



realtimepublishers.comtm

The Administrator Shortcut Guidetm To



Blocking Spam with Sender Validation



SpamLionTM
Anti Spam Gateway

Alan Sugano

Introduction

By Sean Daily, Series Editor

Welcome to *The Administrator Shortcut Guide to Blocking Spam with Sender Validation!*

The book you are about to read represents an entirely new modality of book publishing and a major first in the publishing industry. The founding concept behind Realtimepublishers.com is the idea of providing readers with high-quality books about today's most critical IT topics—at no cost to the reader. Although this may sound like a somewhat impossible feat to achieve, it is made possible through the vision and generosity of corporate sponsors such as SpamLion, who agree to bear the book's production expenses and host the book on its Web site for the benefit of its Web site visitors.

It should be pointed out that the free nature of these books does not in any way diminish their quality. Without reservation, I can tell you that this book is the equivalent of any similar printed book you might find at your local bookstore (with the notable exception that it won't cost you \$30 to \$80). In addition to the free nature of the books, this publishing model provides other significant benefits. For example, the electronic nature of this eBook makes events such as chapter updates and additions, or the release of a new edition of the book possible to achieve in a far shorter timeframe than is possible with printed books. Because we publish our titles in “real-time”—that is, as chapters are written or revised by the author—you benefit from receiving the information immediately rather than having to wait months or years to receive a complete product.

Finally, I'd like to note that although it is true that the sponsor's Web site is the exclusive online location of the book, this book is by no means a paid advertisement. Realtimepublishers is an independent publishing company and maintains, by written agreement with the sponsor, 100% editorial control over the content of our titles. However, by hosting this information, SpamLion has set itself apart from its competitors by providing real value to its customers and transforming its site into a true technical resource library—not just a place to learn about its company and products. It is my opinion that this system of content delivery is not only of immeasurable value to readers, but represents the future of book publishing.

As series editor, it is my *raison d'être* to locate and work only with the industry's leading authors and editors, and publish books that help IT personnel, IT managers, and users to do their everyday jobs. To that end, I encourage and welcome your feedback on this or any other book in the Realtimepublishers.com series. If you would like to submit a comment, question, or suggestion, please do so by sending an email to feedback@realtimepublishers.com, leaving feedback on our Web site at www.realtimepublishers.com, or calling us at (707) 539-5280.

Thanks for reading, and enjoy!

Sean Daily

Series Editor

Introduction.....	i
Chapter 1: Spam and Spam Filtering Methods.....	1
A Brief Summary of Spam	2
Why Does Spam Exist?	2
How Spammers Get Your Email Address	3
Anti-Spam Legislation.....	4
Traditional Spam Blocking Methods.....	4
Keyword Searching.....	4
Advantages.....	5
Disadvantages	5
ORDB Checking.....	6
Advantages.....	7
Disadvantages	7
Whitelists	8
Advantages.....	8
Disadvantages	8
Blacklists.....	9
Advantages.....	9
Disadvantages	9
MX Record Lookup	9
Advantages.....	9
Disadvantages	10
Heuristics and Bayesian Filtering.....	10
Advantages.....	11
Disadvantages	11
Sender Validation.....	11
Advantages.....	15
Disadvantages	16
Summary.....	17
Chapter 2: Sender Validation Solutions.....	18
Client-Based Sender Validation Solutions.....	18
Sender Validation Desktop Software.....	18

Advantages.....	20
Disadvantages	20
Sender Validation Desktop Services.....	21
Advantages.....	22
Disadvantages	22
Sender Validation Integrated with the Mail Account	22
Advantages.....	24
Disadvantages	24
Server-Based Sender Validation Solutions.....	24
Sender Validation Business Server Software	25
Advantages.....	26
Disadvantages	26
Sender Validation Business Services.....	27
Advantages.....	28
Disadvantages	28
Sender Validation Planning and Implementation	28
Evaluating Sender Validation in Your Environment	29
The Right Sender Validation Solution for Your Company	29
Necessary Steps to Implement a Sender Validation Service	33
Necessary Steps to Implement a Sender Validation Dedicated Server.....	34
Summary	35
Chapter 3: Implementing a Sender Validation Solution in Your Company	36
Cost Justification Compared with Other Methods.....	36
Software Cost.....	37
Implementation Benefits and Cost Justification	37
Reduced Storage on Your Internal Mail Server.....	38
Ongoing Administrative Costs.....	39
Minimal User Support Costs.....	39
Implementation Steps for a Sender Validation Service	40
Select a Sender Validation Service	41
Select an Implementation Strategy	41
Train End Users	41
Select an Initial Implementation Date.....	41

Upload a Pre-Approved Senders List	41
Change Your MX Record	42
Follow Up	42
Set Up a Quarantine Period.....	42
Sender Validation Internal Server Implementation Steps.....	42
Select a Sender Validation Package.....	43
Select an Implementation Strategy	43
Plan End-User Training	43
Decide on Sender Validation Server Placement	44
Purchase the Server Hardware	44
Install the OS on the Server	44
Install the Sender Validation Software on the Server	44
Select an Initial Implementation Date(s)	45
Upload a Pre-Approved Senders List	45
Reconfigure the Firewall.....	45
Make Mail Server Modifications	50
Activate the Sender Validation Server.....	50
Test the Sender Validation Server	51
Establish a Quarantine Period.....	51
Backup	51
Post Installation Tasks and Best Practices	52
Dealing with eCommerce and Other Legitimate First-Contact Situations	53
Troubleshooting	53
Summary	54

Copyright Statement

© 2004 Realtimedpublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimedpublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimedpublishers.com, Inc or its web site sponsors. In no event shall Realtimedpublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimedpublishers.com and the Realtimedpublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimedpublishers.com, please contact us via e-mail at info@realtimedpublishers.com.

Chapter 1: Spam and Spam Filtering Methods

Spam—everyone hates it, and it has reached epidemic proportions over the past year. Some estimates list the spam rate as high as 70 percent of all Internet mail traffic. Spam clogs up Internet WAN lines and consumes a significant amount of a user’s day. If you have reached the point at which spam is annoying enough to do something about, this guide will help you do so by focusing on the following topics:

- Existing anti-spam technologies
- Enterprise-wide spam solutions
- Spam filtering topologies
- Spam product selection, implementation steps, cost justification, Return on Investment (ROI), and integration with existing mail packages
- Spam filtering add-ons, estimated costs, and implementation pitfalls

Installing spam filtering software on a single workstation is a fairly simple task; however, implementing an enterprise-wide spam filtering solution requires careful evaluation and planning. You can expect difficulties—particularly false positives—when implementing any spam solution. In addition to being prepared for these considerations, you need to be aware of and plan for ongoing maintenance, which can be a hidden cost when implementing a spam solution.

There are many methods to block spam:

- Keyword filtering
- Open Relay Database (ORDB) checking
- Whitelists
- Blacklists
- Mail Exchange (MX) record lookups
- Heuristics
- Sender validation

However, only sender validation holds the promise of blocking 100 percent of spam. Sender validation has been around for quite some time and has had success in the Post Office Protocol (POP3) market. The concept of sender validation is very simple. If a user that is on your “approved senders” list sends you a message, you get the message. If the user is not on the list, the message is quarantined. Most sender validation spam solutions deal with individual POP3 mailboxes. Although these individual solutions work well, such has not been the case for past network enterprise deployments of sender validation. Sender validation has been criticized as an undesirable solution for fighting spam in enterprise environments. To avoid any problems and benefit from the 100 percent blocking power of sender validation in an enterprise environment, simply select a vendor that has a mature sender validation solution. In this guide, we’ll examine how to avoid the pitfalls of sender validation and implement this solution to cut spam to zero.

Fortunately, implementing an anti-spam solution is one of the easiest IT projects to cost justify. Imagine the productivity savings each user will experience if their spam is cut to zero! Typically even a small company can recoup their investment of an anti-spam solution in as little as 2 months. For larger companies, the cost recovery is even faster. Thus, sender validation is a solution that sells itself. Before we jump into how to begin saving money through sender validation, let's briefly establish a foundation of spam history and terminology.

A Brief Summary of Spam

Everyone knows what spam is; we've all received it. It's the automated mass email of advertisements and other annoying email messages. Just as important as defining what spam is, is to define what spam is not: a virus, identity theft, or instant messaging.

Why Does Spam Exist?

Spam exists because it works. When compared with snail mail junk mail campaigns, spam has significantly lower costs. Consider the example that Table 1.1 shows.

Traditional Mail Campaign	Spam Email Campaign
Cost per piece \$1.37	Cost per piece \$0.001
Mail 10,000 pieces	Email 1,000,000
Total cost is \$13,700	Total cost is \$1000
Hit rate is 2 percent	Hit rate is .02 percent
Total hits of 200 at \$68.50 each	Total hits of 200 at \$5 each


Table 1.1: Traditional mail vs. a spam email campaign.

From this very simple example, you can see that spam campaigns typically cost much less per hit than a traditional direct mail piece. But because the hit rate is much lower (in this example 100 times lower) than a direct mail piece, spammers must send out significantly more pieces to achieve the desired number of hits. Thus the reason that spammers use the "shotgun" approach in their mail campaigns—the cost per piece is almost zero, so spammers can afford to send their message to any email address on which they can get their hands. They don't bother trying to target their lists for specific groups that might be interested in the product. Spam is all about volume; the more messages sent, the better chance of receiving a hit.

Although spammers closely guard their hit ratios, they are making money. However, they must annoy a significant amount of the population to get the desired number of hits—before I implemented a spam solution, I typically received 200 to 300 messages per day. Spam has grown to a point at which both end users and organizations are willing to invest in a solution to stop the spam and recoup valuable lost productivity.

How Spammers Get Your Email Address

Spammers use *email harvesting* to continually get new email addresses. They use harvesting spiders/programs such as Atomic Harvester III, Email Marketing, and Text Bomber that monitor the Internet looking for new email addresses to gather. These programs are capable of gathering email addresses on specific Web sites, can target users in specific geographic areas, can target users in specific newsgroups and chat rooms, and can spoof IP addresses of bulk email servers.

 For more information about the capabilities of spam harvesting programs, check out <http://www.emilemail.com/>.

One of the more covert harvesting programs uses EMAIL_ID, which will capture your address when you simply visit a site by tricking your browser into giving your name and email address. If the security level on Microsoft Internet Explorer (IE) is set to the default level, you should receive a warning message before this information is submitted.


Spammers might also attempt to guess your email address by using a dictionary/directory attack. This type of attack simply runs down a list of names and tries each one until it gets a hit. When a hit is determined, the spammer exploits the entire domain name by following the naming convention (for example, <first_initial><last_name>@<domain_name>) for email addresses in the domain. Dictionary attacks are common on hotmail.com, msn.com, and other widely used email domains because of their mail volume and number of users. Spammers hit these sites continuously 24 hours a day, 7 days a week with dictionary attacks. When a hit is identified, the email address is recorded, and this list is sold to other spammers. These sites are continuously under attack, so you are almost guaranteed to receive spam if you set up a mail account here.

If your Internet connection slows suddenly you might be under a dictionary or Denial of Service (DoS) attack. Examine the log in your firewall to attempt to identify the source of the attack. If possible, use your firewall to block the IP address(es) from which the attack originates, and contact your ISP and ask them to block the IP address at the backbone to prevent further problems.

In a recent Federal Trade Commission (FTC) study, 86 percent of email addresses that were posted on Web pages, chat rooms, and message boards received spam. One email address received spam 9 minutes after a message was posted in a chat room!

 For more information, refer to <http://www.ftc.gov/bcp/online/pubs/alerts/spamalert.htm>.

Thus, never have a direct link from your Web page to a real person's email address. Use a generic email address such as *info@<domain_name>*. Spammers tend to leave these generic addresses alone, and if they do receive spam, the address can be easily changed. Alternatively, your company can create a Web form (rather than use a generic email address).

 I've had mixed success submitting an opt-out request to spam mail. If the spam appears to come from a legitimate source, I've had better luck with the opt-out request. Be aware, however, that replying to a spam mail verifies to spammers that they've reached a real person. Use the opt-out feature at your own risk.

Anti-Spam Legislation

As a result of the spam problem, 30 states within the United States have passed laws that make it illegal to send spam, but enforcing the laws inter-state and even within the same state (not to mention internationally) is difficult if not impossible. In June 2003, the Burns-Wyden bill passed. This bill legislates that spammers can face up to 1 year in prison and a maximum fine of 1 million dollars. Although anti-spam legislation will help, it probably will not solve the problem. Law enforcement has higher priorities within the IT industry such as catching virus creators and cyber terrorists. Thus, rather than wait for a legal remedy to this problem, the only effective solution is to use a spam blocking tool.

Traditional Spam Blocking Methods

There are quite a few anti-spam software packages on the market, most of which use a combination of spam blocking methods to reduce the amount of spam in a user's mailbox. However, spammers are constantly developing new methods of bypassing spam filters. Thus, except for sender validation solutions (which don't necessitate ongoing updates), spam filtering solution vendors must develop additional methods of blocking spam to keep up with the spammers. Let's examine these methods and their advantages and disadvantages.


