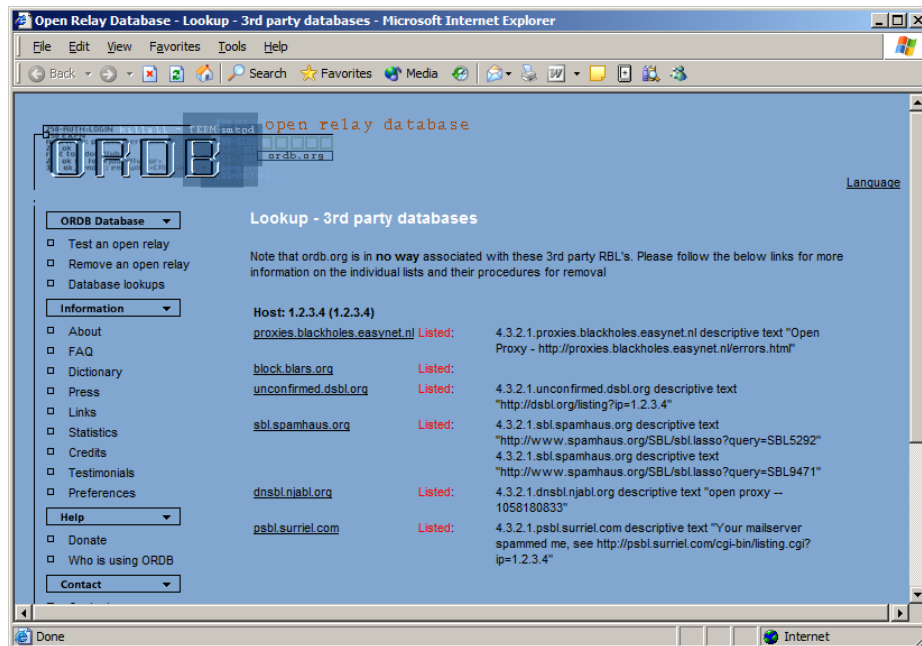


Figure 4: Sample third-party mail abuse listing from ordb.org

## Connection Filters

Exchange Server 2003 has the native ability to use RBL service providers. The Global Settings for an organization has a Message Delivery object with properties that include a Connection Filtering option, as shown in Figure 5. You can define one or more RBL providers and the Exchange server will query each one to see if the SMTP server or domain has been identified with a spammer. You can also insert your own entries in the Global Accept And Deny List configuration and you can stipulate exceptions in an exception list, useful when one of your vendors or customers has been listed by an RBL but you still need to do business with them while they clean up their act.

Once a global filter has been configured, each SMTP virtual server in the organization must be configured to use the filter. This is done for each IP address on each virtual server by making an entry in the Advanced Properties of the virtual server's IP settings.

Exchange 2000 and Exchange 5.5 do not include RBL connection filtering, so you'll need a third-party tool. Just about all of the antispam suites for Exchange include an RBL component along with the other filter mechanisms. If you only want to do RBL filtering, two popular free products are:

ORFilter

([www.martijnjongen.com/eng/orfilter](http://www.martijnjongen.com/eng/orfilter))

GFI MailEssentials

([www.gfisoftware.com/mes](http://www.gfisoftware.com/mes))

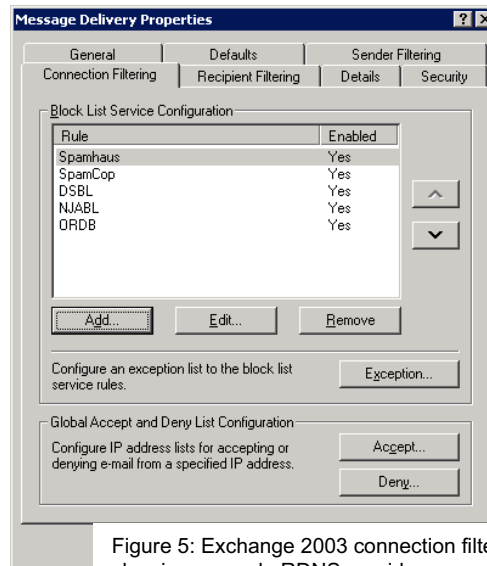


Figure 5: Exchange 2003 connection filter showing example RDNS providers

## RBL Limitations

RBL filtering, in and of itself, can't protect you completely from spam. First off, not all spam comes from open relay and proxies. Second, spammers who use open relays and proxies are getting sophisticated about skipping around from one exploited machine to another so that no one machine sends sufficient junk mail to get the attention of the service bureaus. This is especially prevalent now that so many worms have deposited proxy servers on unprotected home machines. RBL filtering can also result in many false positives if desirable messages come from blocked servers or SMTP domains.

