

CHAPTER 5

BLOCKING SPAMMERS WITH DNS BLACKLISTS



In Chapter 4 we introduced you to DNS Blacklists as one of several means for fighting spam. In this chapter, we will look at popular individual DNS Blacklists, explain how to implement them on a mail server, and help you decide which list is the best one to use. When referring to DNS Blacklists, the shorthand DNSBL is often used, and that's how we'll refer to them throughout this chapter.

Before we talk about specific blacklists and how to implement them, we'll delve into what DNSBLs are and how they work.

UNDERSTANDING DNS BLACKLISTS

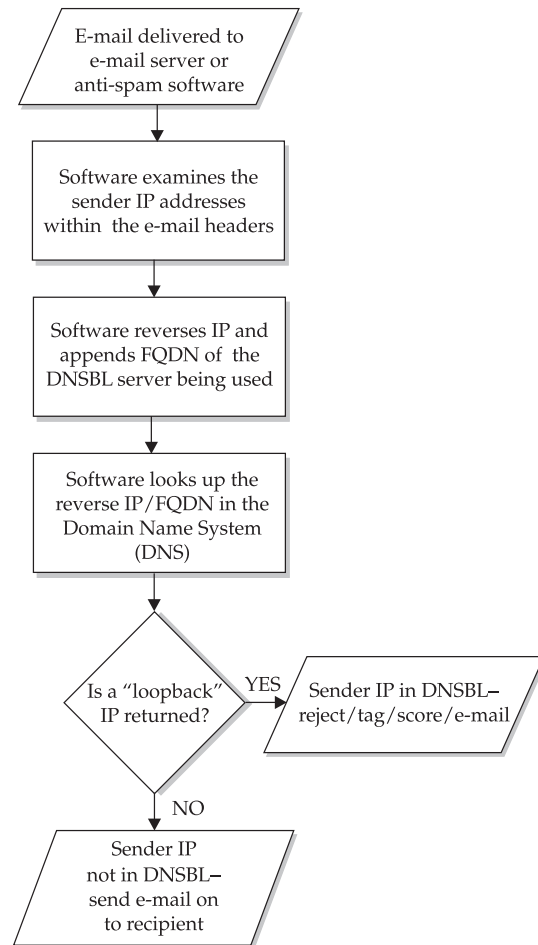
DNSBLs are an integral part of any spam-fighting toolkit. The fact that many, many users on the Internet are updating them means you get the benefit of blocking a spammer before the first piece of spam even hits you. To understand how DNSBLs help, you need to know the types of DNSBLs available and how they work.

Types of DNSBLs

Currently, two different types of DNS Blacklists are used:

- IP-based blacklists
- Domain-based blacklists

The majority of DNSBLs are IP-based, which look at the Internet Protocol (IP) address of the server sending the mail. Every host, including e-mail servers, connected to the Internet has its own unique IP address. This IP address is checked against a database to see whether or not it's for a known spammer, known open relay, or known open proxy. As a rule, IP-based DNSBLs work like this:

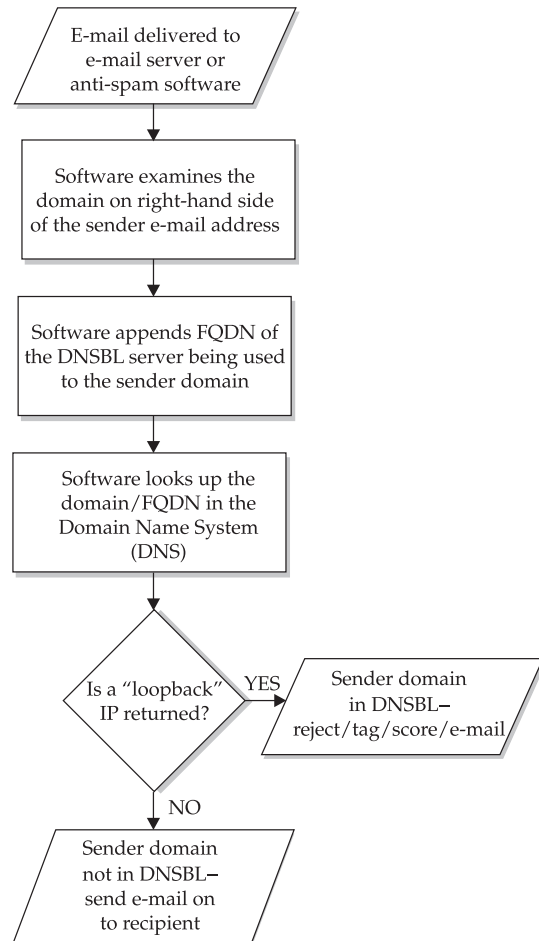


1. An attempted e-mail delivery to your mail server or anti-spam software occurs.
2. Your mail server or anti-spam software examines the IP addresses of the mail servers the e-mail passed through to get to you.

3. Your e-mail server or anti-spam software then reverses the order of the IP address and appends the fully qualified domain name (FQDN) of the DNSBL server being used. A FQDN is the Internet domain in which the server resides (such as ordb.org), plus the hostname of the server (such as relays.ordb.org). So, for instance, if the IP address of the sending server is 192.168.42.6, and the FQDN of the DNSBL server is relays.greatdnsbl.tld, the appended name is 6.42.168.192.relays.greatdnsbl.tld (note that a fictitious private IP address is used in this example).
4. Your server or anti-spam software then looks up the appended name in the Domain Name System (DNS). If an actual IP address is returned, the server exists in the blacklist. Otherwise, it doesn't exist or the blacklist is down. The IP address that gets returned should always be in the special-use "loopback" network (127.0.0.0/8, per RFC 3330). The IP address that gets returned may have some special meaning for an individual DNSBL.
5. Your mail server or anti-spam software uses the response from the DNSBL to decide what to do with the mail. If it determines that the mail has come from a system in the DNSBL, it may reject it, tag it, or score it, depending upon the software and its configuration.

Domain-based DNSBLs are also called *right-hand side blacklists* (RHSBLs). These lists look only at the second-level and top-level domains (for instance, the *.com* and the *yahoo* that comes before it—*yahoo.com*) of a given e-mail address or FQDN. Here's how they work:

1. An attempted e-mail delivery to your mail server or anti-spam software occurs.
2. Your mail server or anti-spam software looks at the domain on the right-hand side of the @ sign. For instance, *spammer.tld* would be parsed from *spam4you@spammer.tld*. Or, in the case of the server *spamserver.spammer.tld*, it would look only at *spammer.tld*.



3. This domain name is then appended to the front of the FQDN of the DNSBL server. For instance, if the parsed domain is spammer.tld, and the name of the DNSBL server is relays.greatdnsbl.tld, the appended name is spammer.tld.relays.greatdnsbl.tld.
4. Your server or anti-spam software then looks up the appended name in the DNS. If an actual IP address is returned (again, from the loopback network range), the server exists in the blacklist. Otherwise, it doesn't exist or the blacklist is down.
5. Again, your mail server or anti-spam software takes the response from the RHSBL and then uses it to decide what to do with the mail.

So why have both IP-based and domain-based DNSBLs? IP-based lists typically consist of systems that have sent spam in the past or that are capable of sending spam (such as open relays). They sometimes also consist of entire networks that are capable of sending spam (such as dial-up lists). They are more specific, because they include only IP addresses or networks that actually performed the behavior that got them onto the list (that is, sending spam), rather than just being part of that domain.

Criteria for DNS Blacklists

Currently, most DNSBLs use one or more of the following criteria to determine whether or not an Internet host belongs in the list. Sometimes a DNSBL organization has a separate blacklist for each criterion, or it might have one or two lists that employ a mix of these.

- Open-relay list
- Open-proxy list
- Known spammer lists
- Dial-up user list

We will briefly discuss each of these.

Open-Relay List

Open-relay lists are extremely popular forms of DNSBLs. A machine is considered an open relay if it allows unauthorized users to e-mail to a third party—that is, neither the person sending the mail nor the person receiving the mail are within domains for which the e-mail system is a mail server. In the early days of the Internet, when spam wasn't such a problem, many systems were configured as open relays. Nowadays, it's no longer necessary to run mail servers as open relays, so systems that remain as such are generally run by administrators who simply lack the time, effort, or knowledge to configure them otherwise. Open relays are attractive to spammers because they allow spammers to use someone else's resources for sending bulk e-mail, and these systems remove spammers one step from the sender—thus making them more difficult to track.