Keyword Searching

For keyword searching, the anti-spam software looks for specific words or phrases in an email. In you're in the market for this type of solution, look for a package that supports keyword phrases, keyword conditions, and keyword searching in either the subject or body of the email message. Keyword searching can reduce the amount of spam by performing a search for words that are likely to be included in the spam message (for example, Viagra, refinance, and mortgage). Phrase searching with conditions will give you more flexibility to search for items such as "need cash" and "refinance." This functionality provides a finer degree of control when searching for keywords and should help reduce false positives. Some spam filter vendors allow you to update your keyword searches based on the most current spam messages on the Internet.

Advantages

If the message you receive has a consistent word or phrase, keyword searching is an effective method of blocking spam. It is very useful for blocking other unwanted messages that contain viruses, such as the Sobig.F worm, that use the several phrases as the following email message shows:

Re: Approved
Re: Re: My details
Re: Thank you!
Re: That movie
Re: Wicked screensaver
Re: Your application
Thank you!
Your details

 Although anti-spam software can block unwanted messages that contain viruses, do not rely anti-spam software as your only virus email scanner. Purchase a virus scan option with the anti-spam package or install a dedicated email virus scanner on your email server.

Disadvantages

Unfortunately, this method requires that you receive at least one email with a consistent keyword before you can block future messages with a keyword search. You must manually maintain the keyword list as new spam messages are received, unless the spam filtering vendor supplies updates for you. In addition, this method has the potential to consume considerable resources on the server—for example, if you perform searches on the message body versus the subject line or add keywords to the search list, more resources will be consumed on the server. On a heavily loaded server, some messages can get through the keyword search. Smart spammers randomize the words in the subject and message body in an attempt to bypass the keyword filter. Finally, keyword searches have the potential to cause many false positives depending on the type of mail your company receives.

ORDB Checking

A mail server configured as an open relay allows spammers to bounce messages off the mail server to send the spammers' messages. Some packages can perform an ORDB check to determine whether a message was received from a mail server that is identified as an open relay (see Figure 1.1).

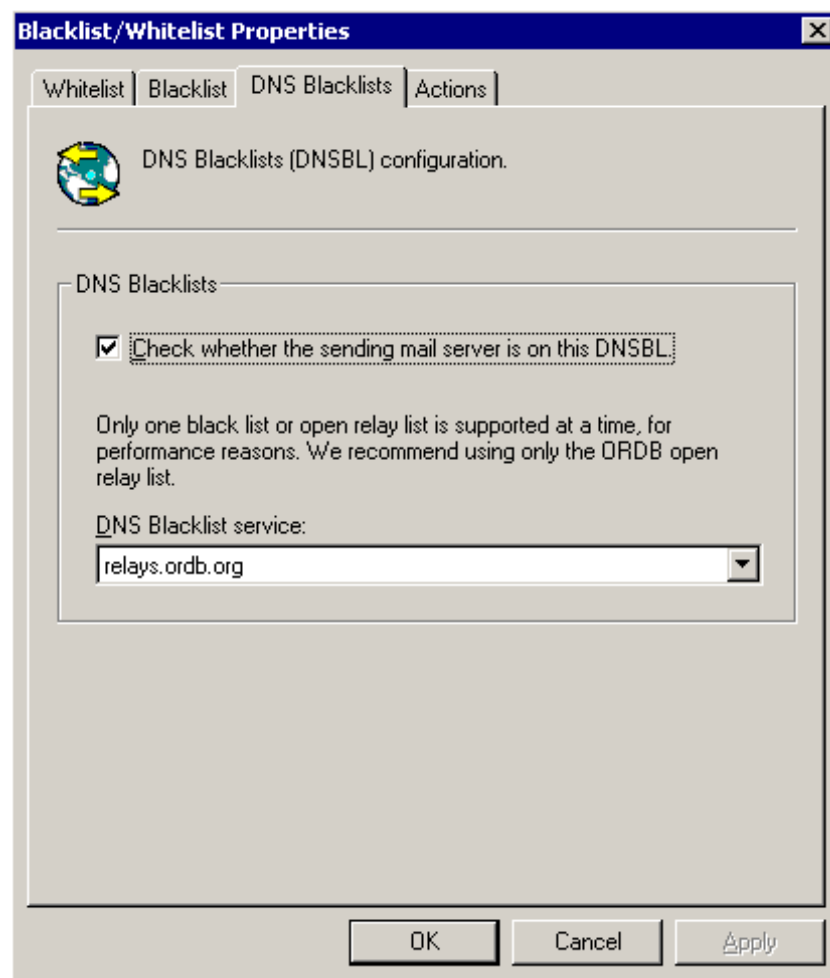



Figure 1.1: A DNS Blacklists screenshot from GFI Mail Essentials.

 If your server is an open relay, it is simply a matter of time before it is listed in one or more of these databases:

<http://abuse.easynet.nl/blackholes.html>

<http://www.delink.net/>

<http://dnsbl.njabl.org/>

<http://dsbl.org/main>

<http://ordb.org/>

For a comprehensive list, refer to <http://www.declude.com/junkmail/support/ip4r.htm> and <http://www.moensted.dk/spam/>.

How to Determine Whether Your Server is an Open Relay

Many of these sites can test whether your server is an open relay. When you bring up a new mail server, it is a good idea to test it to verify that the server is not open. If your server is marked as an open relay, you must first close it, then submit it for retesting. If you are running an earlier mail package (for example, Microsoft Exchange 5.0, Novell GroupWise 5.2) that cannot be shut down as an open relay, take a look at the anti-spam features of your firewall. Some firewalls have anti-spam features built-in to their Simple Mail Transfer Protocol (SMTP) daemons that can help close the open relay. Another option is to upgrade your email software to a version that does not allow open mail relaying. An irony of the ORDB is that it provides a convenient list for spammers to relay their messages—the opposite of what the ORDB is trying to prevent. After the server has been retested and is no longer an open relay, it should be removed from the database. If the server is marked as an open relay, it might be listed in multiple databases. If such is the case, you must submit a removal request from each database on which the server is listed. The response time for a removal request varies depending on the database list.

Sometimes after the server is removed from an ORDB, the server might still have difficulty sending mail messages to one or more domains. If all else fails, you can change the external IP address and MX record of the server to bypass this problem. Typically, the ORDBs only list specific IP addresses rather than ranges of IP addresses. Thus, changing your mail server address is a simple workaround if your mail server is identified as an open relay (even though it is not open anymore).

A quick way to test this workaround is to change the outside address of your firewall, then try to send mail to the problem domain(s). If you are successful, issue an MX record change to the ISP that hosts your domain. If this workaround does not work, don't bother with the MX record change—the sending problem lies somewhere else, possibly with DNS, the mail server, or the message is infected with a virus. Be aware that you will temporarily take down your incoming mail while you run this test until you either update your MX record and it propagates throughout the Internet or change the firewall back to its original address. For this reason, it is a good idea to have extra IP addresses when ordering your DSL, T1, or broadband connection from your ISP even if you plan to use Network Address Translation (NAT) on the firewall. If you decide to use this approach, make absolutely sure that your relay is closed before changing the IP address; otherwise, you will end up on the ORDB again.

Advantages

Checking whether an email message came from a server marked as an open relay can block as much as 50 percent of your spam. Another benefit is that once this method has been configured, there is no on-going maintenance.

Disadvantages

Checking an ORDB consumes bandwidth because a lookup must be performed for each received message. If you use this method, rely on one of the larger ORDBs such as ordb.org. Open relay checking can potentially generate false positives because the relay might already be closed. Unfortunately spammers are getting smarter in their relaying methods. In the past they would find an open relay and exploit it until it was marked as an open relay. Now they hop from server to server and relay a smaller number of messages. This process makes it very difficult to identify the mail server as an open relay. Going forward, this anti-spam method will become less effective.

Some mail filtering services maintain their own “real-time” open relay list that is continually updated. When a mail server appears to deliver spam, the relay is tested periodically to verify that it is still sending spam. Once the server stops sending spam, it is automatically taken off the open relay list. This approach was developed to catch the technically savvy spammer that hops from open relay to open relay to avoid detection.

Whitelists

If a message is received from an email address or domain on a whitelist, the message is delivered to the user. If you’re shopping for this functionality, look for a package that can support entire domains with one whitelist entry such as `*@<whitelist_domain.com>` (instead of separately listing individual users in the whitelist). This feature is very useful for users who correspond with multiple users in another company regularly. Of course, you don’t want to open an entire domain—such as `*@aol.com`, `*@yahoo.com`, and `*@hotmail.com`—from which users will receive spam on a regular basis, but for other domains, this feature can save a lot of administrative time.

Typically, a whitelist entry overrides conflicting configurations. For example, if a message is received from a user that is on the whitelist, but the message originated from a server marked as an open relay, the message is allowed through. Some software packages can automatically add users to a whitelist when an internal user sends mail to that person. However, this feature can be undesirable, especially if a user decides to opt-out of a mailing list. By replying to the mailing list message, the opt-out address is automatically added to the whitelist. If you decide to turn on the auto-add whitelist feature, make sure your users do not reply to such opt-out email messages. Alternatively, IT staff can simply remove an unwanted address from the whitelist. Many solutions offer the choice of per-person or company-wide whitelists, which enable administrators to decide whether users’ auto-add feature will affect other users.

Advantages

Preloading a whitelist of approved senders will reduce the number of false positives when implementing anti-spam software. For this reason, preloading this list is an integral part of any whitelist implementation. At least, give the whitelist system time to “learn” who your users send mail to before turning on the spam-blocking feature. If the whitelist overrides other filter values, you can use the whitelist and blacklist in combination to filter out spam. (I’ll discuss blacklists in the following section.) For example, you can block an entire domain, such as `*@hotmail.com`, in the blacklist, then selectively list email addresses in the hotmail.com domain for messages from senders you want to pass through the spam-filtering software.

Disadvantages

If you do not implement the auto-add whitelist feature, this list must be maintained manually. Even if you preload a whitelist, expect to receive several false positives when implementing anti-spam software. If you’re running Microsoft Outlook, you can export all the email addresses in the contacts list for each user, consolidate the list, format it based on the spam-filtering software requirements, then import the list into the whitelist. The number of manually added whitelist entries should taper off after the package is up and running for a few weeks—especially if you enable an auto-add feature. Both the whitelists and blacklists are responsible for the majority of the ongoing maintenance for anti-spam packages that use these methods of blocking spam.

Blacklists

Blacklists work just the opposite of whitelists—if a message is received from an email address or domain on a blacklist, the message is rejected by the server. Blacklists have the same drawbacks as keyword searching, because you usually have to receive a spam message before you can block it (unless, of course, you already blocked the entire domain).

Advantages

If spam is consistently sent from a single email address, blacklists are an effective spam-fighting tool. However, more than 75 percent of spam is from a one-time use address, blacklists alone will help to protect against only a quarter of the spam. As I previously mentioned, you can use the blacklist and whitelist combination to block an entire domain, then only let selected messages through the spam filter. If you implement a server-based spam-filtering solution, you need to enter the blacklisted address only once on the server; after the address is blacklisted, all mail received from this address is automatically blocked at the server level.

Disadvantages

The biggest disadvantage of blacklists is ongoing maintenance. As new messages appear, the administrator must add the sender's name to the blacklist. Most spammers use “throwaway” email address such as spam123@yahoo.com. Once the email service recognizes the sender as a spammer, the account is deactivated. However, because many spammers don't even bother acquiring an email account in the first place—a recent study showed that more than 76 percent of spam is from nonexistent accounts—deactivation of a spammer's account is of little consequence. Maintaining a list of all these one-time use accounts causes exhaustive and excessive blacklist checking by the server—particularly considering that most spammer addresses are only used once. Thus, with blacklists, you're always one step behind the spammer, so a subscription to a blacklisting company is required to make this method an effective spam-fighting tool.

MX Record Lookup

MX record or a reverse DNS record lookup performs a DNS query on the sender's domain. If the sender's domain matches the MX record IP address of the server, the mail is accepted. If the IP address does not match, the message is rejected.

Advantages

This approach can work well if a sender's domain name has been spoofed by a spammer. In such a case, the server would know that the message is not coming from the legitimate contact.

Disadvantages

This approach has the potential to create many false positives. The following scenarios can cause a false positive:

- Incorrect or missing reverse DNS record—Many companies do not bother to have a reverse record created when establishing the MX record for their mail server.
- Multiple mail servers—Larger companies or ISPs can have multiple mail servers for their domains. When the server performs a reverse lookup, the server might not get all mail server IP addresses for the domain, which can cause a false positive because the IP address of the sender's server might not match the IP address of the reverse lookup.

For these reasons, I suggest using other methods for spam blocking.

Heuristics and Bayesian Filtering

Heuristics and Bayesian filtering is one of the more recent methods developed to block spam. The software gathers statistics about the type of message received, then makes a judgment call about whether the message is spam. To make this determination, some software packages use a point scoring system and others use custom algorithms. This method can be a very effective weapon against spam.