RBLs can also become a sticky web from which to extract yourself should you or one of your colleagues inadvertently set up an open relay or an open proxy on the public side of your production network or a public-facing machine gets hacked or infected.

If you get listed by an RBL, messages from you get bounced by subscribers to the RBL service. Depending on the sender's filtering configuration, you might not be informed of the block. This avoids large-scale disruption of the Internet during worm attacks. It also leaves your users wondering why no one is answering their e-mails.

Once you correct the condition that got you listed on an RBL, contact the provider to get your system removed from the list. RBL providers do not talk to each other in any coordinated fashion, so you may have to communicate with more than one of them. It can take as long as a week, maybe even longer, to clear your name from an RBL database. Don't bother trying to muscle the RBL provider into working faster. They deal with hardcore spammers every hour of every day, something that takes a Jabba the Hutt mentality. They aren't afraid of threats from you or your CEO or your lawyer. If you comply with standard SMTP practices, you will eventually get removed from their system.

## Challenge-Response Blocking

Since RBLs have their fallibilities in discriminating between real human beings and spammers, the next logical step is to make the determination yourself. In his novel *Dune*, Frank Herbert described a rite designed to determine whether or not a person was a human being. The rite involved the use of a neural pain inducer and a poisoned dart called a *gom jabbar*, the theory being that a human being could control the flight from pain if faced with the greater danger of imminent death.

Several antispam solutions use the gom jabbar approach. "Prove you're human," the solution says, "by doing something no machine can do." Examples of vendors who use this approach include SpamArrest, [www.spamarrest.com](http://www.spamarrest.com), and SpamLion, [www.spamlion.com](http://www.spamlion.com).

Both of these solutions involve redirecting a user's incoming and outgoing e-mail to the vendor's e-mail servers. When the vendor gets an inbound message on behalf of a client, the message server checks its database to determine if the sender has been validated. If not, the system holds the message in a queue and returns a response asking for the sender to take some action. SpamLion asks that the sender click a hyperlink, similar to the registration verifications done by Internet mailing list managers such as Majordomo. SpamArrest requires the sender to read a bitmap containing a word then type that word in the reply.

In either case, the onus for filtering messages goes on the sender, not the recipient, unless the sender wants to take the time to go through a list of unverified messages to pre-approve the sender. If a sender decides it's too much trouble to validate, the message eventually gets deleted from the vendor's servers.

Challenge-response systems have enviable filter performance characteristics but they aren't always the perfect solution, especially in situations where a non-human sends a desirable message. Opt-in mailing lists are a prime example, as are vendor bulletins sent to inform customers of vital product information. Very few list managers take the trouble to reply to challenge-response messages, with the result that potentially useful information never makes it to the recipient. If the recipient isn't willing to take the time to pre-approve a mailing list by domain name, then the list's mailings go right to the junk pile.

Another problem with the challenge-response approach is that the customers who use these systems depend on the quality of the vendor's messaging infrastructure. If the vendor's system should go down, or undergo an unforeseen load that makes it unreachable, or get brought down by a globally-based Distributed-Denial-of-Service (DDoS) attack, then all mail to and from the customer grinds to a halt.

That being said, it's tough to argue with 100 percent filter performance, so you should include at least one challenge-response option on your short list of solutions when doing your evaluations.

## Rules-Based Filtering

Between the wholesale filtering of RBLs and the highly targeted filtering of challenge-response applications lies a wide spectrum of antispam products that aim to block unsolicited e-mail without whacking useful or desirable messages in the process. Unlike RBLs or the challenge-response products, these products examine the content of messages then use a combination of keywords, signatures, rules, and content analysis to determine their “spamminess.”

Rule-based analysis is the hallmark of virus scanners, so it’s no surprise to find many of the antivirus vendors coming to market with products that mimic the operation of their virus scanning engines. If a vendor updates the rulesets often, they can be tremendously effective. For example, if a spammer begins an extensive campaign using a particular style of offer and format of message, the vendor can add the signature, send out an update, and block a great deal of the spam.