Open Proxy Lists

Open proxy servers are even worse than open relays. A proxy in the networking world is similar to a proxy in the real world—that is, it's someone, or something, who stands in or acts as a substitute for someone, or something, else. In this case, the proxy is a network device or server that makes a connection to a network resource for an end user, rather than the end user making the connection directly to the resource. For instance, if a user wants to go to the web page <http://www.mcgraw-hill.com>, she types that URL into her browser. If her browser is configured to use an HTTP proxy server, the proxy server gets the request, goes out to <http://www.mcgraw-hill.com>, pulls down the web page, and sends it on to the browser. The user never connects directly to <http://www.mcgraw-hill.com>.

Using a proxy provides a certain level of inherent anonymity. The user's IP address never appears in the McGraw-Hill web server's logs—the proxy's IP address does. Typically, proxies limit what networks can connect to them. Open proxies, however, allow *any* user to connect to them from anywhere, and use them to go anywhere. An open proxy, therefore, allows a spammer to make a Simple Mail Transport Protocol (SMTP) connection to a mail server, but it hides where the mail is coming from. To the system to which they're connecting, it looks like the sender is the proxy—nothing behind that is seen by the mail server. It's yet another powerful tool in the spammers' toolkit.

Known Spammer Lists

Known spammer lists are domains, systems, or networks that are known spammer dens. Unlike open-relay and open-proxy servers, these are in the list not for technological reasons, but because they've actually sent spam.

Dial-Up User Lists

The idea behind dial-up user lists is to prevent what the Mail Abuse Prevention System organization (MAPS, which we'll talk about a little later) calls "spam trespassing," whereby a spammer dials into an ISP (possibly using a forged "trial" account) with a system running a mail server or "ratware" and sends out bulk e-mail. By doing so, the spammer gets around traditional spam detection measures. The spammer uses the ISP's network resources but not its server resources. These lists often also include other ISP accounts with dynamic IP addresses, such as cable modems and DSLs.

Adding or Removing Entries from a DNS Blacklist

Each DNSBL has a method for adding or removing entries from a blacklist. This is most often performed through the DNSBL's web page, although it can also be done through e-mail. Sometimes you have to be a "member" of a DNSBL to add entries to its blacklist, which sometimes means you have to be a paying customer (or donator).

Anyone who wants to remove a server, network, or domain from a DNSBL can generally do this, too. This is done at the DNSBL's web site. The listed item is retested before it is removed. If the listed server, network, or domain is under the control of a known spammer, or if it has multiple strikes against it, it may not be removed at all.

CHOOSING A DNS BLACKLIST

Many DNSBLs are out there—more than 100 public ones, and who knows how many private ones. Organizationally, they typically fall into three categories:

- Nonprofit organizations that are dedicated to spam-fighting. These organizations generally have employees and a semicorporate structure.
- A loose-knit group of administrators who have banded together to fight spam. These groups usually do not have full-time employees, and they resemble open-source projects more than anything else. (For example, though there may be a recognized leader, they follow democratic principles in decision-making.)
- Individuals who have set up their own DNSBLs for their own private use, but who allow others to use them if they like.

In addition to their organizational structure, DNSBLs differ greatly in the way they operate. Areas in which there might be differences include the following:

- Criteria on what constitutes a spammer (or potential spammer)
- The method they use to obtain candidates for their black hole lists
- Their rules for getting removed from their list
- The type of lists they run (open-relay, known spammer, dial-up, and so on)
- Whether or not you have to pay to use the service

All of these influence the effectiveness and ease-of-use of a particular DNSBL.

Because the lists are dynamic in nature and structured around the particular philosophies of their organizers, none can be considered the be-all and end-all authority on who is and who is not a spammer. All have the potential to give false positives and false negatives.

In general, nonprofit organizations have stricter rules as to which systems or domains end up in their blacklists, followed by the loose-knit groups, and then individuals. Therefore, if you're concerned about false positives, an organization is the better way to go.

The DNSBL that is right for you fits in well with your personal or organizational policy on spam. How tolerant are you of a little spam getting through? How tolerant are you of legitimate mail being rejected? Are you willing to pay (or donate) to use a blacklist?

An extensive list of blacklists can be found at the Declude web site: <http://www.declude.com/junkmail/support/ip4r.htm>. However, we'll overview several of the most popular and effective ones in the following sections.

MAIL ABUSE PREVENTION SYSTEM (MAPS)

Mail Abuse Prevention System, LLC (MAPS for short), is one of the biggest, oldest, most controversial, and most well-known DNSBLs around. Formed in 1997 by a small group that included Internet developer Paul Vixie (author of the BIND software found on the

majority of Internet DNS servers), MAPS is a nonprofit corporation based in California. MAPS main web site can be found at <http://www.mail-abuse.org>.

Vixie's reputation and knowledge has given MAPS a lot of respect among system administrators and a lot of profiling in the press. Unfortunately, MAPS's methods (which really aren't all that different from many other blacklists), its high profile, and its ubiquity have made it the legal target for numerous bulk e-mail senders who feel that MAPS is unfairly persecuting them. MAPS even collects money online for its legal defense fund!

How MAPS Works

MAPS has one of the widest assortments of DNSBLs available. A different group within the MAPS organization operates each one, and each has its own documentation, policies, FAQ, and tools. Following is a brief rundown of the available lists.

Nonconfirming Mailing List (NML)

<i>List type:</i>	<i>IP-based</i>
<i>DNSBL server:</i>	<i>nonconfirm.mail-abuse.org</i>

The Nonconfirming Mailing List is MAPS newest offering. The NML contains a list of reported IP addresses that send e-mail to mailing lists without fully verifying that the recipient actually requested the information they're sending out. In other words, the senders didn't verify that the e-mail was solicited. Other things that could get someone on the list:

- Not fully disclosing the terms and conditions of the list
- Not using lists for their original intended purpose
- Not obtaining a separate verification for each list to which they add a subscriber

Opt-out clauses, by the way, do not count.

Dial-Up User List (DUL)

<i>List type:</i>	<i>IP-based</i>
<i>DNSBL server:</i>	<i>dialups.mail-abuse.org</i>

The DUL contains a list of reported IP address ranges that are part of an ISP's dial-up network or some other dynamically assigned range. Some ISPs work with MAPS to have their own dial-up networks added to the lists, so they won't be the unwitting accomplices of spammers.

Of course, some legitimate computer hobbyists and end users run SMTP servers on their dial-up accounts. To them, MAPS suggests that they do not send mail directly from their SMTP server, but instead send it through their ISP's SMTP server.

Relay Spam Stopper (RSS)

<i>List type:</i>	<i>IP-based</i>
<i>DNSBL server:</i>	<i>relays.mail-abuse.org</i>

The RSS contains a list of IP addresses reported to have sent bulk e-mails, most of which are open relays. Just because a system is an open relay, however, doesn't mean that it will be in the RSS—it actually has to have sent spam. In this way, this list differs from the ORDB's list, which we'll discuss later in this chapter in the section "Open Relay Database (ORDB)."

Realtime Blackhole List (RBL)

<i>List type:</i>	<i>IP-based</i>
<i>DNSBL server:</i>	<i>blackholes.mail-abuse.org</i>

This is the granddad of all MAPS lists. It includes networks or hosts that fall under the following criteria:

- They have been reported to send spam.
- They are an open relay.
- They are an open proxy.
- They provide support services to spammers, such as web hosting, software, e-mail drop boxes, and more.

Because they include the support-services aspect, the chance of wanted e-mails being blocked is higher than it might be with other, more singular lists.

Realtime Blackhole List+ (RBL+)

<i>List type:</i>	<i>IP-based</i>
<i>DNSBL server:</i>	<i>blackholes.mail-abuse.org</i>

The RBL+ is MAPS's premier subscription service and combines most of its lists, including the RBL, DUL, and RSS. It also includes a list that's available only to RBL+ subscribers, called the Open Proxy Monitor (OPM). The OPM is similar to the RSS, but it contains IPs of systems that are open-proxy servers and have been known to send spam. The RBL+ provides a robust blacklisting method, but you also have a higher chance of blocking wanted e-mails when you use it.