Heuristics and Bayesian filtering works like a blackjack player who is counting cards. A card counter knows that the deck is in his or her favor when a series of low cards appears because this means that the deck is “ten rich,” increasing the probability that the dealer will bust if the dealer must draw a card. Heuristics and Bayesian filtering similarly looks at words in email messages that are already marked as spam, then compares how often key words appear in an incoming email to estimate the probability that the message is spam. Generally, more recent data is more heavily weighted and email keywords are continually updated with new and current information. This system gives heuristics and Bayesian filtering the advantage of becoming somewhat self-maintaining.

If you're considering an heuristics and Bayesian filtering solution, consider a filter that looks at outgoing email to reduce the amount of false positives. For example, if you work for a refinance company and the word mortgage appears quite frequently in your outgoing emails, you want to ensure that messages that contain mortgage aren't blocked. In this particular case, the word mortgage will not have such a heavy weight for incoming mail because it occurs quite frequently in the company's outgoing mail. This analysis of outgoing email will reduce the amount of false positives.

Because heuristics and Bayesian filtering typically takes the whole message into account, it can usually catch misspelled words such as s*e*x or v-i-a-g-r-a. In fact, these misspelled words almost guarantee that the message is spam because a legitimate email will most likely never spell words in this manner.

Advantages

The biggest selling point for heuristics and Bayesian filtering is that this solution is very low maintenance. Heuristics and Bayesian filtering constantly gathers information about incoming mail and updates statistics on an ongoing basis. Because it typically only looks at mail sent and received by the company, the statistics are custom-tailored for the company's email. Usually these statistics are more heavily weighted on the most recent data. Some companies claim they can block out as much as 99 percent of spam with a very low percentage of false positives by using heuristics and Bayesian filtering.

Disadvantages

Heuristics and Bayesian filtering is only as good as the engine/algorithm making the spam judgment call. Typically, the entire message is evaluated, which results in an additional load on the email server assuming the heuristics and Bayesian filtering engine is installed on the same machine as the mail server. On a heavily loaded server, this spam-blocking method can cause performance issues.

In addition, after the heuristics and Bayesian filtering analysis, each message is typically assigned a probability ranging from 0 to 100 percent that the message is spam. This probability must be fine-tuned over time. Set the threshold too high, and too much spam gets through. Set the threshold too low, and you generate many false positives. Refer to the software documentation for a recommended initial setting, then fine-tune this setting based on your company's requirements. Because every company's email is different, you must use trial and error to determine the best setting for your company. Also, because heuristics and Bayesian filtering has the potential for generating false positives, look for a package that also supports a whitelist or some other method of receiving a legitimate message that was incorrectly marked as spam.

Sender Validation

At a basic level, sender validation works by letting mail through if the sender is on an approved list and rejecting the mail if the sender is not on the list. Think of sender validation as an "intelligent whitelist." Once a sender is placed on the approved list, the mail server will accept mail from this address. The concept is simple, but it is the management of the approved list and a smooth validation process that are keys to a successful sender validation anti-spam solution. Most corporate sender validation packages work like the flowchart that Figure 1.2 shows.

Sender Validation Process




Figure 1.2: A flowchart that illustrates the sender validation process.

Sender Validation Solutions of the Past

Historically, desktop/POP3 versions of sender validation have had the most success in the spam world. They are the easiest to implement and evaluate for the individual user; however, they are generally impractical within a business environment. The advantage of such solutions is that, compared with an internal enterprise solution, the software development cycle for a desktop sender validation solution is relatively short. It is fairly easy to evaluate and install because the sender validation server and software are usually located off-site.

However, some sender validation POP3 solutions are not online all the time, which can cause mail delivery problems. Depending on the solution, the validation process can be cumbersome and take a long time (days) to complete. It is possible to encounter a deadlock situation when both the sender and receiver have a sender validation spam protection for their mail. Sometimes the sender validation solution will not send an NDR to a legitimate sender, so the sender assumes their mail was delivered, but it wasn't. Make sure you can manually add senders to the approved list to avoid a deadlock situation. Some of the earlier attempts at sender validation were built by amateurs and lack the stability and features of a proper corporate sender validation solution.


Early sender validation enterprise solutions (circa 1995) lacked the stability and functionality that corporate users required. The advantage of such solutions is that the sender validation server is internal, so it is always online. However, many of the early attempts at sender validation were immature products, which resulted in a bad reputation for sender validation. These immature products were costly to evaluate because they typically required a dedicated internal server and complete implementation of the package just to evaluate it (a significant investment in both hardware and time for IT staff to purchase, install, configure, and sometimes develop the software for the sender validation server). With some early sender validation solutions, it was possible to encounter a deadlock situation between companies. Like the POP3 solution, you sometimes had the problem of lost NDRs, so a legitimate sender assumed you received their message when you actually had not. Some of the earlier sender validation solutions were developed in-house by internal IT departments or individuals, which had mixed success rates. Some of these sender validation systems have matured and evolved into today's systems.


If you're considering a sender validation solution, look for the following features to ensure a successful implementation:

- Auto-learn outbound communication—Look for a sender validation package that monitors outbound communication. Ideally the package provides the option to auto-add a user to the approved list when an outgoing mail is sent to the user. This feature can dramatically reduce false positives. However, even with the auto-learning feature, expect false positives with any “first-contact” message sent through an sender validation solution. Typically, this happens only the first time a user purchases an item online and the new vendor sends a receipt or other notification of the purchase. If the vendor's address is not in the approved sender's list, the user will receive a false positive. The false positive typically occurs because the e-commerce vendor does not go through the validation process, and the user has not done prior business with that vendor. Once the first interaction has been manually approved by the user, this one-time false-positive situation does not recur.

- **Pre-load approved list**—The sender validation software should allow an administrator to upload a predefined list of approved senders prior to package implementation. Doing so will reduce the number of false positives and messages in quarantine when the package is first implemented. This reduction in false positives makes this feature mandatory for any successful sender validation implementation. Alternatively, look for a solution that provides an auto-learn feature, which could be enabled for a period of time (a few weeks to a month) with spam protection off. As I mentioned earlier, such an opportunity would allow the solution to “learn” who the company communicates with before activating the protection (although loading a customer and vendor contact list provides for a much faster deployment).
- **Removal of deadlocks**—Ensure that the sender validation solution offers some method for removing deadlocks (for example, through a user override of a sender’s address). A deadlock can occur if two users within companies that are using sender validation software send each other a message. Both systems send and wait for the other system to respond to the mail, creating a deadlock situation. Make sure that the deadlock removal process has been thoroughly tested in a corporate mail environment.
- **Approved list management**—Users should have the option to manually add and delete items from the approved list.
- **One time validation**—Once a sender is on the approved list, they should not have to validate again.
- **Approved list**—Any sender on the approved list should have their message delivered.
- **Flexible quarantine**—To reduce maintenance costs, the package should have the option to auto-delete a message in quarantine after a user-defined period of time if the sender does not validate. This feature will also reduce the storage requirements on the sender validation server. Messages that receive a non-delivery report (NDR) during the validation process should automatically be deleted from quarantine.
- **Flexible validation**—The validation process varies from package to package. Some sender validation packages require the sender to click an HTTP link, require a reply with certain text in the message, or require the user to enter a pass phrase to get on the approved senders list. Regardless of the method, look for a package that has flexibility in its validation process. Make sure that the validation process is compatible with any mail client and mail server. The validation process should make it extremely difficult for an automated mail system to complete the validation process. The user should be notified that certain mails are pending validation and have the option to override the sender validation filter. The entire validation process should be easy to use and understand by the sender to reduce the sender’s confusion and the number of false positives.
- **Backup flexibility**—The sender validation approved list is a key component to this spam-blocking method, so ensure that the sender validation package configuration can be easily backed up on a regular basis to allow for a graceful recovery in the event of a hardware failure. For larger implementations, look for flexibility in the database engine (such as the ability to use SQL Server or Oracle as the database back end). This feature is especially important for implementations in which the sender validation server will handle a very large number of users (20,000+ users).

- Load balancing/fault tolerance—Very large companies and those for which email is a mission-critical application should look for a package that supports load balancing and/or fault tolerance/failover between multiple servers. Even without this feature, make sure that it is easy to bypass the sender validation filtering server (typically an IP address change on the firewall) in case of a complete hardware failure on the server.

 Fault tolerance is a concern with any anti-spam solution. Using a secondary MX record to the ISP's backup-relay server is an excellent solution for fault tolerance.

 Vendors such as SpamLion and MailFrontier provide sender validation packages that offer all these necessary features.

Advantages

There are quite a few advantages of sender validation, especially when compared with other spam-filtering methods:

- More effective—Some sender validation implementations experience 100 percent reduction in their automated spam. Even if you don't achieve 100 percent reduction, sender validation will be significantly more effective than other spam-filtering methods.
- False positives—Sender validation packages that have an auto-learn feature will result in a lowered false positive rate compared with other methods. In addition, once a sender is on the approved list, a false positive will not be repeated. This functionality places a light load on the server for approved senders because the server simply performs a lookup on the sender's email address rather than a battery of tests to determine whether the message is spam.
- Low maintenance—Once sender validation is implemented, there is a lower maintenance rate than with other filtering solutions (particularly compared with a whitelist and blacklist implementation). Maintenance is lowered even further if a sender list is preloaded as part of the implementation process and the approved sender's list remains relatively constant.
- Easy deployment—Some sender validation solutions eliminate the need to deploy additional software at the desktop level; thus, if there is no deployment on the desktop, there is no ongoing software maintenance necessary at the desktop level. This feature makes future upgrades easier, because you only need to upgrade the server.
- Available as a service—Some sender validation corporate solutions are available as a service rather than an internal dedicated server solution, giving you the flexibility to implement the solution as a service or on a dedicated internal server.
- No compatibility issues—Most of the sender validation solutions are compatible with all types of mail servers. Their validation process has been well tested for a variety of users and corporate environments.

- User flexibility—Most sender validation solutions can be enabled for all or only a portion of users. Thus, these solutions can coexist with non-protected people or other spam-filtering solutions.
- Security and serviceability—Internal sender validation solutions do not expose the internal mail server to the Internet. This feature allows the IT staff to patch the sender validation server (the “external server”) without taking down the internal mail server.
- Client independent—Most sender validation solutions work at the server level, so the end user gets the benefit of sender validation regardless of the client (MAPI, wireless, Web-based, POP3) used to access the mail server.

Disadvantages

To truly benefit from sender validation, you need to be aware of the disadvantages of this spam-blocking method:

- Significant initial cost—The sender validation solutions that have the most flexibility and features typically require a dedicated server, while other spam-filtering solutions can be installed on an existing mail server. Dedicated server sender validation solutions require more setup time because the OS must be installed on the dedicated server. Some sender validation vendors are in the process of developing solutions that install on an existing mail server. Depending on the current load of your internal mail server, you might want to implement a dedicated server solution anyway. With increased mail traffic, storage requirements, virus scanning, instant messaging, and advanced groupware features, your server might already be severely taxed. On a cost-per-user basis, sender validation is less expensive for larger companies because the cost of an internal server is spread out over a larger number of users. For smaller companies, a service-based sender validation solution might be less expensive than a dedicated internal server solution. Although sender validation can have a higher initial cost for larger companies, the total cost of ownership (TCO) over the first year is lower than other solutions because of the lower ongoing maintenance costs and greater overall effectiveness.
- One-time false positives—Many e-commerce vendors do not respond to customer inquiries, which can lead to a false positive of the sender validation sender when conducting an initial purchase. An easy workaround for this issue is to train end users to add the e-commerce vendor’s email address to their approved list when they first purchase an item from a new vendor. Of course, these vendors only need to be validated the first time, so this disadvantage is only an issue for purchases from new vendors.
- Legitimate senders don’t validate—Some email users are not familiar with the sender validation process and therefore do not validate, causing their messages to bounce back. If the mail is legitimate, usually the skeptical sender will call the recipient to see why the message bounced. The recipient can then ensure the legitimacy of the validation process or simply add the senders name to the approved list. Because of this issue, the validation process should only be necessary once, be simple, and easy to understand even for the novice user.

- End-user training—Typically, sender validation solutions have a Web-based interface to manage messages in quarantine and the approved list. Most of these interfaces are easy to use; however, some resources should be budgeted for end-user training. This training can be a short seminar or a simple instruction manual about how to use the sender validation software. End users are usually open to learning about how to manage their message quarantine and approved list rather than having to deal with an overwhelming amount of spam.
- Spammers can validate—Although highly impractical and unlikely for the spammer, it is theoretically possible for them to go through the process of validation and get on a user's approved list. If such should occur, users should have the ability to remove an address from their approved lists.
- Dedicated server—If the sender validation implementation requires a dedicated internal server, this server is one more resource that must be maintained by the IT department. It must be backed up on a regular basis and receive the same care as any other server on the network.
- Constant new mail senders—If your company constantly receives mail from different users rather than repeat customers, sender validation is probably not the correct solution for your company. In such cases, ongoing maintenance will probably be higher with a sender validation solution than with other spam-filtering methods.