Rules-based filtering has a problem, though, due to the highly variable nature of spam content. For example, you’ve seen a metamorphosis of spam subject lines over the last year to include nonsense in response to rule-based scanners that look for consistent subject content for a particular spam. A rules-based antispam engine keyed to look for the phrase “You Can Last Twice As Long” might pass over “You ,Can La\$t Twice /As Long mxplxct”. Spammers also place extra words in the body of the message to fool a rules engine, then play games with the font colors so that the words don’t appear when the user reads the message.

## Bayesian Filtering

To find a way to counter the slippery content of spam messages, many antispam vendors employ a decision process based on the work of an 18<sup>th</sup> century mathematician and minister named Thomas Bayes. In a document titled “Essay Towards Solving a Problem in the Doctrine of Chances,” Rev. Bayes proposed that it’s possible to infer the probability that an event will occur based on the number of times it has occurred in the past.

This type of decision process mimics, in some ways, the way people learn. For example, let’s say you know nothing about fast cars and you take a motor trip on a long, lonely stretch of good old American blacktop. A red car speeds past you. You pass a green and a blue car. Another red car whizzes by. And another. Pretty soon, if this happens enough, you start to develop a belief that all red cars are fast, or at least a lot faster than you. If enough red cars zoom past, even if you encounter one or two slow red cars, you may still be convinced that at least *most* red cars are pretty darned fast.

A spam filter that uses Bayesian decision processes must be taught the difference between good e-mail and spam, and you or your users must do the training. After a while, if it has been trained correctly, the Bayesian filter can recognize a good e-mail with content such as, “I desperately need a good headache cure. Do you have one?” from a bad e-mail saying, “Need a sure cure for your headaches? We have one.”

Bayesian filters break the contents of a message into tokens then analyze the frequency of the token usage compared to known spam examples. This makes them particularly adept at identifying messages containing word games. A sentence such as “Kure ba l dne\$\$ with hare im-plant\$” might would fool a rules-based filter but a Bayesian filter would know that tokens with odd characters nearly always indicate spam and rate the message accordingly.

## SpamAssassin and Its Descendants

Investigations done a few years back initially determined that naïve Bayesian filters required too much training to be useful against spam. This attitude changed thanks to innovative ideas from a developer named Paul Graham in papers titled “Plan for Spam” and “Better Bayesian Filtering,” available at [www.paulgraham.com](http://www.paulgraham.com), and the introduction of a wildly popular Bayesian-based open source, Perl-based spam filter called SpamAssassin.

Following on the tremendous success of the open source version of SpamAssassin, its creator and members of the development team chose to contribute to commercial products based on the core operation of SpamAssassin.

Network Associates ([www.nai.com](http://www.nai.com)) acquired a Windows implementation of SpamAssassin by purchasing a company called Deersoft and now use the code as the basis of their SpamKiller product line. MessageLabs ([www.messagelabs.com](http://www.messagelabs.com)), a service bureau that specializes in message content filtering, incorporates SpamAssassin code plus extensions written by a SpamAssassin developer into filters used to analyze and divert unsolicited messages. Stata Labs ([www.statalabs.com](http://www.statalabs.com)) has a product called SPro that uses SpamAssassin extensions written by members of the development team.

Other companies have shown that Bayesian filtering is a superior technology for handling spam. A company called Audiotrieve ([www.audiotrieve.com](http://www.audiotrieve.com)) uses Bayesian filters based on their extensive experience with language handling to produce an Outlook plug-in called InBoxer. There are many other examples of antispam products that incorporate Bayesian filters along with their own proprietary heuristics. (The word “heuristics” means “I won’t tell you how this works, but as long as you’re satisfied with the output, then it’s a success.”)

## Limitations of Filter Technology

As you evaluate the filter technology used by a particular product, whether rules-based or Bayesian, take a good look at the features the vendor plans on implementing over the next year to 18 months to get an idea whether the vendor recognizes the sophisticated nature of spammers.

Carefully crafted messages won’t trigger a negative response from a word analyzer. Filters can be fooled by spam messages that use a mixture of ASCII and URL translations. Spammers increasingly avoid using words entirely, choosing instead to embed graphics that can’t be subjected to Bayesian analysis. To respond, the next generation of spam scanners must learn to deal with graphics either by looking for words or offensive material within the bitmaps themselves or by learning the patterns of graphics within the message.