Subscribing to MAPS

After years of being free, MAPS converted to a fee-based service in July 2001. The service is still free for individuals and hobbyists, but all others must pay a yearly subscription fee. Payment levels depend on the type and size of your organization:

- Nonprofit/Educational
- Small Business (less than 100 users)
- Standard (everyone else)

Each list has its own price, with the RBL+ being the most expensive and the DUL being the cheapest. Each list also has different prices for a DNS query versus a DNS zone transfer. The Standard pricing level also has a per-user cost.

The RBL+ also has the capability for a Border Gateway Protocol (BGP) routing feed. BGP is the standard routing protocol on the Internet “backbone.” Put another way, it’s how Internet traffic knows how to get from point A to point B. If you’re an ISP, MAPS can insert routes into your BGP routing table that will prevent spammers from reaching your network. This is a relatively extreme measure, as it blocks *all* Internet traffic from the spammer, not just mail traffic. BGP implementation is beyond the scope of this book.

No matter which level you are, even if you’re a hobbyist or individual, you must mail MAPS a signed agreement for the list to which you’re interested in subscribing. These agreements are available on the MAPS web site and are, on average, a dozen pages long. In this agreement, you agree that (among other things) you’ll hold any information MAPS gives you confidential, and that you understand that just because you’re a subscriber doesn’t mean you’re exempted from getting listed in a MAPS blacklist yourself.

In the agreement, you also put the IP address of the DNS or mail server you intend to use for querying or zone transferring with MAPS (or the address of your router if doing BGP). MAPS uses filters that prevent anyone who hasn’t signed an agreement from accessing their lists.

If you’re not a hobbyist or individual, pricing for MAPS’s services is across the map (pun intended). They can range from \$50 per year for a zone transfer of the DUL list for nonprofit pricing, to \$1500 (or more) per year for a query against the RBL+ under standard pricing. For MAPS’s latest Schedule of Annual Fees, visit its web site (<http://mail-abuse.org/feestructure.html>).

SPAMCOP

SpamCop is a popular DNSBL that has been around since 1998. SpamCop itself is based in Seattle and is run by Julian Haight (who wrote the code) and many contributors. SpamCop has a unique method of keeping its list fresh and removing sites that are no longer spamming in a timely manner, and that appears to make it one of the “fairest” DNSBLs around.

Apart from being a DNSBL, SpamCop also offers filtered POP/IMAP/web-mail accounts. In this chapter, we're not going to go into this service in-depth, but we'll briefly describe it since it's a major component of their business. The SpamCop service currently costs \$30 per year per e-mail account. You can either forward your existing e-mail account to the service or use it as a new account under the domain spamcop.net (that is, your e-mail address would be *your_account@spamcop.net*). The service checks incoming mail for viruses, and then it uses SpamCop's DNSBL to decide whether or not the message is spam. If it decides it is spam, it drops the mail into a Hold folder for you to peruse (or not) later. You can check your mail using your existing mail client via either POP or IMAP, or you can use SpamCop's web-mail page.

How SpamCop Works

<i>List type:</i>	<i>IP-based</i>
<i>DNSBLserver:</i>	<i>bl.spamcop.net</i>

Unlike MAPS, SpamCop is a single list. SpamCop doesn't care whether you're running an open relay, open proxy, or mailing from a dial-up IP—it makes no checks of the site's configuration. The only way a site gets on the list is by having someone report it to SpamCop as being a spam sender. Anyone can sign up for SpamCop's free reporting service (though it does require a verifiable e-mail address).

SpamCop uses a ratio-based scoring system to decide how long to keep a site on the blacklist. The score is based on several factors, including how fresh the spam is, how long it has been on the blacklist, and how SpamCop was notified of the spam. We'll go into details on individual scores in the list that follows.

What follows is a breakdown of the submission process and how mail hosts are automatically added or removed from the blacklist:

1. A registered SpamCop user forwards suspected spam (with full headers!) to SpamCop as an attachment. SpamCop's site has a good list of criteria for what it considers spam and what therefore should be forwarded (viruses, hoaxes, and chain letters don't count, for instance).
2. If this is the first spam from that server to hit SpamCop's database, the server won't be immediately listed. If it's the second, it will be listed for 24 hours, unless more suspected spam messages arrive.
3. Suspected spam sites are weighted by freshness. The more recently a spam was received, the higher the score the site that sent it gets. Fresh spam gives the site a premodified score of 4:1, with it sliding down until it reaches a score of 1:1 after 48 hours. Reports older than one week are ignored when calculating weights.
4. SpamCop has a number of addresses called "spamtraps," which are e-mail addresses created for the sole purpose of collecting spam. Because these spamtraps aren't real people, they've never subscribed to any legitimate lists

and do not receive legitimate mail. Reports coming from these spamtraps have their scores multiplied by 5. For instance, if a site has two reports against it from spamtrap sites, it receives a total score of 2×5 , or 10:1.

5. To avoid blocking legitimate mail from large sites that might also squeeze out some spam every once in a while (such as AOL), the spam reports are balanced against the total number of legitimate messages sent. This is done by monitoring selected third-party sites. Whenever one of these third-party sites checks against the SpamCop blacklist for an IP address that isn't in there, that host is given a nonspam point. After 1000 points, each subsequent point that a host receives counts for only half of its full value. For instance, a host that gets 3000 nonspam reports will only have 2000 total points.
6. If no spam reports have come in for a listed site within 48 hours, that site is removed from the list.

It's important to note that SpamCop does not block sites that support spammers, such as web sites or e-mail drop boxes; only those that actually sent the spam are blocked.

Subscribing to SpamCop

Unlike using their POP/IMAP/web-mail service, using SpamCop as a DNSBL is free. The company does, however, request a donation to help it keep the service running. If your site has 1 to 10 users, it requests \$50 per year; for 11 to 100 users, it requests \$150 per year; for sites with more than 100 users, it requests \$1 per user per year. The suggested minimum donation for a site larger than 100 is \$150; the minimum for a site of 10,000 users is \$1000 per year.

SpamCop doesn't allow DNS zone transfers because its list is too dynamic to make zone transfers efficient. Instead, it allows you to mirror the list to your own DNS server using the rsync and SSH utilities. This service, however, costs \$1000 per year.

You can make a donation (or pay for the mirroring) using PayPal on SpamCop's web site, or you can mail the company a check at its physical address.

OPEN RELAY DATABASE (ORDB)

The Open Relay Database (ORDB) lists open relays on the Internet and has done so since 2001. ISPs and organizations large and small use ORDB. It's a nonprofit headquartered in Denmark, but it has contributors and users the world over.

ORDB is simply a listing of reported and tested open relays. It's different than MAPS's RSS, in that ORDB doesn't care whether or not the site has actually been used to send spam before, only that it's *technically* an open relay. Therefore, if you use it, there's a chance of blocking mail from legitimate mail servers, even if those servers aren't currently being used by a spammer. The thought is that it's only a matter of time before a spammer finds the open relay and *does* use it, so it's better to go ahead and block the site. This also encourages systems administrators to close their open relays if they know their servers might get blocked. The company's web site is at <http://www.ordb.org>.

How ORDB Works

List type:	IP-based
DNSBL server:	relays.ordb.org

Open-relay candidates are tested when someone enters the IP address of the servers they wish tested on the ORDB web site. The person requesting the test must put in a valid e-mail address, because a reply is required before testing commences. This adds at least some accountability and reduces abuse of the system. Testing can take up to 72 hours, depending on the size of the queue. If a site fails the test and is found to be an open relay, that site will be added to the ORDB blacklist.

Removal of a site occurs in the same way. The IP addresses are entered into the web form, and e-mail is sent to which the submitter must reply; then the site is rechecked. If it passes and isn't found to be an open relay, the site is removed within 72 hours. If the test still detects an open relay, the site stays on the list.

Subscribing to ORDB

ORDB does not require payment, nor do you sign an agreement. Anyone is free to use it, and the organizers believe that use is the best way to support ORDB. It does, however, accept donations of any size on its web site via PayPal, or you can mail a check or money order to the postal address on the site.

DISTRIBUTED SERVER BOYCOTT LIST (DSBL)

The Distributed Server Boycott List (DSBL) is a group of administrators and users who have banded together to fight spam. They are primarily concerned with spam sources that are open relays or open proxies. The DSBL web site is at <http://www.dsbl.org>.