Summary

The spam problem becomes more of an issue everyday—spam exists because spammers make money doing it. The cost per piece of spam is dramatically lower than a traditional direct mail campaign. However, as we explored in this chapter, there are many spam-blocking methods available and being developed to combat this growing problem.

Each of these spam-fighting strategies has advantages and disadvantages. Often a combination of these strategies can provide a satisfactory solution for blocking spam. Among all of these methods, only one can potentially eliminate 100 percent of spam—sender validation. Although sender validation got a bad rap in the 1990s as a result of homegrown systems that lacked features and functionality and had poorly designed user interfaces, the current crop of sender validation solutions are ready for the corporate environment and have been fully tested and refined.

Spam robs a tremendous amount of time and resources from end users, IT staff, mail servers, WAN links, and storage requirements. The good news is that almost any solution will save your company money. Obviously, you want the best solution, and sender validation is a good fit for most companies. Once you've decided that sender validation is the right technology for your company, you must evaluate the advantages and disadvantages of each sender validation package. In the next chapter, we'll take a look at finding the right sender validation package for your company as well as how to evaluate the package you choose to ensure that it is the best solution for your environment.

Chapter 2: Sender Validation Solutions

In the first chapter, we explored the growing crisis of spam as well as the spam-blocking methods available and being developed to overcome this problem. In this chapter, we'll delve deeper into sender validation solutions—exploring both client-based and server-based solution options. Although I will briefly cover sender validation solutions for individual use, the chapter will focus on business solutions.

The first step in implementing a sender validation solution in your organization is to select a sender validation service or dedicated server. But which solution is right for your environment? We will identify the critical considerations to help you determine the answer. In addition, we'll explore the steps necessary to implement the sender validation solution you select. Let's start by taking a look at client-based sender validation solutions.

Client-Based Sender Validation Solutions

Quite a few sender validation solutions are available for the desktop. Your options in this arena include both services and software that integrates with your existing desktop mail software. The following sections explore these solutions.

Sender Validation Desktop Software

You install sender validation desktop software on each workstation, as this sender validation solution is targeted more towards the end user—rather than a business that has its own internal mail server. Make sure that you pick a package that is compatible with your mail client software. Sender validation desktop software is best suited for POP3 and individual mail users. Figure 2.1 shows how sender validation filtering is performed on each workstation.

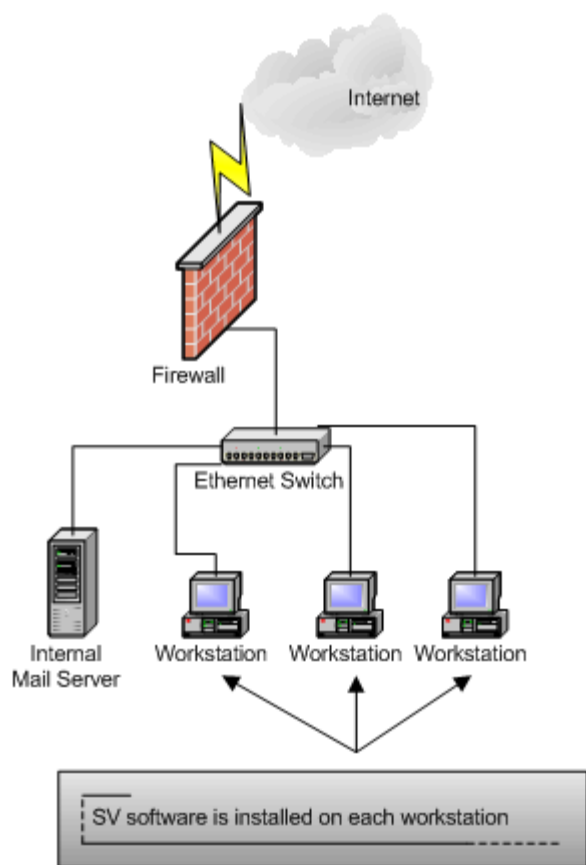



Figure 2.1: Sender validation filtering on each workstation.

The following list highlights desktop packages that use sender validation to fight spam:

- DigiPortal ChoiceMail (<http://www.digiportal.com/>) spam filtering software
- MailFrontier Matador (http://www.mailfrontier.com/products_matador.html) works with Outlook and Outlook Express
- MYmailSAFE (<http://www.mymailsafe.com>) offers both a consumer and business anti-spam solution
- Bongosoft (<http://www.bongosoft.com/>) is a newcomer in the anti-spam market
- Mail Wiper (<http://www.mailwiper.com>) provides a pop-up blocker with the purchase of the anti-spam software
- Spam Bully (<http://www.spambully.com>) integrates with Outlook and Outlook express

 Remember to pick a package that is compatible with your mail client!

Some of these packages use sender validation in conjunction with other spam filtering methods—such as blacklists, whitelist, and Bayesian filtering—to combat spam.

Advantages

Sender validation desktop software is a useful solution for individual users that do not connect to a corporate mail server and use a single email package to access their mailboxes. It is also useful for users who always access their mail from the same email client. This option is a fairly low-cost solution for the individual user.

Disadvantages

Sender validation desktop software has the following disadvantages:

- It offers no central management.
- Changes/upgrades must be installed individually on each workstation.
- Not a good solution for more than a few users—server-based solutions are usually more cost effective for larger installations.
- Many client-based sender validation solutions require that your desktop computer be running 24 × 7 so that the software can immediately notify the sender of questionable mail; otherwise, this notification is delayed until sender validation is restarted. This setup can cause validation delays of as many as 7 days if a user connects only once a day to send/receive mail.
- There is no “master list,” so each user must individually validate mail senders.

In addition to these drawbacks, the package must be compatible with your existing mail client—because this software is client specific, it does not work with a different client to access mail. This requirement is commonly a problem faced by users of Outlook Web Access (OWA) because the spam filtering software is not active when users access their mail through OWA. If such users decide to change their email client, they might need to also change their sender validation software.

Sender Validation Desktop Services

Some companies offer sender validation as a service. These services are also targeted more towards the end user. Figure 2.2 shows how a typical sender validation service is implemented on a workstation. Most services charge by mailbox and/or by megabytes of mail storage.




Figure 2.2: Sender validation filtering service on a workstation.

The following list highlights companies that provide sender validation desktop services:

- Mailblocks (<http://www.mailblocks.com>) can consolidate as many as 10 mailboxes
- USOpt (<http://www.usopt.com/>) offers spam and virus filtering services
- Spam Arrest (<http://spamarrest.com/products/individuals.jsp>) works with any POP3 mail account

Advantages

These services are very easy to set up. In addition, they require minimal cutover planning and are usually easy to disable if you don't like the service.

Disadvantages

A drawback to this solution is that sender validation desktop services are often more costly in the long run: the typical break even point is 2 to 3 years compared with installing sender validation software on your workstation. In addition, your sender validation list is off-site and is dependant upon the service's and your Internet connection. Also, a service provider will sometimes charge extra for more storage because they have to track your approved sender's list and mail. As a minor security concern, the sender validation service has the addresses to and from which you send and receive email.

Sender Validation Integrated with the Mail Account

Just like some ISPs provide integrated antivirus scanning as part of the email service, some ISPs and application service providers (ASPs) offer sender validation integrated into their email service. This option is a good fit for individual users, but typically does not work well for a business. Figure 2.3 shows how sender validation is integrated with an existing ISP-provided mail account.

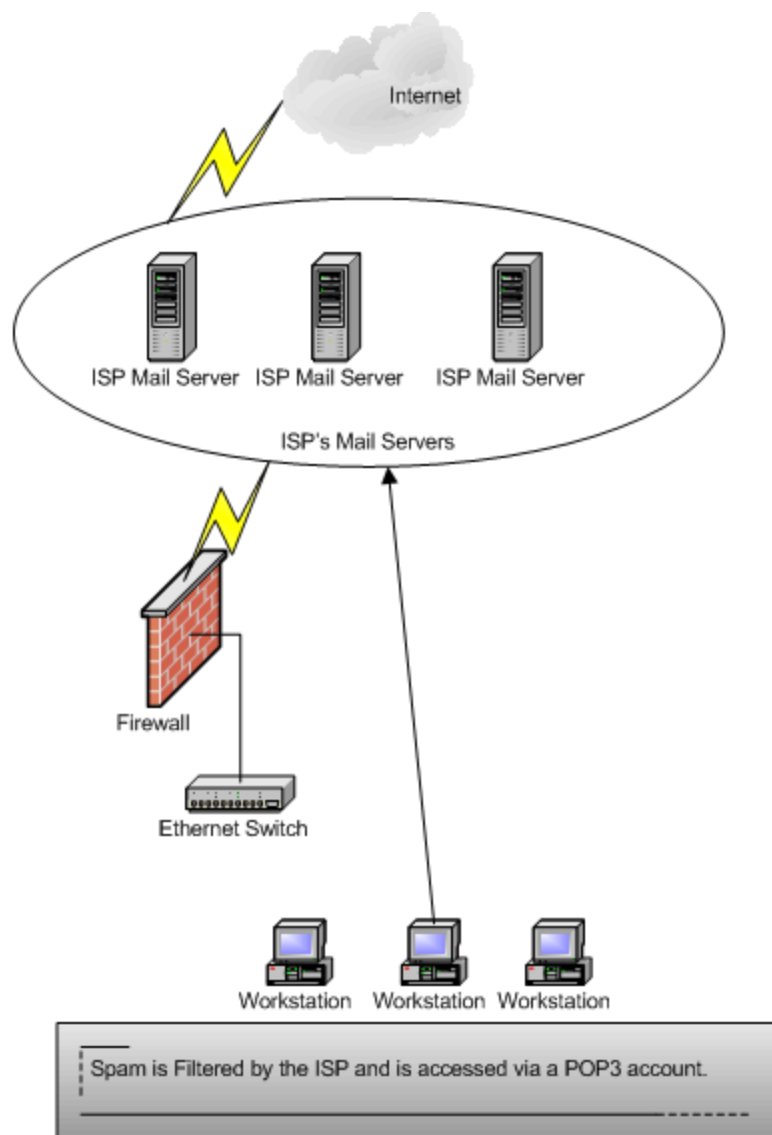


Figure 2.3: Integrated sender validation filtering performed by an ISP.

The following list highlights vendors and ISPs that provide sender validation solutions that are integrated into a mail account provided by the ISP:

- SpamLion (<http://www.spamlion.com>) offers an ASP version of its software designed to work with an ASP or ISP; this company's service enables you to keep your domain name and share protection among a group
- GoodbyeSpam (<https://www.goodbyespam.com/>) provides a package specifically designed to work with an ISP's mail service
- EarthLink spamBlocker (<http://www.earthlink.net/>) offers sender validation features: a sender's email address must be in the recipient's address book in order to for the recipient to receive the message; otherwise, the message is placed in quarantine

Advantages

One of the benefits of this solution is that no mail server is required. In addition, some vendors offer Web-based email services that include sender validation protection. Typically the mail is accessed via a POP3 account.

Disadvantages

With the exception of SpamLion, this sender validation solution is for single users—not for business users—and has most of the same disadvantages as the desktop software and service solutions. In addition, typically, you must use the provider's domain name and cannot receive mail under a personalized domain.

Server-Based Sender Validation Solutions

If you're evaluating a sender validation package for your company, consider a server-based solution. For most companies, the server-based sender validation solution is more cost effective and far more practical to use than installing a sender validation solution on each workstation or using a sender validation service. Server-based sender validation solutions are easier to implement in an enterprise environment because no additional software is necessary on the workstation. They're mail client independent, so if a user accesses the mail with a different client—wireless, OWA, or some other method—the anti-spam software still works because it operates at the server—not desktop—level. Figure 2.4 displays how server-based sender validation is integrated into an existing network.