The real test of any antispam product goes well beyond its filtering technology. It must integrate well with your current e-mail platforms. It should either integrate fully with your current antivirus technology; or, at the very least, the two products must play well together. The client interface must require minimal training for your users. The configuration must exhibit flexibility in the face of your own unique message traffic patterns. And like the ancient Roman doctor once said, above all it must do no harm. False positives must occur infrequently and users must be able to recover blocked messages quickly and simply.

## Edge Filters

Some antispam solutions are designed to work outside the boundaries of your private network, either on a host in your DMZ or at a service bureau. The solution monitors SMTP traffic to and from your organization and diverts or tags suspicious messages.

These “edge” filters have several advantages. They can be configured to block obvious spam and thus keep the messages out of the message stores on the private side of the network. The savings in storage alone will probably pay for the software or service bureau fees. Because the edge filter sees all traffic, it learns about spam patterns more quickly. And if you purchase a product that depends on rules sets or signatures that must be downloaded from the provider, having one or two edge filters simplifies distribution of the support files.

You do not need to run an Exchange server in your DMZ to take advantage of edge filtering. Some solutions rely on an appliance you plug in at your DMZ. Once installed, you do a bit of configuration at the appliance, a bit more at your firewall, then point your Exchange servers at the appliance using the Smart Host feature in the SMTP connector and start watching the filtered mail for false positives. An example of an appliance is the Message Management System from Tumbleweed Communications ([www.tumbleweed.com](http://www.tumbleweed.com)).

If you prefer not to learn yet another set of operating instructions for an appliance, you can set up a standalone server in your DMZ running Windows 2000 or Windows Server 2003 and install an antispam solution that integrates with the SMTP service in IIS. Products in this category either store blocked messages or simply tag messages with a spam score and let the Exchange servers or Outlook clients filter for the scoring attribute. Exchange 2003 comes with a set of custom event sinks that antispam products can use for tag filtering. Outlook 2003 has filters already in place to triage messages based on tags. Third-party products are available that perform these same actions for earlier versions of Exchange and Outlook.

## Store Filters

If your budget doesn't have room for additional servers/appliances, tiered firewalls, or filtering service bureaus, then your public-facing Exchange server probably sits behind a moderately priced firewall that does port forwarding for incoming SMTP traffic. In this configuration, you could use a store-based antispam solution. Most the major antivirus vendors also have an antispam product, so that would naturally be the first place you should start your evaluation.

The primary objection to store-based antispam and antivirus products is the possibility that they could make your Exchange server unstable. Vendors work closely with Microsoft to attempt to develop products that are 100 percent compatible with Exchange, but every engineering effort represents a compromise and antispam/antivirus solutions are no different. You stand the best chance of combining stability and performance by using products from the same vendor.

Still, if another vendor's product gives you better performance or filtering capabilities, then by all means use it once you've done sufficient testing to ensure no conflicts. The last thing you want is a tussle between an antispam and an antivirus application over a message that looks like spam but has a virus payload. Pay particular attention to the quarantine strategies used by the antispam and antivirus applications. Review pertinent Knowledgebase articles in TechNet that reference the vendor's product or design strategies and use Google to search newsgroups for potential conflicts.

## Client Filters

Some users object to edge filters and store filters because of concerns over false positives. Other users simply don't want to give up control over their e-mail. "Don't mess with my messages," they say, "Just make it easier for me to weed through the junk."

Even if you decide to implement edge or store filters, you might want to think about deploying client filters, as well. Users often supplement their corporate e-mail with POP mail obtained from outside mail servers. By deploying client-based filters, you provide a service to these users. You can combine antispam filtering with your antivirus desktop solution, with the same caveats for compatibility and stability as for store-based filters.

Also, if your edge or store filter simply tags messages as potential spam, then you'll either need a client-side application that handles the tagged messages or you'll need to deploy your Outlook clients with filter rules that deal with the tagged messages.

When evaluating client-side filters, check to see where blocked mail gets stored. If a client-side filter simply moves messages from the user's inbox to another server-side folder, then you've done a service to your users but not to yourself. If possible, find a client-side filter that will either do a hard delete on the junk mail after a short period of time or will shunt the junk mail into an alternate repository that resides on the local desktop.