How DSBL Works

The DSBL does not test sites itself. Anybody can submit a site to the DSBL, but two types of users actually do so: *untrusted* and *trusted*. An untrusted user is anyone who reports spam to the DSBL. A trusted user is someone who has requested an account with the DSBL and provides a rationale as to why they should be a trusted user. The DSBL then gives the user a provisional account that can be revoked if the user violates reporting standards the DSBL sets forth.

List

List Type:	IP-based
DNSBLServer:	list.dsbl.org

This list contains only spam sources verified by users trusted to the DSBL staff. Because of this, it has a lower incidence of false positives than the other DSBL lists.

Multihop

List Type:	IP-based
DNSBL Server:	<i>multihop.dsbl.org</i>

This list contains only multihop relayed spam sources that are verified by users trusted to the DSBL staff. Even though the users are trusted, this list still may cause false positives because it catches all the hops in a spam's path.

Unconfirmed

List Type:	IP-based
DNSBL Server:	<i>unconfirmed.dsbl.org</i>

Additions to this list are made by untrusted users, so there's a high likelihood of false positives. It's best to use this list sparingly or with some other program, such as SpamAssassin, that just scores or tags likely spam rather than rejecting it outright.

Subscribing to DSBL

No subscription or form is required for any of the DSBL lists. Simply start using them.

In addition to regular query lookups, you can also download the entire zone file using the `rsync` utility or via HTTP at DSBL's web site.

SPAMHAUS

The operators of Spamhaus believe that 90 percent of all spam in Europe and North America is sent by less than 200 known spammers, which they keep up with in their Registry of Known Spam Operators (ROKSO). By knowing their enemy and tracking their movements from one ISP to another, the Spamhaus Blacklist (SBL) has become a popular and effective DNSBL. Its web site is at <http://www.spamhaus.org>.

How Spamhaus Works

List type:	IP-based
DNSBL Server:	<i>sbl.spamhaus.org</i>

Spamhaus is updated around the clock by an international team of administrators who are on the watch for known spammers and spamming in progress. Updates are made to the SBL once per hour.

The SBL contains a list of known spammers. It does not contain a list of open proxies or open relays, so Spamhaus suggests using its list in combination with a good open relay/open

proxy list. It uses the following criteria to decide whether a given IP address will be listed in the SBL:

- **Spam Sources** Spammers sending bulk e-mail from an IP address directly under the spammer's control
- **Spam Gangs** Netblocks of known spammers listed in the ROKSO
- **Spam Services** Web servers, mail servers, DNS servers, and other services used by spammers
- **Spam Support Services** IP addresses that knowingly provide hosting, spamware, or other support services for spammers

IP addresses that end up in the SBL remain there until the spamming source has been removed and Spamhaus has been notified. Spamhaus itself does no further checks on the IP addresses. However, to keep entries from getting old, Spamhaus expires records from the SBL. The timeouts are two, seven, or fourteen days for unidentified spam sources; six months for persistent spammers; or a year or more for well-known spammers such as those in the ROKSO.

Subscribing to Spamhaus

Spamhaus doesn't cost anything to use, and you don't have to fill out any forms to use it. You simply set it up as you would any other DNSBL. It even allows DNS zone transfers for free if you're from a large enough organization (such as an ISP, university, or large company). To request this service, you must contact Spamhaus directly.

NOT JUST ANOTHER BOGUS LIST (NJABL)

Not Just Another Bogus List (NJABL) is run by a group of e-mail administrators who are frustrated with the policies and uptime of existing DNSBLs. They decided to take matters into their own hands and created a blacklist that is almost entirely supported by the e-mail administrators who use it.

How NJABL Works

<i>List type:</i>	<i>IP-based</i>
<i>DNSBL server:</i>	<i>dnsbl.njabl.org</i>

The NJABL has only one list. The NJABL includes any IP address in its list that meets the following criteria:

- The system is an open relay, open proxy, or is running an open web-to-mail gateway.

- The IP address belongs to a dial-up or dynamic range. This information is received through American Registry for Internet Numbers (ARIN) records or by ISPs reporting such ranges to them directly.
- The system has been used to send out spam directly.

NJABL tests for open proxies by scanning individual servers for them. It then does an open-relay test by scanning the SMTP port. It takes about four weeks to remove a system from the list.

Subscribing to NJABL

NJABL is basically run by the administrators who use it. You don't need to pay anything to subscribe, nor do you have to sign up anywhere. You can simply start using it whenever you like. The only thing they ask is that you subscribe to list@njabl.org to receive the latest announcements.

You can also contribute to NJABL by feeding IP addresses that connect to your mail server. You change your connection method so that anyone connecting to your server consents to be scanned as an open relay. Later, NJABL tests each of those servers for open relays.

NJABL is normally used in query mode, but you can also get an rsync zone transfer by e-mailing help@njabl.org.

RFC IGNORANT (RFCI)

RFCI sets itself apart from the rest of the DNSBLs because it is not concerned whether or not a site is, or could be, a spammer. In fact, instead of worrying about spam, RFCI is more concerned about whether or not a domain or IP network block's administrator is a good "Netizen" (that is, citizen of the Internet). The domain and network blocks in the RFCI have been placed there because their owners are deemed "RFC ignorant."

You may have heard the term *RFC* before and are vaguely aware of it being related to Internet standards (numerous companies tout their products as being "RFC compliant"). *RFC* stands for *Request for Comments*, which is the common name for the rules and best practices ratified and published by the Internet Engineering Task Force (IETF). Many RFCs are like technical blueprints, but some RFCs are human protocols—best practices for configuring and running networks and servers on the Internet. Domains are listed in the RFCI because their owners have refused or ignored these policies and procedures.

How is this related to spamming? Spammers try to hide as much information about themselves as possible. Most of the RFC practices that RFCI checks for relate to actually getting into contact with a human being. If a spammer has registered a domain, he's likely to provide bogus contact information so that you can't track him down and complain (or, possibly, litigate). Hence, a good portion of RFCI's denizens are likely to be spammers.

RFCI can be found at <http://www.rfc-ignorant.org>. The online RFC papers can be found at <http://www.ietf.org> under "RFC Pages."

What Makes Someone RFC Ignorant?

RFCs aren't laws, so no criminal penalty is enforced for breaking them. Since the beginnings of the Internet, however, administrators have chosen to shun unrepentant violators (some software programmatically does this). For inclusion in the RFCI, a site has to demonstrate that its owners failed to implement one or more of the RFC guidelines listed in the following sections.

Delivery Status Notification (DSN)-Related

<i>List type:</i>	<i>Domain-based</i>
<i>DNSBL server:</i>	<i>dsn.rfc-ignorant.org</i>
<i>Relevant RFCs:</i>	<i>821, 2821, 2505, 1123</i>

Inclusion in this list occurs if the mail exchanger (MX) record for the sender's domain does not accept mail with the originating address given as <> (blank). For instance, if we send mail to the MX server `exampledomain.tld` using these SMTP commands,

```
MAIL FROM <>
RCPT TO <postmaster@exampledomain.tld>
```

and the server doesn't accept the message, `exampledomain.tld` is added to the list. All MX records for a domain are checked, and having one that doesn't accept mail with a blank sender address is cause for inclusion. Domains with MX records that contain private or reserved IP addresses (for instance, a loopback address or networks in RFC 1918, such as 10.0.0.0/8) are also listed.

Postmaster-Related

<i>List type:</i>	<i>Domain-based</i>
<i>DNSBL server:</i>	<i>postmaster.rfc-ignorant.org</i>
<i>Relevant RFC:</i>	<i>2821</i>

The postmaster address is designated for reporting problems with mail servers. Inclusion in this list occurs if the sender's domain (with an MX record) does not have a valid postmaster e-mail address. An example is `postmaster@exampledomain.tld`. And, yes, e-mails to the postmaster must ultimately go to a human being, even if an auto-response is sent first.

Abuse-Related

List type:	Domain-based
DNSBL server:	abuse.rfc-ignorant.org
Relevant RFC:	2142

The abuse address is designated to report spamming, fraud, and other mail-related incidents by a given domain's users. Inclusion in this list occurs if the sender's domain (with an MX record) does not have a valid abuse e-mail address—for example, abuse@exampledomain.net. As with the postmaster, the abuse address must ultimately go to a human being.