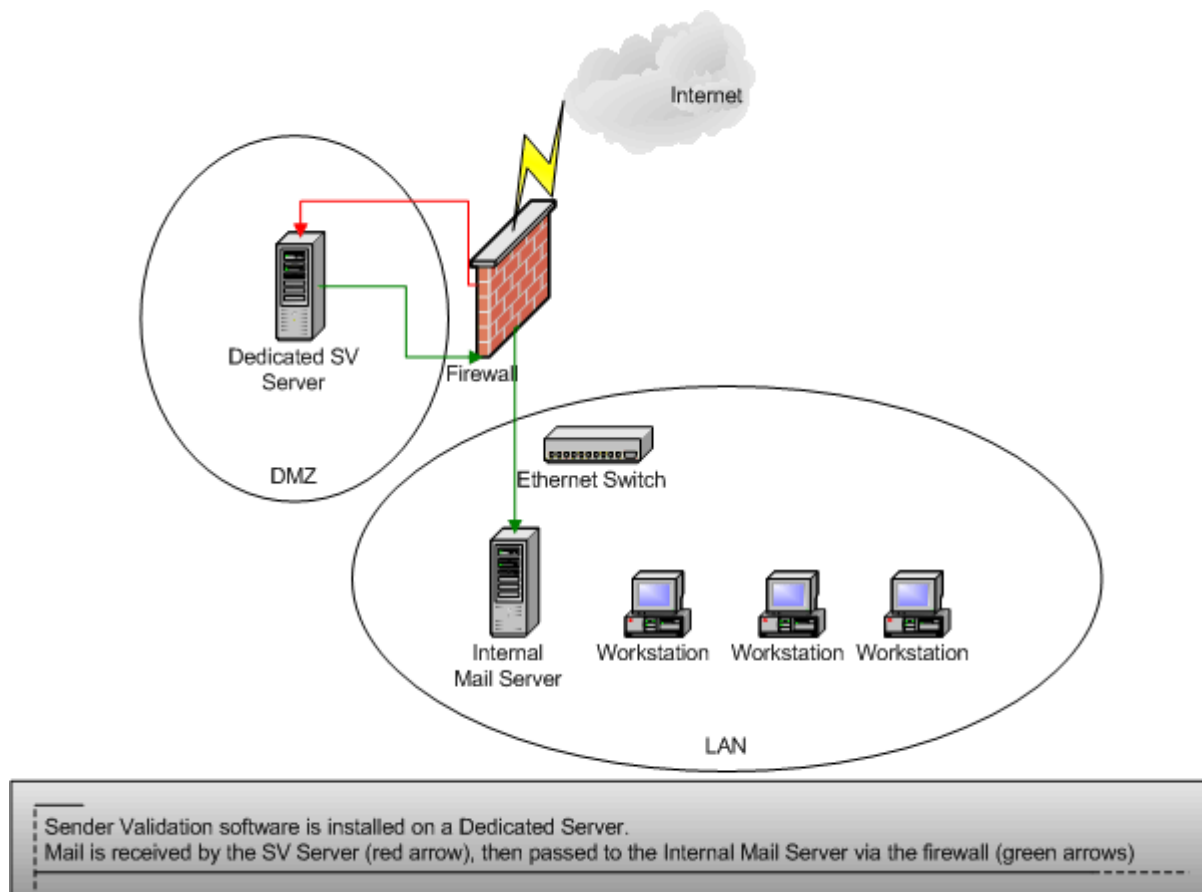


Figure 2.4: Sender validation filtering on a dedicated server.

Sender Validation Business Server Software

The following list highlights vendors that offer server-based sender validation solutions for a business environment. These packages are compatible with any mail server that uses SMTP:

- SpamLion (<http://www.spamlion.com/CorporateEdition.asp>) is typically installed on a dedicated server and can handle a range from ten to thousands of users
- MailFrontier (http://www.mailfrontier.com/products_asg.html) can be installed on a Windows- or Solaris-based computer; in addition to sender validation, MailFrontier uses blacklists and Bayesian filtering to catch spam

Advantages

If your company has a policy of not allowing mail to be handled by a third-party, a software-based server solution is a good fit. They are locally managed, and, typically, the approved sender list can be backed up with existing tape backup software. In addition, server-based sender validation solutions offer the following advantages:

- There is little maintenance after the system is up and running.
- Depending on your mail volume, these packages are not too resource intensive, so a “super” server is not necessary. Some packages will run on your existing mail server and some will run on Linux and other free OSs. The dedicated server will align with other servers (mail, firewall, file and print) just like any other infrastructure server on your network.
- Software that runs on a dedicated server is typically compatible with any mail server that supports SMTP. Some packages can coexist on an existing Web server for reduced startup costs.
- Most sender validation solutions include a Web-based user interface that ensures compatibility with almost any mail client and OS.
- If you’re running Microsoft Exchange Server, some solutions have enhanced features that integrate specifically with Exchange Server.
- For the enterprise environment, some solutions offer load balancing and failover to reduce the probability of downtime. If they do fail, you can simply redirect traffic on the firewall to bypass the sender validation server.

Disadvantages

Most sender validation packages run best on a dedicated server, so the startup cost is higher than that for software that installs on an existing mail or Web server. In addition, some end-user training may be required for users to manage their quarantine and approved senders lists.

For a dedicated server solution, the OS must be installed before the sender validation software can be implemented. The sender validation will require some maintenance (service packs, bug fixes, software updates, and son on), but will not require blacklist, keyword, or message pattern updates.

Some packages require an annual license and do not offer one-time purchase options. And some packages require ISP support using a secondary email relay server to fully implement their failover solution.

Sender Validation Business Services

Some vendors provide sender validation as a service. A sender validation service offers the advantages of a faster implementation and lower initial investment. However, using a sender validation service will probably cost your company more money in the long run. That said, this option is a good alternative if you're short on IT staff and have less than 50 email users to protect. Figure 2.5 shows how a sender validation service is integrated into an existing network.




Figure 2.5: Sender validation filtering service.

The following list highlights companies that offer sender validation services for businesses:


- SpamLion (<http://www.spamlion.com/SpamLionService.asp>) offers a service that has an expandable 10-user license
- Hushmail (<https://www.hushmail.com/services.php?PHPSESSID=475a528134b3e892a14b3c083f161107&subloc=identity>) offers, in addition to sender validation services for business, 2048 mail encryption and digital signing of mail messages

Advantages

This sender validation solution offers shorter implementation time and lower startup costs. In addition, there is little maintenance after the system is up and running. All of the sender validation services are hosted off-site, so you don't have to worry about maintaining another server: no backup issues or maintenance.

Disadvantages

If your company has a policy of not allowing mail to be handled by a third-party, sender validation business service is not an option. In addition, a service may cost more in the long run: typical payback for a dedicated server is less than 2 years. If the service goes down, you will have no mail or spam filtering until the service comes back up. If the service stays down for an extended amount of time, you must change your MX record in order to receive mail. Most services charge by the user, so each additional user will increase your sender validation filtering expense.

 Need more information? Refer to the following sites about sender validation solution and other spam related topics:


 <http://www.spamcon.org/>

 http://www.1stopspam.com/stopspam/Anti_Spam_Organizations/

 <http://www.spamanti.net/en/news/news200310.php>

Sender Validation Planning and Implementation

When considering a sender validation solution for your company, there are a few factors to consider before making your final purchase decision. The following sections explore these additional items that you should consider before purchasing and implementing a sender validation solution in your company. From this point forward, we will focus on sender validation for a business.

 If you're considering a sender validation solution for personal use, take a look at one of the products mentioned earlier.

Evaluating Sender Validation in Your Environment

For any large sender validation implementation, consider a “proof of concept” test to determine how sender validation will work in your environment. Such a test will answer the following questions:

- How will sender validation work with the type of mail my company receives and sends?
- How will my users react to managing their approved senders and quarantine lists?
- How will the sender validation solution handle the anticipated mail load?
- If the sender validation solution is installed on an existing mail server, how will the mail server be able to handle the additional load?
- How will the sender validation solution integrate with my firewall?
- How will the sender validation solution integrate with my mail server?

Consider each of these questions before making a final sender validation purchase decision.


The Right Sender Validation Solution for Your Company

Should you go with a sender validation service or a server-based solution? There are some important factors to consider when making this decision. The following items should help you determine which sender validation solution (service or server-based) is right for your company:

- Number of email users—Sender validation services are a good fit for companies that have 10 to 50 email users. For a larger number of users, it’s generally more cost effective to implement a sender validation server-based solution. With a server-based solution, the larger number of email users, the shorter the payback period, compared with a sender validation service. For example, the typical payback period for 100 users using a sender validation server-based solution is just over 1 year, compared with using a sender validation service. If your company is small but growing rapidly, consider implementing a dedicated server because you will save money in the long run. In general, the greater the number of users, the more cost effective the dedicated server solution will be.
- Company email policy—Some companies have a policy that email must be handled internally. If your company has this requirement, an internal sender validation server-based solution will be your only option.
- Integration with existing mail server and infrastructure—Most sender validation packages integrate with existing mail servers as long as the servers support SMTP. Because the packages work at the server level, you don’t need to worry about email client compatibility. If you’re running products—such as GalSync, Alias Identification, Blackberry Enterprise Server, and Auto NDR—check with the sender validation solution’s vendor to see how well the solution integrates with these or any other email–infrastructure–related packages. This item is an important consideration regardless of whether you select a sender validation service or sender validation server-based solution.

- Product evaluation/demo—For this category, consider the following questions:
 - Does the company offer an evaluation?
 - Is it possible to see a running demonstration of the software package?
 - Is it possible to have the demonstration run with samples of your email?
 - Can you try before you buy?
 - How do the users like the interface?
 - How powerful is the administration interface?
 - Is the package easy to use?
 - Does the sender validation solution offer different ways to handle the validation process?
 - How easy is the validation process for email senders?
 - How easily can a spammer get around the validation process?
 - Will the sender validation solution handle your current and future email load?
 - Is the product scalable?
 - What type of fault tolerance does the solution provide?
 - How easy is it to bypass sender validation filtering in case there is a problem with the sender validation service or sender validation server-based solution?
 - How long has the sender validation solution been around?
 - How do companies that use the sender validation solution like it?
 - How stable is the product (what is the service’s historical uptime percentage)?
- Funds—Do you have the budget for a dedicated server, sender validation software, and installation costs? Make sure that you have enough backup capacity on an existing tape drive or consider purchasing a dedicated tape backup and software for the sender validation server. If you’re short on capital, the initial startup costs are less with a sender validation service compared with a sender validation server-based solution. The more mail users that you have, the faster you can recoup the cost of a dedicated server. You might want to budget for a consultant to ensure that the sender validation solution is properly integrated with your existing DNS, STMP routing, and firewall configuration.
- Internal IT resources—If your company is short on internal IT resources, a sender validation service might be a better fit. The sender validation service can be implemented with fewer internal resources—both in the initial setup and ongoing maintenance. If you are short on IT resources but still want to implement an internal sender validation server-based solution, budget for consulting work to assist in the implementation to save time spent on incorrect configuration and problem solving.

- Virus scanning—An important item to remember is email-based viruses. Sender validation does not check for viruses! Although less than 99 percent of all email viruses come from spammers, a best-of-breed antivirus solution is a good idea. Check whether your current antivirus solution has a dedicated email virus scanner. Some sender validation services have the option of scanning for viruses, which I highly recommend you purchase. Regardless of which sender validation solution you ultimately decide on, you should still maintain virus protection on your servers. If the sender validation service goes down, you might need to bypass the sender validation/virus scanning and accept email directly into your server. For this reason, you should still maintain virus scanning on the server. If you allow users to check their email from POP3 accounts, you still have virus exposure from this source.

 Don't think you are immune to viruses just because you have sender validation filtering!

- Availability—If you're leaning towards the sender validation service, ask the provider about the number of servers, number of users per server, average server utilization, peak server utilization, type of connection, connection redundancy, historical uptime, and a service level agreement (SLA) if the service goes down. In general, you have more control over an internal sender validation server-based solution; you are at the mercy of the sender validation service if the line or server goes down (most sender validation services provide redundant connections, so your email will queue rather than not be delivered). Talk to other companies to see how happy they are with the level of service that the sender validation service vendor provides.
- Licensing—Make sure that the sender validation solution vendor offers product licensing that is a good fit for your company. Find out the following licensing information so that you can compare among vendors:
 - How is the product licensed (per user, per server, per megabyte of storage, per month, per year, one-time license, renewable annually)?
 - Is a maintenance contract required?
 - Are upgrades included in the purchase price?

- Product support—The answers to the following questions are more important if you select a sender validation server-based solution; however they are still important for the sender validation service:
 - Does the company offer 24 × 7 support?
 - Is there an additional cost for extended support?
 - How well does the vendor know their product?
 - Does the vendor have expertise in implementing this package in an enterprise environment?
 - Does the vendor have a staff of implementers that can provide onsite help if necessary?
 - How much does onsite support cost?
 - Does the vendor have experience with your mail server and firewall?
 - What is the average size of their implementations?

☞ Particularly for a server-based sender validation solution implementation, modifications to the firewall, existing mail server, and sometimes MX records must take place in order for the sender validation software to work correctly. Selecting a vendor that is familiar with your environment will prevent problems during the implementation.

- Sender validation service cancellation policy—Obtain the service’s cancellation policy in case you want to move to another solution. It is better to know the answers to the following questions before signing up for a service so that there won’t be any surprises if you decide to cancel the service and move to a dedicated server solution:
 - Before purchasing a solution, get answers to the following questions:
 - Does the vendor need advance notice to cancel the service? If so, how long?
 - If you decide to move to a server-based sender validation solution, can you transfer the approved senders list to an internal server?