Small shops with no Exchange servers can use the junk mail filtering capabilities of Outlook. Outlook 2000 and later have canned rules for recognizing junk mail and pornography. Outlook 2003 exposes this ruleset with a Junk Mail button in the e-mail configuration tab in the Option menu. You can pre-populate the filter with lists obtained from the Internet. For example, the GazNET site ([www.gaznet.com](http://www.gaznet.com)) has a downloadable list that can be placed directly in an Outlook profile.

## Final Thoughts

In the final analysis, the decision to limit or eliminate spam from entering your network should not be "Should I do it?" but "When will I do it?" and "How much do I want to spend?" It's galling to shell out money for spam prevention because those dollars come from other projects that actually further your organization's mission, but if trends continue, the current volume of unsolicited messages will rise from its current volume of 50 percent of all Internet message traffic to a whopping 90 percent in just a couple of years. Don't wait until your Exchange server stores burst at the seams before taking action.

# Closing the Door on SPAM

## At a Glance

### Industry

Global IT security products and services

### The Network Environment

Nearly 4,000 users worldwide  
Windows 2000 Exchange server

### The Challenge

Protecting the organization from unsolicited junk e-mail as well as associated virus risks.

Finding a solution that would scale effectively, be customizable, and work across multiple domains with fault tolerance and load-sharing capabilities.

### The Solution

McAfee WebShield e500 and e1000 appliances with SpamKiller for WebShield appliances.

### The Benefits

Fault tolerant, load shared e-mail spam and virus protection.

Easy-to-use interface for management and reporting.

### The Customer

The client is a global IT security company that provides solutions and services designed to secure, protect, and manage IT infrastructure.

## Introduction

Spam, now known globally as one of the biggest threats to the Internet and source of untold annoyance to end users and IT managers alike, is on the rise. The number of spam e-mails flooding into corporate networks is a serious problem. Users spend more and more time deleting unwanted e-mail from inboxes and run the risk of inadvertently being exposed to inappropriate content contained in spam messages, which can have a serious legal liability impact to corporations.

The problems of dealing with spam effectively are many and complex—from recognizing what is and what is not unsolicited commercial e-mail, to the fact that spam is extremely difficult to trace back to its source, as spammers invariably take steps to hide their origins and identities.

When a leading IT security organization wanted to protect itself from spam and associated problems, it needed a solution that would be robust, scalable, and fault tolerant. As a global organization with nearly 4,000 users, e-mail is a critical business tool and downtime is not an option. To solve the problem, the company turned to McAfee® Security and its range of WebShield® appliances and SpamKiller™ products.

*“SpamKiller for WebShield is currently identifying over 1.7M unsolicited e-mails per month. In the future, we estimate we will be blocking another 400,000 e-mails, which will be 2.1M unsolicited e-mails that won’t reach our users. That has to be some kind of record.”*

*The IT Security Director at a  
Global IT Security Organization*

## The Network Environment

Over twenty-four Microsoft® Exchange 2000 servers deployed globally serving in excess of 8,000 mailboxes across two separate exchange organizations

Server operating systems deployed are Microsoft NT 4.0, Microsoft Windows® 2000, Sun Solaris Ver 8+, and IBM AIX.

Client operating systems are Windows XP Pro, 2000, and NT 4.0

Wide-area network infrastructure utilizes frame-relay and ATM protocols.

## The Customer’s Challenge

As one of the world’s leading IT security companies, success can bring its own problems; open to all sorts of attacks from the outside world and, like any other connected business today, attracts inordinate amounts of unsolicited e-mails.

The company in question needed to conduct business online and protect its users from risks; and with a large and varied user base, the technology solution had to be scalable and easily customized.

The IT Security Director describes the problem: “We receive spam and viruses on a hourly basis and are engaged in a constant battle with hackers and virus writers. Since we have multiple Microsoft Exchange servers with thousands of people connected to them, the amount of traffic that we send and receive can vary enormously and with e-mail being an essential tool we needed a solution that would be fault tolerant at an enterprise level.

“The key requirement for us was to have a robust Internet gateway solution. Rather than scanning messages for spam or viruses at the e-mail server level we wanted scanning for spam and viruses to be on the perimeter of our environment, thereby dealing with the problem before delivery to our exchange or mailbox servers. An implementation at the gateway would ensure that spam messages would not enter the network or reach end-users, significantly reducing network storage and legal liability issues and increase user productivity. With the number of users we have, dealing with any potential threats and unsolicited e-mails at the inbox would be a never ending task, hence the need to close them down at the ‘front door’ of the network.”