WHOIS-Related

List type:	Domain-based
DNSBL server:	postmaster.rfc-ignorant.org
Relevant RFC:	954

The WHOIS database contains contact information for owners of a domain, including e-mail addresses. This applies to both the generic Top Level Domains (or gTLDs, such as .com, .org, or .net), as well as the country-coded Top Level Domains (ccTLDs, such as .us, .uk, or .de). The registrars for these domains run their own WHOIS servers to provide this information.

Many domain registrars, however, don't check up on the information their customers provide, so it's possible to include bogus points of contact, leave some information blank, or let information get stale.

A domain is included in this list if it meets the following criteria:

- Information is obviously bogus or wrong (such as a U.S. phone number of 555-1212 or address of 1600 Pennsylvania Ave., Washington, D.C.—for any user other than the White House).
- Information is provably wrong, such as phone lines that are out of service, e-mail addresses that bounce, or returned snail mail.
- The TLD for that domain does not have an operating WHOIS server, referred to by the root WHOIS server at whois.iana.org.
- The domains in the e-mail contact addresses do not have valid MX records, or have records that are obviously bogus (such as a loopback address or an RFC 1918 reserved address such as 10.10.10.10).

A domain's record does not need to have a valid fax number, as not everyone has a fax. However, the domain site is expected to have valid voice phone numbers, e-mail accounts, and postal addresses that someone actually checks and responds to.

IPWHOIS-Related

List type:	IP-based
DNSBL server:	postmaster.rfc-ignorant.org
Relevant RFC:	954

IPWHOIS is much like WHOIS except it relates to the registry contact information for IP address network blocks, rather than domains. The policy for inclusion in this list is practically the same as for the WHOIS databases. If any bogus information is included in the IPWHOIS record in that netblock's Regional Internet Registry (RIR), the netblock will be included. If a large netblock is included in the list, such as 192.168.0.0/16, all subnets under that network are also included (for instance, 192.168.5.0/24). Thus, a problem with the parent network causes problems with the child. In other words, if your ISP turns out to be noncompliant, your network isn't compliant either.

It's important to note that IPWHOIS is the only RFCI service that is IP-based (that is, the criteria is based on the IP address of the e-mail's sender, rather than the domain).

Subscribing to RFCI

Use of RFCI is completely free, and you don't need to sign an agreement to use its services. Just put the listing servers you want to use in your DNSBL-capable mail server or client, and that's it. RFCI provides a web interface to report RFC ignorance and e-mail address to have a domain or IP netblock removed. You can also subscribe to a mailing list for community-based support.

Now that you've been introduced to some of the blacklists, we'll show you how to implement them. We focus on three mail server packages: Sendmail and Postfix for Linux and Unix, and Exchange for Microsoft Windows.

IMPLEMENTING DNSBLS WITHIN SENDMAIL

Sendmail is the most venerable mail transfer agent on the Internet and runs on most Unix and Unix-like operating systems. (In fact, almost every major distribution of Linux comes with Sendmail.) Sendmail added direct DNSBL support with version 8.9, and changed the syntax slightly in 8.10. Support was also possible in version 8.8, but you had to hack your `sendmail.cf` configuration file to do it. However, due to security vulnerabilities in earlier versions of Sendmail, including major buffer overflows discovered in the spring and fall of 2003 (see CERT Advisories CA-2003-12 and CA-2003-25), we highly recommend that you run the latest stable version. You can find Sendmail at <http://www.sendmail.org>.

Configuring Sendmail for IP-Based DNSBLs

By Sendmail standards, configuring Sendmail to use IP-based DNSBLs post version 8.10 is quite straightforward. Edit the `sendmail.mc` file, or equivalent `.mc` file, that you're using. If you are unfamiliar with `.mc` files, they are macro definitions compiled with the `m4` application to create `sendmail.cf` configuration files. Check Sendmail's documentation or web site for more information. DSNBL support is configured as an `m4` Feature.

The easiest DNSBL to use is the MAPS RBL, as support is preconfigured. All you have to do is edit your `sendmail.mc` file and add the following line *before* the MAILER section:

```
FEATURE(`dnsbl')
```

Note that the first single quote is actually the *grave accent* (a backward apostrophe that's typically on the upper-left side of your keyboard).

Then, back up your current `sendmail.cf` file and compile your `mc` file to create the `.cf` file:

```
# m4 sendmail.mc > sendmail.cf
```

If it's successful, you won't get any messages. Then restart Sendmail using the new `.cf` file.

To use other DNSBLs, you have to specify the server within the `FEATURE` command. Here's the basic syntax:

```
FEATURE(`dnsbl', <zone server>, <rejection message for remote server and logs>)
```

Here's the ORDB's DNSBL for an example:

```
FEATURE(`dnsbl', `relays.ordb.org', `550 Email rejected due to sending server ↵
misconfiguration - see http://www.ordb.org/faq/#why_rejected"')dnl
```

For multiple entries, add a `FEATURE` line for each relay you wish to use.

Configuring Sendmail for Domain-Based RHSBLs

Using RHSBLs under Sendmail requires that you add a different `FEATURE` command to your `sendmail.mc` file. This feature is not part of the standard Sendmail distribution. You can download Derek J. Balling's `rhbl.m4` file from http://www.megacity.org/software_downloads/rhbl.m4. Put this file in the `./cf/feature` directory of your Sendmail source tree, or wherever you have stored your `.cf` configuration files. Then add the RHSBL to your `.mc` file using this syntax:

```
FEATURE(rhsbl, <zone server>, <rejection message for remote server and logs>)
```

For example, here's how to use the RFCI's DSN domain-based blacklist:

```
FEATURE(rhsbl, `dsn.rfc-ignorant.org', `550 Mail from domain " ${RHS} "
refused. ↵
MX of domain do not accept bounces. This violates RFC 821/2505/2821 - see ↵
http://www.rfc-ignorant.org/"')
```

As with the DNSBLs, add a `FEATURE` statement for each RHSBL you want to use.

IMPLEMENTING DNSBLS WITH POSTFIX

Postfix is another popular Sendmail replacement for Unix and Unix-like operating systems, designed by Wietse Venema while he worked at IBM. IBM released it to the public in 1998 and Postfix is credited with instigating Big Blue's open-source strategy. Postfix was designed with security and speed in mind. We're going to cover only Postfix 2.x here. Postfix can be had at <http://www.postfix.org>.

Configuring Postfix for IP-Based DNSBLS

Postfix is probably one of the simplest mail servers to configure for DNSBLS. All you have to do is edit your `main.cf` file (usually in `/etc/postfix`) and modify (or add) the `smtpd_client_restrictions` line. This line controls many spam-related functions, including anti-relaying provisions.

To configure Postfix 2.x to reject mail included in a DNSBL, make the following modifications:

```
smtpd_client_restrictions =  
    reject_rbl_client <zone server>
```

where `<zone server>` is the DNSBL you want to use. For instance:

```
smtpd_client_restrictions =  
    reject_rbl_client sbl.spamhaus.org
```

To add multiple listings, separate each `reject_rbl_client` line with a comma:

```
smtpd_client_restrictions =  
    reject_rbl_client sbl.spamhaus.org,  
    reject_rbl_client relays.ordb.org,  
    reject_rbl_client lists.dsbl.org
```

Restart Postfix to implement the changes. Rejected e-mails show up in your log file.

Configuring Postfix for Domain-Based RHSBLS

Postfix 2.x has built-in RHSBL support. You configure it much the same way as you do DNSBLS, except that you use the `reject_rhsbl_sender` command. Configure it like this:

```
smtpd_client_restrictions =  
    reject_rhsbl_sender <zone server>
```

Here is an example with an RFC Ignorant list:

```
smtpd_client_restrictions =  
    reject_rhsbl_sender dsn.rfc-ignorant.org
```

To add multiple RHSBLs (or combine them with DNSBLs) simply separate them by commas:

```
smtpd_client_restrictions =  
    reject_rbl_client sbl.spamhaus.org,  
    reject_rbl_client relays.ordb.org,  
    reject_rbl_client list.dsbl.org,  
    reject_rhsbl_sender dsn.rfc-ignorant.org,  
    reject_rhsbl_sender postmaster.rfc-ignorant.org,  
    reject_rhsbl_sender whois.rfc-ignorant.org
```

Restart Postfix after making changes. Rejected e-mails will show up in your log file.