Necessary Steps to Implement a Sender Validation Service

Once you've decided to go with a sender validation service, follow these suggested implementation steps (I'll walk you through implementation steps for a server-based sender validation solution in a moment):

- **Select a sender validation service**—Make sure users are comfortable with the user interface offered by the service. Ideally, test the service with your company's email to make sure the service is a good fit for the company. Carefully review the administration capabilities of the sender validation service before making the final selection. During the evaluation phase, some sender validation companies have the ability to turn on sender validation filtering for certain email accounts with any remaining accounts set to a "bypass" mode.
- **Consider a phased rollout of sender validation**—Consider at least a two-phased approach when rolling out the sender validation service. Phase one will be a select number of power users. After phase one, you can fine tune sender validation, address any issues that arise, and provide additional training before rolling out sender validation for the entire company. The last thing you want is a bunch of upset email users as a result of your failure to anticipate all of the potential problems prior to the rollout.
- **Train end users**—Make sure that your users understand the validation process and can add/delete users from the approved senders list.
- **Change your MX record**—You must redirect your MX record to the servers of the sender validation services. The service's servers will redirect your email to your internal server. You might want to add a backup MX record to your ISP's mail server in case the sender validation service's mail goes down (so they can hold the mail for you). You can set a backup MX record directly to your internal server; however if the service does not respond in a timely basis, mail will be directly delivered to your mail server bypassing the sender validation service. In addition, if a spammer discovers that the backup MX record points directly to your mail server, the spammer can use the backup MX record to completely bypass your sender validation service. If you decide to use a backup MX record that points directly to your internal mail server, create a rule on your firewall to only accept incoming mail from the sender validation service's servers and your ISP servers. That way, the firewall will prevent anyone from sending mail directly. If the sender validation service's servers go down, simply disable this rule to bypass the sender validation servers and allow incoming mail to flow directly to your internal mail server. It's much easier to disable a rule on the firewall than calling the ISP to change the MX record for your domain, and waiting for the change to replicate across the entire Internet.
- **Mail server modifications**—Some sender validation services have the capability of auto-adding outgoing email addresses to the approved senders list. To use this functionality, you must redirect your outgoing mail to the sender validation service's mail server. Alternatively, some sender validation services require installing a "learning" module on your existing mail server. If you don't plan to use the learning feature, no modification of the mail server is usually necessary. Without the learning module, everyone must complete the validation process, which is a big inconvenience. This inconvenience was a major reason why earlier attempts at sender validation were rejected. I strongly suggest turning on the auto-learn feature.

- Follow up—Expect some issues to arise when turning on the sender validation service. Some users will require help adding/deleting users on their approved senders list, and senders might have difficulty completing the validation process.
- Quarantine period—You might want to increase or decrease the quarantine period for your messages based on your company’s requirements.

Necessary Steps to Implement a Sender Validation Dedicated Server

If you have more than 50 users, consider a sender validation server-based solution. It’s usually more cost effective to run an internal sender validation server compared with the cost of a service. The dedicated server will handle the sender validation process. Consider the following suggestions when implementing a dedicated server sender validation solution:

- Select a sender validation package—As with the first step of implementing a sender validation service, make sure users are comfortable with the user interface. In addition, obtain answers to the following questions:
 - Does the software package have a planning guide to assist in the implementation?
 - What type of support does the vendor offer during the transition period?
 - Does the vendor offer support after the initial transition is complete?
 - Does the software vendor have expertise only in the software package they support? (Ideally the vendor should be familiar with your mail server, firewall, ISP, and MX record changes.) The more familiar a vendor is with your network, the smoother the transition should be.
- Consider a phased rollout of sender validation—This consideration follows the same guidelines as with a sender validation service implementation. The phased approach provides a safety net with limited exposure in case unforeseen issues arise.
- Train end users—Depending on the level of experience of your user base, it might be necessary to train your users on how to use the sender validation software. Ideally, the training should take place just before their mailboxes are cut over to the service so that they will retain as much of the training as possible.
- Decide on sender validation server placement—You can place the sender validation server behind the firewall, in the DMZ, in front of the firewall/off-site. If you do not have a firewall, I strongly suggest purchasing one before implementing the sender validation server. It will be more difficult and ultimately take more time if you have to install a firewall after you install the sender validation server. If your firewall has the capability, I suggest installing the sender validation server in the DMZ. Doing so allows the sender validation server to become the “sacrificial lamb” in case the server is attacked by hackers. This setup increases security because there is no direct email communication with your internal email server and the Internet for incoming mail. The downside of placing the sender validation server in the DMZ is that the firewall configuration is the most complex with this topology.

- DNS changes—Usually no DNS or MX record changes are necessary to implement the server if you redirect mail by altering your firewall configuration. An MX record change pointing to the sender validation server might be necessary if you want OWA users to retain the same IP address.
- Firewall changes—To set up the sender validation server, you must redirect incoming mail from your internal mail server to the new sender validation server. If you are running Web-based validation, you must redirect port 80 traffic to the sender validation server. Doing so can be problematic if you're self-hosting an existing Web site—although some sender validation servers can coexist with an existing Web site. If you want the sender validation solution on a separate server, you can either set up an additional external IP address to handle the Web-based validation or use port redirection on the firewall to handle inbound Web-based validation. If you are running any VPNs, make sure to test the sender validation server with the remote locations to ensure that they can use the sender validation server.

Summary

Sender validation packages come in various forms: desktop software, desktop services, integrated with an ISP-provided email account, dedicated business server, and business service account. As we explored, each of these sender validation solutions has advantages and disadvantages, and there are many vendors to choose from for each of type of solution.

We also looked at the critical issues to help you select a package that meets your company's needs. A little homework up front will help you with a successful implementation. After addressing the questions in the evaluation section, you should have a good idea which sender validation route to take. Once you decide on the right type of solution, you can follow the suggestions in the related implementation steps section.

In the next chapter, we'll take a look at cost justification of the sender validation solution. Don't worry—sender validation is probably one of the easiest IT projects to cost justify. We'll take an in-depth look at the implementation steps necessary to ensure a successful sender validation rollout for your company. We'll also discuss ongoing support issues, upgrades, maintenance, and how to support additional users. Finally, we'll explore what the future holds for sender validation.

Chapter 3: Implementing a Sender Validation Solution in Your Company

Sender validation is the only spam solution I know of that has the potential to eliminate 100 percent of a company's spam. There are many ways to implement a sender validation solution in your company. Sender validation can be implemented as a service or a dedicated server. The dedicated server can be installed in your company's DMZ, local area network (LAN), or in a co-location facility. Each configuration has its advantages and disadvantages. Your email environment will dictate the best sender validation configuration for your company.

In this chapter, we'll take a look at the cost justification of a sender validation solution, the estimated startup and maintenance costs, and the implementation steps necessary to set up a sender validation solution as a service or dedicated server. In addition, I'll provide best practices to follow after the sender validation solution is up and running.

Cost Justification Compared with Other Methods

Ideally, a sender validation solution should be installed on a dedicated server for the best performance and stability. If you have more than 100 users in your company, you probably want to implement any server-based anti-spam solution on a dedicated server.

Remember that any anti-spam solution will increase the load on an existing server. Often mail servers are already heavily loaded, and installing an anti-spam solution on the server will just make the situation worse. By installing a sender validation solution on a dedicated server, you ensure compatibility with other SMTP mail servers and reduce the risk of having the sender validation software conflict with existing mail services, such as antivirus software.

Expect to pay \$1000 to \$3000 for a name brand (such as Hewlett-Packard and Dell) server capable of handling a server-based sender validation solution for more than 100 users. Make sure to budget the time and cost to install the operating system (OS) on the anti-spam server. In addition, make sure you have adequate backup capacity to backup this additional server, an available port on your Ethernet switch, and the physical room to hold the server.

Software Cost

Of course, you must factor in the cost of the sender validation software. Most vendors sell their software by the number of users. For larger installations, many vendors offer a quantity discount. You can contact the software vendor directly to get the best deal. Expect to pay \$5 to \$30 per user for a sender validation solution. This option is less expensive than a spam filtering service but more than a non-sender validation spam solution.

The software cost is typically a one-time purchase cost. However, some sender validation vendors require an annual license renewal. This fee usually includes free technical support and upgrades during the maintenance period. Expect to pay 15 to 25 percent of the purchase price for maintenance and upgrades.

Implementation Benefits and Cost Justification

Fortunately, a sender validation solution is one of the easiest IT projects to cost justify, especially if you have more than 100 email users in your company. The return on investment (ROI) is usually less than a year, and in some cases, as little as several weeks.


An ROI Example

The biggest selling point of a sender validation solution is the time savings secured by the end users. Consider the following simple example: A company with 200 email users implements a sender validation solution that costs \$8000. On average, each employee of the company is paid \$25 per hour. After implementing the sender validation solution, each employee saves 10 minutes per day because the employees do not have to wade through junk mail, are not constantly interrupted by “ding—new message” every 15 minutes, and are less likely to accidentally delete valid email messages. Based on these assumptions, the company’s savings per day is:

$200 \text{ users} \times \$25 \text{ per hour} \times (10 \text{ minutes}/60 \text{ minutes}) = \$833.33 \text{ savings per day}$

If you take the total cost of \$8000 for the sender validation implementation and divide it by \$833.33, it takes roughly 10 days to pay for the sender validation solution. Let’s assume that this company works 5 days per week. In this example, the company will have an ROI of 2 weeks! It is difficult to identify any IT project that has a shorter ROI period.

Based on this simple example, the short ROI period should cost justify any sender validation project. The elimination of undesirable spam messages (pornographic ads) should also reduce coworker tension, and reduce the likelihood of a lawsuit issued against the company.

 The ROI for a company implementing a sender validation spam solution can be as little as 2 weeks.

In addition to the increased user productivity, there are other benefits of a sender validation solution. Reduced storage on your internal mail server is one such benefit.

Reduced Storage on Your Internal Mail Server

A sender validation solution will reduce the storage requirements on your mail server because the junk mail will never reach the server in the first place. Reducing the amount of mail on the server has the following benefits:

- Better mail server performance—Reduced mail storage increases performance of the mail server—especially when performing searches.
- Reduced mail store size—A significant percentage of a company’s mail store contains spam. Although spam might reside in the deleted items folder, it still takes up space on the server before it is permanently deleted. Failed non deliverables (NDRs) and bad mail items can consume significant space on a mail server. Some companies estimate that spam takes up half of their current mail storage space.
- Reduced backup/restore times—The backup and restore times on any mail server will be reduced because of the decreased mail store size.
- Reduced backup storage requirements—If your tape backup is close to capacity, implementing an anti-spam solution might eliminate the need to purchase a higher-capacity backup system. Even if a company must back up the sender validation server, the company will still have a net reduction in backup capacity requirements.
- Reduced mail store maintenance time—With a reduced mail store size, defragmentation of the mail store and mail store repair utilities will run faster. In the event of a mail store corruption, downtime will be reduced because repair utilities do not have to deal with a large volume of junk mail on the server.
- Reduced stress on your WAN links—Sender validation reduces the amount of spam traffic on your WAN links to internal remote mail servers.
- Sender validation has the potential to eliminate 100 percent of spam—Sender validation can be implemented as your company’s first anti-spam solution or as an upgrade to replace an outdated anti-spam solution.

Ongoing Administrative Costs

Although there are many benefits that enable an organization to easily cost-justify a sender validation solution implementation, to truly compare sender validation solutions, you must consider the ongoing administrative costs. The following list highlights these considerations:

- Sender validation service solutions require annual re-licensing—Some sender validation solutions require annual re-licensing of as much as 100 percent of the product cost per year. Although this cost can be very expensive in the long run, it might be the best solution for your company—such is especially true of smaller companies that use a sender validation service and do not have an internal server.
- Sender validation dedicated server annual maintenance costs—For most sender validation server solutions, expect to pay 15 to 30 percent in maintenance fees for ongoing technical support and software upgrades.

However, after the sender validation product is in place, there is very little ongoing maintenance. The most important maintenance item on the sender validation server is the backup of the sender validation database. This database contains all of the company's "approved senders." As long as this database is backed up, you should be able to quickly recover from any hardware problem.

Minimal User Support Costs

Most of the sender validation solutions include an integrated junk mail box or Web interface to manage their junk mail. The initial training on a sender validation solution should be minimal. An hour training session to explain how sender validation works and how a user manages junk email should be more than adequate. Some users might not require any training.

After the initial end-user training costs, ongoing training costs are minimal:

- New hires might require a brief training session.
- Existing users might require some training on a new release of the software so that they can take advantage of any new or enhanced features of the software.
- Help desk support and desktop costs are minimal because everything is managed at the server level.

Implementation Steps for a Sender Validation Service

Figure 3.1 illustrates a typical sender validation service implementation.

