

Introduction to Innovative Mail Filtering with AXIGEN

How to secure your email from spam and viruses using AXIGEN Mail Server's Milter protocol-based filtering with SpamAssassin, ClamAV, MIMEDefang

❑ The problem with electronic mail today

Email is not what it was supposed to be. Its inventors have not foreseen the dangers associated with this new form of communication. Means of protection were added later on, but do little to protect people's mailboxes. As a result, most of the mail traffic these days is unwanted: spam, malware, phishing.

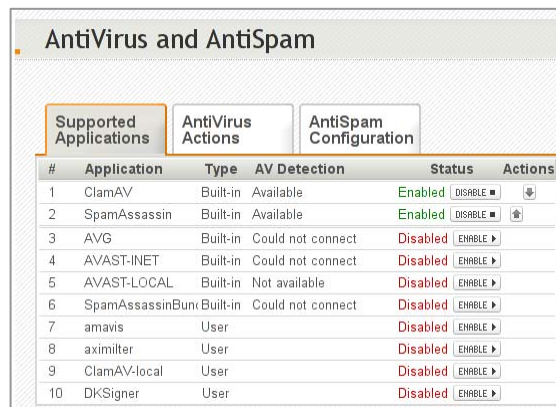
Important resources are allotted by companies to deal with this situation and the results are usually inadequate. Anti-spam filters are not only incapable to detect all the spam, but they also happen to flag legitimate mails as spam. Antivirus products are basically a game of playing catch-up, with vendors adding more and more signatures to their products in response to new threats in the wild. Phishing has also become a huge problem lately, as gangs of underground cyber-criminals become more and more successful in luring new victims in their nets.

How can we use AXIGEN and existing free anti-spam, antivirus and anti-phishing software and technologies so that the result is more than the sum of the parts?

❑ Traditional filtering in AXIGEN

From its very beginnings, AXIGEN has integrated open-source filters such as SpamAssassin and ClamAV and has used the fastest way to process mails through them by interfacing with their daemons directly (*spamd* and *clamd* in our case). The downside was that mail messages were scanned in the queue, after having been accepted by the mail server.

In such a scenario, a number of actions are possible based on the results from content filters such as SpamAssassin and ClamAV:



AntiVirus and AntiSpam					
Supported Applications		AntiVirus Actions		AntiSpam Configuration	
#	Application	Type	AV Detection	Status	Actions
1	ClamAV	Built-in	Available	Enabled	DISABLE ▾
2	SpamAssassin	Built-in	Available	Enabled	DISABLE ▾
3	AVG	Built-in	Could not connect	Disabled	ENABLE ▶
4	AVAST-INET	Built-in	Could not connect	Disabled	ENABLE ▶
5	AVAST-LOCAL	Built-in	Not available	Disabled	ENABLE ▶
6	SpamAssassinBunc	Built-in	Could not connect	Disabled	ENABLE ▶
7	amavis	User		Disabled	ENABLE ▶
8	aximilter	User		Disabled	ENABLE ▶
9	ClamAV-local	User		Disabled	ENABLE ▶
10	DKSigner	User		Disabled	ENABLE ▶

- **Quarantining:** used especially for malware, this requires human intervention for double-checking the results and for rescuing false positives from the quarantine when someone notices a malfunction of the filters.
- **Tagging:** specific headers are added to messages in order to mark them as unwanted (e.g. SpamAssassin adds X-Spam headers). AXIGEN is capable of filtering such unwanted traffic in a specific sub-folder (typically named "Spam") which may be inspected for false-positives by the receiver.
- **Bouncing:** also known as "backscattering", a form of collateral spam, because most unwanted mails have forged message envelope senders. Bouncing mails to those envelope senders is the wrong choice, as many of them are innocent persons whose mail addresses have been harvested by spammers.

- **Discarding:** a more doubtful choice (no filter is perfect) and a big no-no in many ISP and corporate environments.