IMPLEMENTING DNSBLs WITH MICROSOFT EXCHANGE

While Sendmail and Postfix are still the darlings of many ISPs, data centers, and Unix-based shops, Microsoft's Exchange Server is the most popular corporate e-mail system. While many companies no doubt front-end their Exchange Server with a server running one of the three implementations mentioned earlier, without a doubt many Exchange servers also process mail on the front-end.

Exchange 2000 (assuming most of you have upgraded from Exchange 5 and 5.5 by now) had little in the way of anti-spam features. Exchange 2000 required third-party plug-ins or front-end servers to deal with spam. Exchange 2003 has rectified the situation by introducing its own spam-fighting features. We'll discuss this progression as it pertains to DNSBLs.

Exchange 2000

Exchange 2000 requires that you use a third-party application to implement DNSBLs. Almost any anti-spam package for Exchange supports DNSBLs, including those we cover in Chapter 11. You can also use two freeware (but not open-source) programs with Exchange 2000.

The first is a freeware version of GFI Mail Essentials, which can be found at <http://www.gfi.com>. The freeware version lacks a few of the anti-spam features of its full version, but for DNSBLs it should work just fine. If an e-mail is identified as coming from a spammer, GFI lets you tag it and optionally shunt it off into a public folder. Note that GFI also requires that you run Internet Information Services 5 (IIS 5) SMTP service.

Another piece of free DNSBL software for Exchange is ORFilter, available at <http://www.martijnjongen.com/eng/orfilter/>. You can configure ORFilter to block spam or to tag it. ORFilter runs on Exchange 2000 or the Microsoft SMTP service.

Exchange 2003

Microsoft Exchange 2003 has the built-in ability to use IP-based DNSBLs. RHSBL support may come in the future, but it is not currently available. As might be expected in a Windows environment, a nice graphical user interface (GUI) is available for managing DNSBLs, blacklists, and whitelists.

To add DNSBLs to Exchange 2003, open the Exchange System Manager by selecting Start | Programs | Microsoft Exchange (if you've installed it elsewhere, go to the appropriate menu item). The Exchange System Manager should appear. Under the Tree tab on the left-hand side of the window, open the Global Settings folder, as shown in Figure 5-1.

Right-click the Message Delivery item underneath the Global Settings folder and select Properties. Then click the Connection Filtering tab at the top. You will see a Message Deliver Properties window like that shown in Figure 5-2. The Block List Service Configuration box includes two columns: Rule and Enabled. If you've never configured a DNSBL on your server before, these areas will be empty.

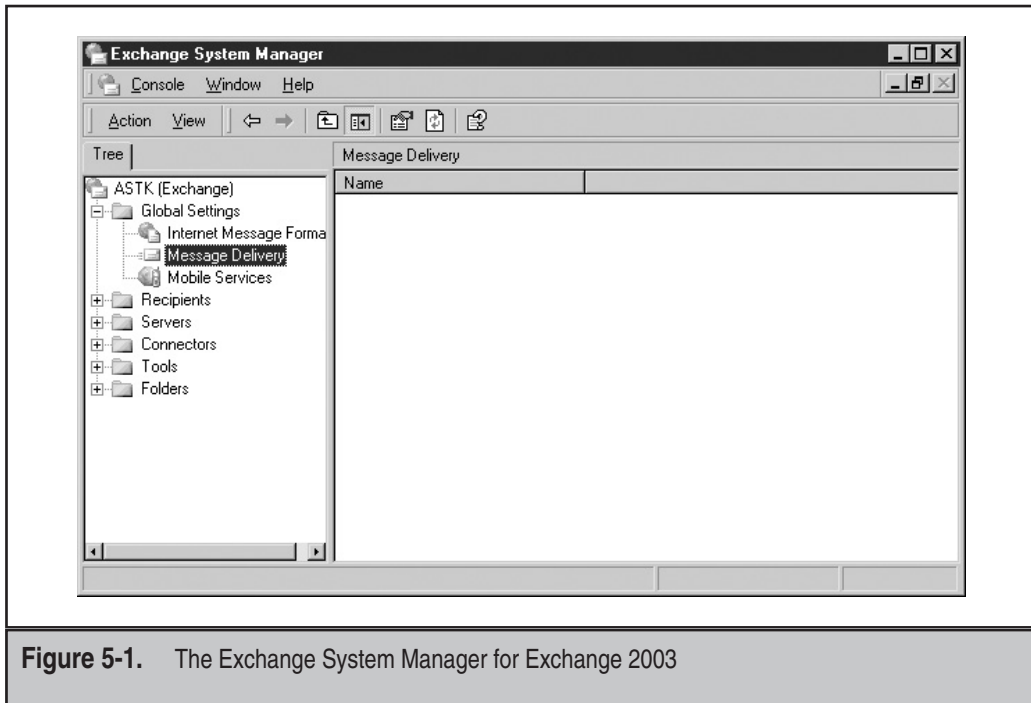


Figure 5-1. The Exchange System Manager for Exchange 2003

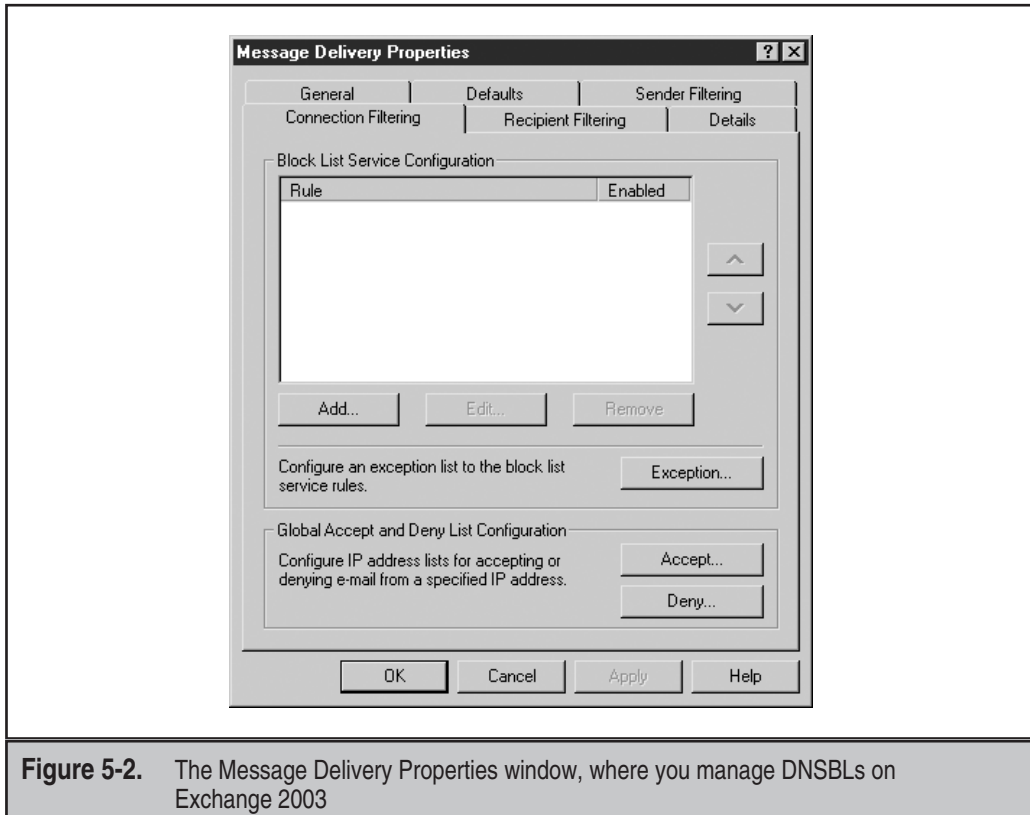


Figure 5-2. The Message Delivery Properties window, where you manage DNSBLs on Exchange 2003

To add a DNSBL, click the Add button. A Connection Filtering Rule dialog box pops up. This is where you'll add your DNSBL information. The first field is Display Name, where you can type in any description you want to give this particular rule. Our first DNSBL will be the Spamhaus SBL; therefore, we'll call our first rule, SpamHaus Black List. The second field, DNS Suffix Of Provider, is equivalent to what we've been referring to as the DNSBL Server in this chapter. Spamhaus' zone server is *sbl.spamhaus.org*, so we'll add that in. This is really all you need to get started, and it will look like Figure 5-3.