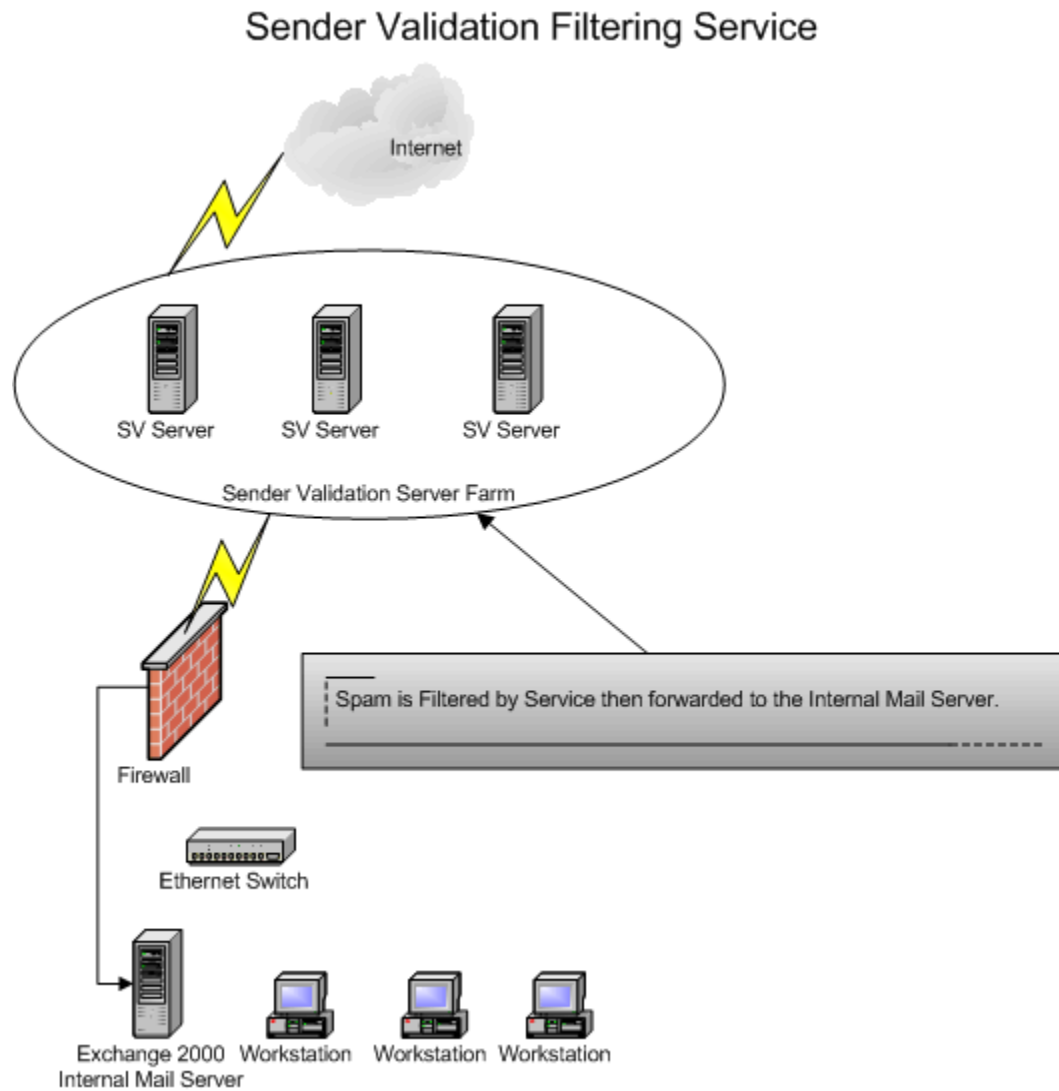


Figure 3.1: Sender validation filtering service implementation.

The following sections walk you through the sender validation service implementation steps.

Select a Sender Validation Service

Decide on the service to use. Make sure users are comfortable with the user interface. Ideally, test the service with your company's email to make sure it is a good fit for the company. Carefully review the administrative capabilities of the sender validation service before selecting a solution. During the evaluation phase, some sender validation companies can only filter certain email accounts with any remaining accounts set to a bypass mode. This setup allows the company to try before they buy so that they can identify any issues that might arise during the trial phase.

Select an Implementation Strategy

Decide whether you want to move the entire company over to the sender validation solution all at once or roll out the sender validation service in phases. The larger the company, the more likely you'll use a phased approach.

Train End Users

Depending on the skill level of your users, it might be necessary to train your users on how to use the sender validation service. Ideally, the training should take place just before their mailboxes are cut over to the service so that they will retain as much of the training as possible before using the actual service.

Select an Initial Implementation Date

Select an initial cut-over date—ideally, on a weekend. This selection should allow enough time for the MX record change to propagate over the Internet. Before you make the change, consider reducing the time to live (TTL) for your existing MX record so that the change will propagate faster throughout the Internet.

Upload a Pre-Approved Senders List

If the service allows an upload of pre-approved senders, I suggest using this feature to reduce the number of messages in quarantine. Alternatively, if the service has a “learning mode,” you can redirect all outgoing messages to the sender validation service and have any sender addresses automatically added to the approved senders list before turning on the sender validation service. Some services require a module on your existing server to forward the email addresses to their server for the approved list. This learning strategy will work better if your company sends mail to a set number of users on a frequent basis. I strongly suggest employing either one of these methods to upload a pre-approved sender list; otherwise, each sender must be manually validated. Manual validation is one of the reasons why early sender validation solutions developed a bad reputation.

Change Your MX Record

You must redirect your MX record to the sender validation service's servers. The sender validation service's servers will then redirect your email to your internal server. I suggest setting up a backup MX record to your ISP's mail server so that the ISP can hold your mail if the sender validation service goes down. You can set a backup MX record directly to your internal server; however, if the service does not respond in a timely manner, mail will be directly delivered to your mail server, bypassing the sender validation service. In addition, if a spammer figures out that the backup MX record points directly to your mail server, the spammer can use the backup MX record to completely bypass your sender validation service.

Follow Up

Expect some issues to arise when turning on the sender validation service. Some users will require help adding users to their approved senders lists, and senders might have difficulty completing the validation process. Prepare your users for a dramatic reduction in mail messages. Some users think their mail might be broken because they don't receive any messages in their Inboxes.

Set Up a Quarantine Period

You might want to increase or decrease the quarantine period for your messages based on your company's requirements.

Sender Validation Internal Server Implementation Steps

If you have more than 100 users, consider a sender validation server-based solution. Running an internal sender validation server inside the company is usually more cost effective than paying a recurring monthly or annual fee for a sender validation service. The following list highlights steps for a sender validation internal server implementation:

Select a Sender Validation Package

To do so, you will need to answer the following questions:

- Does the sender validation solution offer a product evaluation?
- Is it possible to evaluate the product for some users in your company or is it an all-or-nothing implementation?
- Does the product evaluation have any time or user limits?
- Does the sender validation solution offer the product as a service with the option to migrate to an in-house server in the future?

Answers to these questions should help you determine the correct sender validation solution for your company. In addition, make sure users are comfortable with the user interface. Carefully review the administration capabilities of the sender validation package before selecting a solution:

- Does the software package have a planning guide to assist in the implementation?
- What type of support does the company offer during the transition period?
- Does the software vendor have expertise only in the software package they support?

Ideally the vendor is familiar with your mail server, firewall, ISP, and MX record changes. The more familiar the vendor is with your network, the smoother the transition.

Select an Implementation Strategy

Decide whether you want to move the entire company over all at once or rollout the sender validation server in phases. The larger the company, the more likely you will use a phased approach.

Plan End-User Training

Depending on the skill level of your user base, it might be necessary to train your users to use the sender validation software. Ideally, the training should take place just before their mailboxes are cut over to the service so that they will retain as much of the training as possible.

Decide on Sender Validation Server Placement

You can place the sender validation server behind the firewall, in the DMZ, in front of the firewall, or in a co-location facility. If your firewall has the capability, install the sender validation server in the DMZ. Doing so allows the sender validation server to become the “sacrificial lamb” in case the server is attacked by hackers. This setup increases security because there is no direct email communication with your internal email server and the Internet for incoming mail.

Purchase the Server Hardware

Based on the sender validation software recommendations, order the server hardware that will safely support your company’s email load and number of users. Don’t forget to include a provision to backup the server, either with a dedicated tape drive or existing backup resource. Make sure you have an open Ethernet port on your switch and the physical space to accommodate the server.

Install the OS on the Server

Install the OS according to the sender validation solution vendor’s recommendations. Some packages are sensitive to OS version and service pack levels. Install any critical security patches on the server to protect it from hackers. Make sure that the OS configuration matches the sender validation requirements.

Install the Sender Validation Software on the Server

Install the sender validation server software on the server. Be sure to follow any special requirements during the installation process. After you install the software, test the sender validation server to make sure it’s not an open relay. You can use the testing tool at <http://www.ordb.org/submit/> for verification. Performing the open relay test will ensure that the new sender validation solution is not entered into an open relay database.

Select an Initial Implementation Date(s)

Select an initial cutover date that is on a weekend. Doing so will give you more time to reconfigure your firewall and test the sender validation server implementation. It also gives you more time to restore the firewall and servers to their original configurations in case something goes wrong. Make sure that the cutover dates coincide with end-user training.

Upload a Pre-Approved Senders List

Use the pre-approved senders list feature to reduce the amount of messages in quarantine. Alternatively, if the service has a “learning mode,” you can redirect all outgoing messages to the sender validation service and have any sender addresses automatically added to the approved senders list before activating the sender validation server. Some services require a module on your existing server to forward the email addresses to the server so that the addresses can be added to the approved list. This learning strategy will work better if your company sends mail to a set number of users on a frequent basis. Without a pre-approved senders list, each sender must be validated manually. Manual validation is one of the reasons why early sender validation solutions developed a bad reputation. Create the list well in advance of uploading it to the server so that you have enough time to ensure the list is complete and accurate.

Reconfigure the Firewall

Before making any changes to the firewall, make sure you have a good backup of the firewall configuration. Doing so will allow you to quickly restore your current mail configuration if issues arise during the implementation. Figure 3.2 illustrates a sender validation server implementation in the DMZ.

Spam Filtering on a Dedicated Server in the DMZ

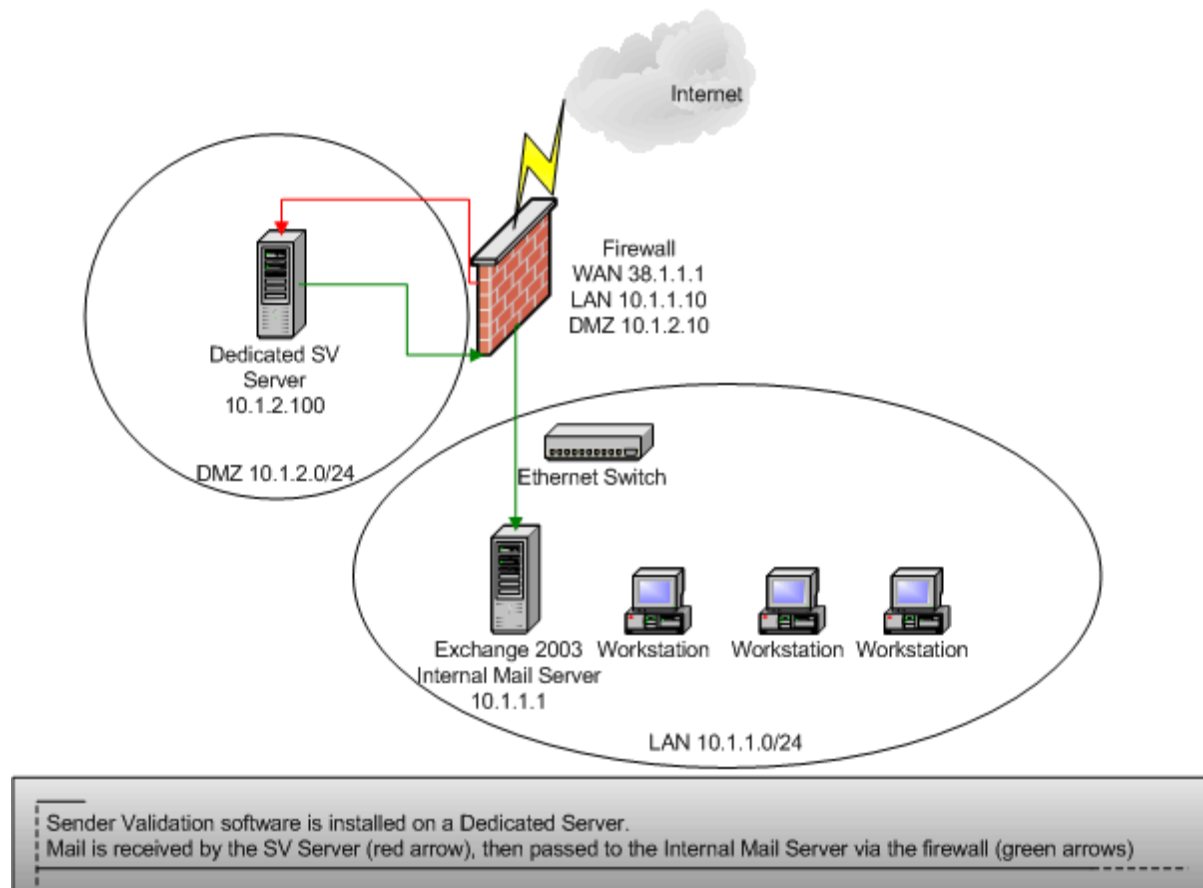


Figure 3.2: Sender validation dedicated server in the DMZ.