#### **The Network Associates Solution—Business Values and Benefits**

The spam problem stems from the fact that it costs very little money to actually send e-mail in bulk, and the technology available today means that vast numbers of e-mails can be sent very easily. This makes the spam problem a numbers game. With the number of unsolicited e-mails sent to people increasing every day, the risks associated with those e-mails also increase.

With the key requirements being a configurable gateway product, the McAfee solution consists of a combination of its WebShield e500 and e1000 appliances with SpamKiller for WebShield appliances. The WebShield appliances combine hardware and award-winning McAfee anti-virus software in integrated appliances, offering easily configurable devices that plug into virtually any existing network and can scan SMTP, HTTP, FTP, and POP3 traffic for hostile code or malicious content. In addition, the WebShield appliances can filter messages by content, blocking specific file names or message subjects preventing users from receiving unsolicited messages.

The IT Security Director explains the solution: “At our Internet gateways we have SMTP relays, which receive and deliver all of our internal and external messages. We use two relay boxes, with fail over capability, so that if we have an issue with one relay we can use the inbound and outbound capabilities of the other. We have implemented a realtime blackhole list at the relay level to further reduce instances of spam breaking through. After a daily zone transfer to our DNS server, our main relay box looks at the DNS server and checks to see if a message comes from an address that is on the blacklist. If it is, it is blocked.

“These relays deliver to our WebShield appliances which have been installed in transparent bridging mode. Transparent bridging allows the appliances to be installed as network bridges and as such require minimal network changes.”

McAfee SpamKiller for WebShield appliances component allows a great degree of flexibility and configuration, and is based on the award-winning SpamKiller application. SpamKiller is powered by the McAfee SpamAssassin™ engine and scoring system. The SpamAssassin scoring system is based on an extensive rule set, which determines if a particular e-mail message is spam or not. The rules are run against every e-mail and each

rule is associated with a score, positive or negative. Rules with negative scores indicate legitimate e-mail and rules with positive scores indicate unsolicited e-mail. When added together, these scores give each e-mail an “overall spam rating.”

A WebShield appliances solution with SpamKiller for WebShield provides five levels of protection for spam—by utilizing integrity analysis, heuristic detection, content filtering, administrator-defined Black and Whitelists, and support for third party RBLs—to provide up to 95 percent effective “out of the box” protection.

The IT Security Director further stated “As part of the process to determine ‘what is spam’ for our organization we currently are not blocking or quarantining messages, simply labeling them as spam and delivering them. We don’t tag anything as spam until it receives a score of five points, in the future we will continue this process as well as blocking anything with a ten-plus score. The WebShield appliances with SpamKiller connect to our Microsoft Exchange servers and they complete the system. From start to finish, it has been a painless rollout; we installed the WebShield appliances and SpamKiller for WebShield software over two days and went through a pilot of twenty users for two weeks. The whole system was installed, tested and implemented for nearly 4,000 users in less than a month.

“SpamKiller for WebShield is currently identifying over 1.7M unsolicited e-mails per month. In the future, we estimate we will be blocking another 400,000 e-mails, which will be 2.1M unsolicited e-mails that won’t reach our users. That has to be some kind of record.”

**McAfee Security** 3965 Freedom Circle, Santa Clara, CA 95054, 408.988.3832 main, [www.mcafeesecurity.com](http://www.mcafeesecurity.com)

Network Associates® products denote years of experience, and commitment to customer satisfaction. The PrimeSupport® team of responsive, highly skilled support technicians provides tailored solutions, delivering detailed technical assistance in managing the success of mission critical projects—all with service levels to meet the needs of every customer organization. McAfee Research, a world leader in information systems and security, continues to spearhead innovation in the development and refinement of all our technologies.

Network Associates, [LIST OF PRODUCTS] are registered trademarks or trademarks of Network Associates, Inc. and/or its affiliates in the US and/or other countries. Sniffer® brand products are made only by Network Associates, Inc. All other registered and unregistered trademarks herein are the sole property of their respective owners. © 2003 Networks Associates Technology, Inc. All Rights Reserved.

For more information on products, worldwide services, and support, contact your authorized Network Associates sales representative.