At this point, you can go back and add other IP-based DNSBLs in addition to Spamhaus.

Another option in this window is to add a custom error message by entering it into the Custom Error Message To Return field. This is an optional field and is not required for the DNSBL to work. It is useful, however, if you want the mail to be rejected with some specific error message (no taunting, please!).

Finally, you can also modify what return status code from the DNSBL the Exchange server uses by clicking the Return Status Code From Provider Service button. You will then see the Return Status Code dialog box, as shown in Figure 5-4. Remember that most DNSBLs respond with "codes" (actually IP addresses) within the loopback range of

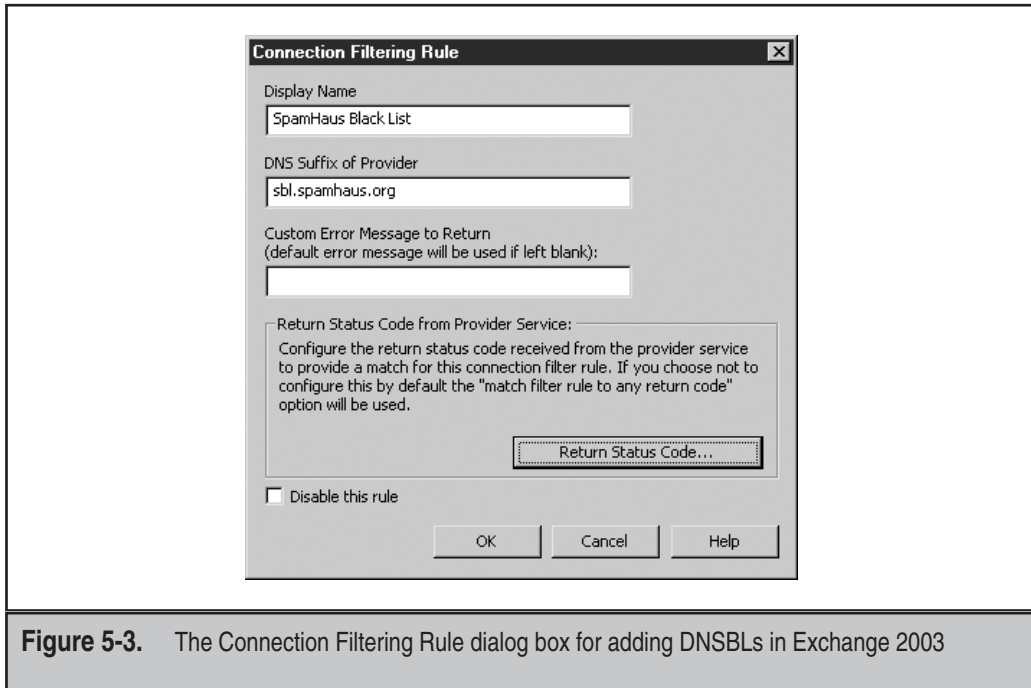


Figure 5-3. The Connection Filtering Rule dialog box for adding DNSBLs in Exchange 2003

127.0.0.0/8. The default (and first radio button) option is to use any response from the DNSBL as a positive response. The second radio button option lets you use a specific mask for the loopback range. The third option lets you select a specific value for the response. The last two options are useful for DNSBLs that have one Zone Server but specific “lists” underneath that server. For instance, a DNSBL’s open relay list might return a 127.0.0.5, their open proxy list 127.0.0.6, and their dial-up list 127.0.0.7. This lets you create separate rules for each one of those lists.

Returning to the Message Delivery Properties window (Figure 5-2), you will see a button labeled Exception. Clicking this will bring up the Block List Service Configuration Settings dialog box, as shown in Figure 5-5. Here you can add local e-mail addresses to which you do not want the blacklists to apply. This is especially useful for addresses that must accept mail delivery, such as postmaster or abuse, and for addresses that concern you because they might lose mail—a sales address, for example.

Other options on the Message Delivery Properties window include the ability to allow and deny certain IP addresses or IP subnets from sending you mail. These are equivalent to local whitelists and blacklists and are configured in dialog boxes that open by clicking the Accept and Deny buttons, respectively. Note that these settings override the DNSBL settings, and that addresses or networks under the Accept list will override those under the Deny list.

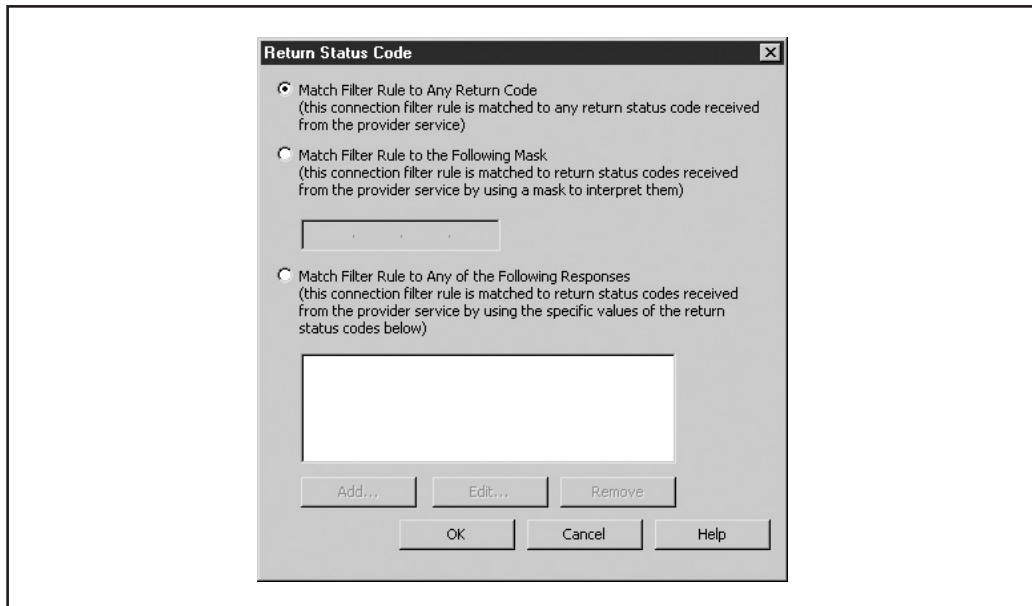


Figure 5-4. The Return Status Code dialog box lets you modify expected responses from DNSBLs under Exchange 2003.



Figure 5-5. The Block List Service Configuration Settings dialog box lets you exclude local users from having their inbound mail checked against the blacklists in Exchange 2003.

Once you configured your rules and have clicked OK, you're not quite done yet. You still need to apply the filters to your virtual SMTP servers. Back under the Tree tab of the Exchange System Manager, open the Servers folder and go to Servers | <Your System Name> | SMTP | Default SMTP Virtual Server, as shown in Figure 5-6.

Right-click Default SMTP Virtual Server and select Properties. You will then see the Default SMTP Virtual Server Properties window, as shown in Figure 5-7. Click the Advanced button.

The Advanced button will take you to the Advanced window, which lists the IP identities for the virtual server and tells you whether or not it's filtered. This is shown in Figure 5-8. Select the identity of the server to which you want to apply the filter and click the Edit button.