Create the following NAT rule on the firewall for the sender validation server (assume that the MX record points to 38.1.1.1 and the internal address of the sender validation server in the DMZ is 10.1.2.100):

Public Address: MX record for the mail server (38.1.1.1) on the WAN

DMZ Address: Private address of the sender validation server (10.1.2.100) in the DMZ

Description: One-to-one NAT rule to translate the public IP address of the mail server (MX record) to private IP address of the sender validation server.

Create the rules on the firewall for the sender validation server that Table 3.1 shows (assume the IP address of the internal mail server is 10.1.1.1).

Source	Destination	Port(s)	Port and Description
Any from the WAN	Sender validation server in the DMZ (10.1.2.100)	25 (SMTP) and 80 (Web)	Allow incoming mail and Web traffic (Web-based validation) to the sender validation server from the public Internet
Sender validation server in the DMZ (10.1.2.100)	Internal mail server (10.1.1.1) on the LAN	25 (SMTP)	Allow the sender validation server to send mail to the Internal mail server
Any from the LAN	Sender validation server in the DMZ (10.1.2.100)	80 (Web)	Allow users to manage their quarantined mail using a Web-based interface
Internal mail server on the LAN (10.1.1.1)	Sender validation server in the DMZ (10.1.2.100)	25 (SMTP)	Allow internal mail server to send mail to the sender validation server in the DMZ
Sender validation server in the DMZ (10.1.2.100)	Any on the WAN	25 (SMTP) 53 (DNS) 80 (Web)	Allow sender validation server to send out mail to the Internet

Table 3.1: Rules on the firewall for the sender validation server.

These rules might vary depending on the type of firewall and the specific requirements of your sender validation server package. These rules assume that the outgoing mail is forwarded by the internal mail server to the sender validation server so that the sender validation server can inspect the outgoing addresses and add them automatically to the approved senders list. Figure 3.3 shows a sender validation server on the LAN.

Spam Filtering on a Dedicated Server on the LAN




Figure 3.3: Sender validation dedicated server on the LAN.

Create the following NAT rules on the firewall for the sender validation server (assume that the MX record points to 38.1.1.1 and the internal address of the sender validation server in the DMZ is 10.1.1.2):

Public Address: 38.1.1.1 on the WAN interface

LAN Address: 10.1.1.2 on the LAN interface

Description: One-to-one NAT rule to translate the public IP address of the mail server (MX record) to the private IP address of the sender validation server.

Create the rules on the firewall for the sender validation server that Table 3.2 shows.

Source	Destination	Port(s)	Port and Description
Any from the WAN	Sender validation server on the LAN (10.1.1.2)	25 (SMTP) and 80 (Web)	Allow incoming mail and Web traffic (Web-based validation) to sender validation server from the public Internet
Sender validation server on the LAN (10.1.1.2)	Any on the WAN	25 (SMTP) 53 (DNS) 80 (Web)	Allow sender validation server to send out mail to the Internet

Table 3.2: Rules on the firewall for the sender validation server.

These rules might vary depending on the type of firewall and the specific requirements of your sender validation server package.

If the sender validation server is in a co-location facility, establish a VPN connection to the server. That way, all traffic will be encrypted between your internal mail server, LAN, and the sender validation server. Figure 3.4 shows a sender validation server in a co-location facility.

Spam Filtering on a Dedicated Server on the WAN

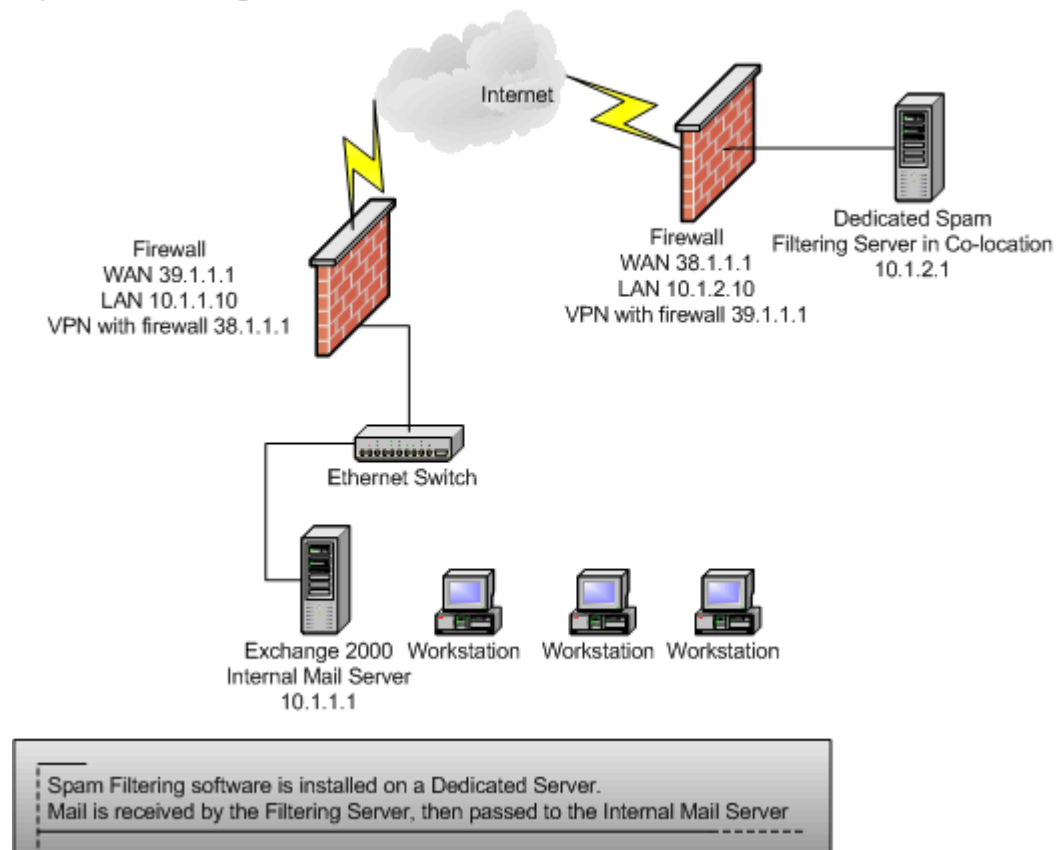



Figure 3.4: Sender validation dedicated server in a co-location facility.

Create a VPN between the co-location firewall and the company firewall. It might be necessary to change your MX record to point to the public IP address of the sender validation server. Make sure to add a Reverse (PTR) record for your sender validation server to avoid mail delivery problems with other mail servers. Some mail servers perform a reverse lookup from the receiving mail server as a way to combat spam. If an MX record change is necessary, make sure to allow enough time for the MX change to propagate throughout the Internet.

As with the previous configurations, these rules might vary depending on the type of firewall and the specific requirements of your sender validation server package.

 Any of these firewall changes can cause serious disruption in service and/or security holes if they are not created properly. If you do not feel comfortable with these changes, budget some time to have the firewall vendor, sender validation server vendor, or a qualified consultant assist you in the sender validation server implementation. Make sure to get a good backup of any firewall configuration before you begin so that you can quickly restore the original configuration if necessary.

Make Mail Server Modifications

Depending on the sender validation software requirements, you might have to re-route your outgoing mail through the sender validation server—especially if you want to take advantage of any auto-learning features. Some sender validation servers require that a module is installed on the mail server to enable this functionality. Check the sender validation server documentation for any other mail server modifications.


Activate the Sender Validation Server

Activate the sender validation server for the desired number of users. Unless the company is very small, I suggest a two-phase approach. Turn on the sender validation software for a select number of users, then fine-tune the system. Create a document to address any of the questions. Before you activate the sender validation software for the remaining users, distribute this document to your end users to reduce the number of Help desk calls. When you activate the remaining users, the system should be fine-tuned to your company's email environment.

Test the Sender Validation Server

Have someone send you mail from a test account to test the following scenarios. You can use the same test mail account by simply adding/removing the email address from your approved senders list:

- **Preloaded list test.** If you decide to upload a pre-approved list of senders, make sure the test account is in the list and mail is accepted from anyone on the pre-approved senders list. Remove the test account from the approved senders list after this test is completed.
- **Manual validation test.** Have an end user manually add the test account to the approved senders list. Doing so will set up the next test.
- **Approved senders list test.** Make sure that a test message from the user added in the previous step is delivered properly.
- **Manual deletion from the approved senders list test.** Manually delete the test user from the approved senders list.
- **Auto-add to the approved senders list test.** If you are using the auto-add to the approved senders list feature, send a message to the test account and verify that the email address is auto-added to the approved senders list. After you verify this addition, delete the test account to prepare for the next test.
- **Sender validation test.** Send a message from the test account and make sure that you receive the validation request. Have the test user complete the validation process and verify that the mail is delivered correctly.
- **Quarantine test.** This test cannot be done initially—ensure that the messages in quarantine are deleted when they reach the proper age. For testing purposes, you can reduce the quarantine period to a shorter time to ensure that the messages are properly deleted from quarantine.

 Proper testing of the sender validation server will ensure a successful implementation.

Establish a Quarantine Period

You might want to increase or decrease the quarantine period for your messages based on your company's requirements.


Backup

Make sure to check the backup status of the server to ensure that the server and sender validation databases are properly backed up.


Post Installation Tasks and Best Practices

After installation is complete, there are some additional tasks to be done. If you use best practices for these tasks, you will enjoy a successful sender validation implementation:


- Document firewall changes—Make sure to document any firewall changes. Record the IP addresses of the sender validation server and mail servers. You might want to keep a copy of the original approved senders list just in case you need to restore it. Create contact information for any key personnel involved in the installation.

 Save the pre-sender validation firewall configuration in case you must bypass the sender validation server in the future.


- Back up the sender validation server—Make sure to incorporate a backup of the sender validation server into your existing backup strategy.

 Ideally the entire sender validation server should be backed up; however if you're short on backup space, the approved senders database is the most critical information on the sender validation server.

- Develop a failover plan—In the event of a hardware failure of the sender validation server, create a failover plan to temporarily bypass the sender validation server. This plan might require firewall reconfiguration, mail server reconfiguration, and MX record changes depending on your environment.
- Perform end user follow up—Survey a sample of users, gain feedback, and address any outstanding issues. Train end users how to send a test mail from a test account or other “approved sender” so that they can verify whether the mail is still working.

 Some users might experience such a dramatic decrease in spam they may think their incoming mail is not working.

- Check quarantine daily—For the first 30 days of the sender validation activation, check the quarantine location daily. Even with a preloaded approved senders list, users might still have a few senders in their quarantine list that they might want to manually add to the approved senders list. After 30 days, users can probably reduce the quarantine check to once every few days. Even with daily quarantine checking, users should still save a significant amount of time with the sender validation solution. When reviewing quarantine, train users to read the sender’s name rather than looking at subject line.

 Typically a user’s “mental filter” looking at quarantine is 10 times faster than looking through an Inbox filled with spam.

- Monitor logs, performance, and data backups—Review the logs on a regular basis and address any issues that arise. It’s a good idea to become familiar with the logs so that you get a feeling for what is normal and what is an exception.

Performance of the sender validation server is critical to ensure that mail is delivered in a timely manner and that the validation process is working. Make sure that the sender validation server has adequate disk space, processing power, and memory. Identify any bottlenecks by using tools such as the Windows Server Performance Monitor, and address bottlenecks as necessary.

In addition, review the backup logs to ensure that the sender validation server is properly backed up.

- Install upgrades, patches, and hotfixes—Install maintenance releases of the software when necessary.

Dealing with eCommerce and Other Legitimate First-Contact Situations

eCommerce and other first-contact situations can lead to false-positives because the senders typically do not respond to these emails. Train users to pre-approve the sender as well as to manage their quarantine to manually add these types of first contact senders to the “approved list.” Some sender validation vendors are working on enhanced features to automate this process in the future.

Troubleshooting

In rare instances, a user might get a spam message. Some sender validation solutions can track who or how a sender was validated. This feature is a good place to start when tracking down how a spammer was added to the approved senders list. Make sure that your internal mail server will only accept incoming mail from the sender validation server; otherwise, a spammer can bypass the sender validation server to deliver spam.

Summary

Some sender validation solutions might support installation directly on the mail server in the future. Of course, you must be running a mail server that is supported by the sender validation solution. This option is attractive for companies with fewer users, because such companies do not have to purchase a dedicated server in order to implement a sender validation solution.

Another development on the horizon is the coexistence of a sender validation solution with firewall or anti-virus software. This develop will eliminate the need for a dedicated server for the sender validation solution. However, installing any service on top of a firewall will make the firewall less secure. It is less fault tolerant because you will lose multiple services if the firewall fails. If your firewall is already heavily loaded, a dedicated server is still the best solution.

Finally, a sender validation lite solution might become available on the market. Such a solution simply verifies that an email address belongs to a valid user and domain. It does not require a response from the sender. Although this option requires less processing power, it is easier to get messages past the spam filter by simply spoofing the sender's address.