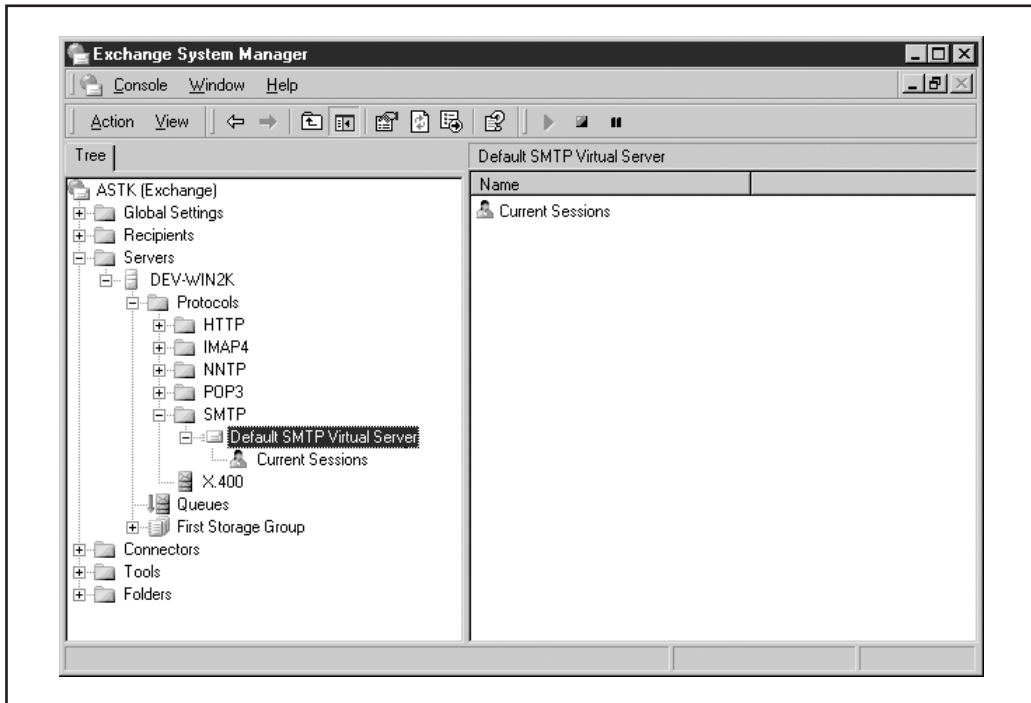


Figure 5-6. Location of Default SMTP Virtual Server in the Exchange System Manager on Exchange 2003

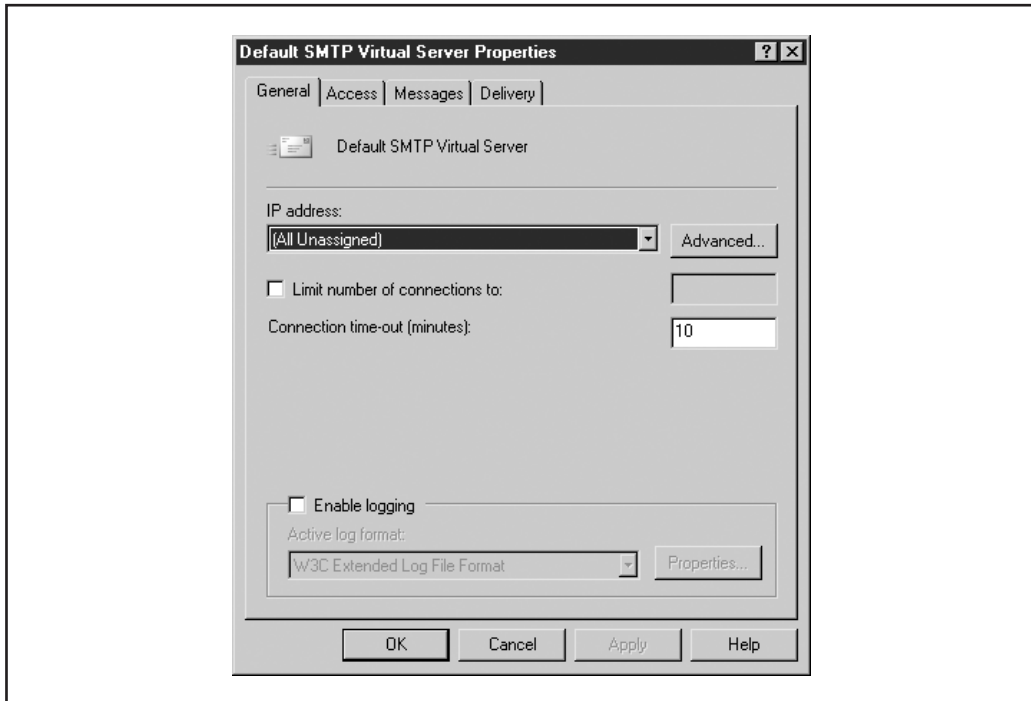


Figure 5-7. The Default SMTP Virtual Server Properties window in Exchange 2003

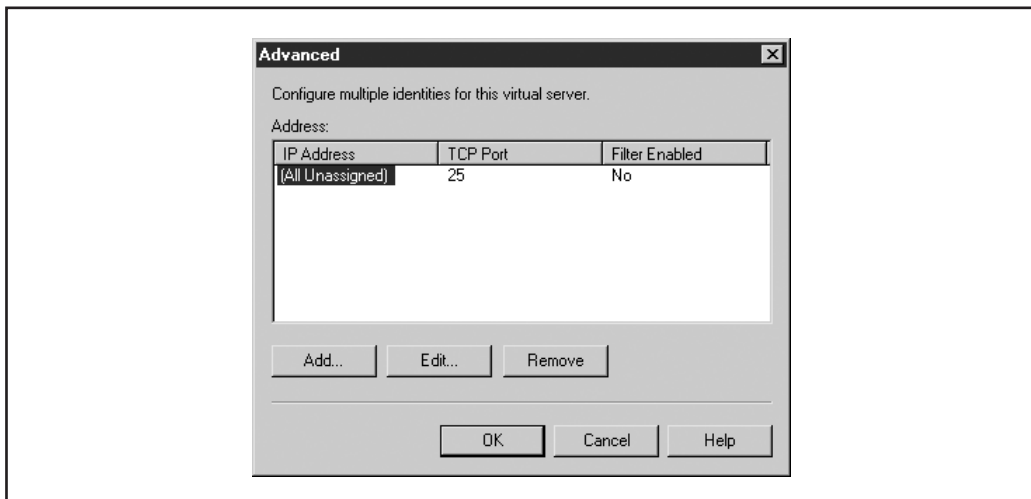


Figure 5-8. The Advanced properties window tells you the filtering status of your virtual server identities in Exchange 2003.

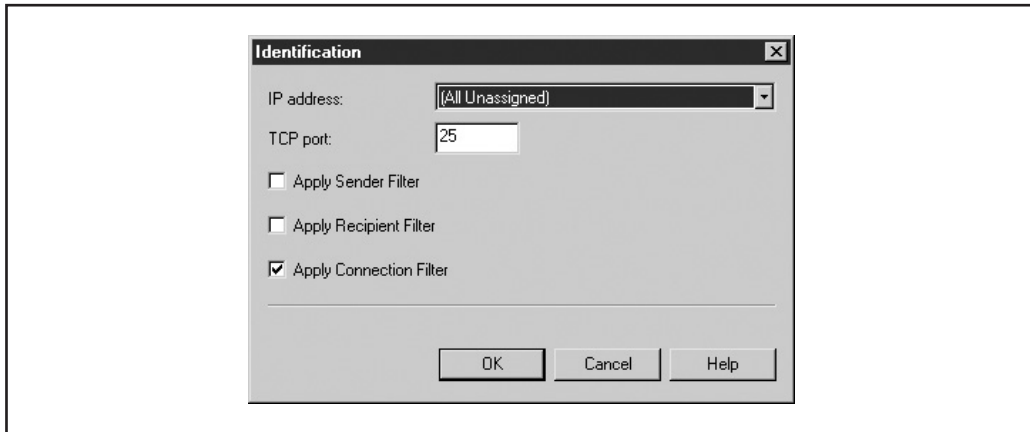


Figure 5-9. The Identification dialog box is where you apply DNSBL filters to a virtual server in Exchange 2003.

After clicking the Edit button, the Identification dialog box will appear. The only item that concerns the filters on this dialog is the bottommost check box, Apply Connection Filters, as shown in Figure 5-9. Check this box, and then click OK to apply this setting. Verify it is enabled in the Advanced properties window. You must follow these steps for each virtual server to which you want to apply your DNSBL filters.

SUMMARY

DNS Blacklists help you reduce spam by allowing you to reject, tag, or score e-mail that comes from known or potential spam sources. It gives you the benefit of the experience and resources of others. The primary pitfall of a DNSBL is the relatively high chance of gross false positives—more than most other anti-spam solutions. Entire networks of e-mail could potentially be lost should a major mail server or service reach a DNSBL. Likewise, you're placing a high degree of trust in the people running and contributing to the DNSBLs.

By understanding the technology and philosophy behind individual blacklists, you'll be able to choose those that suit your needs. Most e-mail servers and anti-spam software support DNSBLs either natively or through third-party add-ons, so there's no reason why DNSBLs shouldn't be a part of your anti-spam strategy. Just don't rely on them as your primary means to thwart spam.