As you can see, every action presents some disadvantages. The only ones that we recommend for a typical setup are the first two: **quarantining** for malware and **tagging** for spam. However, quarantining usually requires administrative intervention in case of false positives. Tagging is also problematic, especially for high-profile mail users who receive tons of spam and don't have the resources to double-check the filtered messages that land in the "Spam" folder. If users do not double-check and they suspect that a false positive has occurred, they have to look it up in a large pool of unwanted messages. Another downside is the fact that "Spam" folders tends to grow indefinitely and waste space in the server storage and backup mediums.

A better solution would be to scan incoming messages before accepting them, during the SMTP conversation. This is possible using the standardized Milter protocol, a communications interface between a mail server and a content filtering service. However, the first implementation of the Milter protocol in AXIGEN, in the form of the **Aximilter** module, was very much like the other AXIGEN filters, requiring the processed messages to be placed in the queue before scanning them.

□ A new approach to email filtering

Starting with version 6.2.2, AXIGEN can integrate with a Milter filter at the SMTP level, enabling scanning of the incoming SMTP connections.

The possibility to scan a message before receiving it opens up new perspectives. In regards to unwanted traffic, it enables us to refuse a message if the content filters strongly indicate that the scanned message is unwanted.

How does this happen and why is this a better choice? Let's follow a hypothetical SMTP conversation between a bot sending malware and an AXIGEN server configured to not accept mails with viral attachments:

```
~$ telnet a.mx.axigen.com 25
Trying 213.233.121.10...
Connected to a.mx.axigen.com.
Escape character is '^]'.
220 This is AXIGEN
ehlo example.com
250-node10 AXIGEN ESMTP hello
250-PIPELINING
250-AUTH PLAIN LOGIN CRAM-MD5 DIGEST-MD5 GSSAPI
250-AUTH=PLAIN LOGIN CRAM-MD5 DIGEST-MD5 GSSAPI
250-8BITMIME
250-BINARYMIME
250-CHUNKING
250-SIZE 10485760
250-STARTTLS
250-HELP
250 OK
mail from: <test@example.com>
250 Sender accepted
rcpt to: <support@axigen.com>
250 Recipient accepted
data
354 Ready to receive data; remember <CRLF>. <CRLF>
Date: Wed, 28 Jan 2009 11:57:40 +0200
From: test@example.com
```

```
To: support@axigen.com
Subject: test eicar
Mime-Version: 1.0
Content-Type: multipart/mixed;
  boundary="Multipart=_SqSsnwYqGdp5K"
```

This is a multi-part message in MIME format.

```
--Multipart=_SqSsnwYqGdp5K
Content-Type: text/plain; charset=US-ASCII
Content-Transfer-Encoding: 7bit
```

just a test

```
--Multipart=_SqSsnwYqGdp5K
Content-Type: text/plain;
  name="eicar.txt"
Content-Disposition: attachment;
  filename="eicar.txt"
Content-Transfer-Encoding: 7bit
```

```
X50!P%@AP[4\PZX54(P^)7CC)7}$EICAR-STANDARD-anti-virus-TEST-FILE!$H+H*
```

```
--Multipart=_SqSsnwYqGdp5K--
```

```
.554 5.7.1 Rejecting because message seems to carry VIRUS Eicar-Test-
Signature
quit
221-node10 AXIGEN ESMTP is closing connection
221 Good bye
Connection closed by foreign host.
```

You may see in the above SMTP conversation that the message was not accepted by the AXIGEN server because the Eicar-Test-Signature virus was detected during the SMTP session. How was that possible? Well, after receiving the final dot in the DATA stage of the SMTP connection, AXIGEN passed the received data to a Milter filter which disassembled the multi-part message and scanned the eicar.txt attachment with an antivirus. The antivirus detected the Eicar signature and told the Milter filter to signal to AXIGEN that the message should be rejected, which it did, with a verbose 5xx error: "Rejecting because message seems to carry VIRUS Eicar-Test-Signature".

Why would this be a better choice than quarantining or tagging? For one thing, we do not accept the message and no further resources are allocated to this mail: processing, storage, backup, double-checks etc. For unwanted traffic, this is a very good thing as it minimizes your problems. However, what happens if, unfortunately, the refused message is a legitimate mail? Let's compare the three valid choices:

- **Quarantining** means that the unwanted mail would end up in a rather large quarantined space. Suppose we only do this for malware as detected by an antivirus content filter: viruses, worms, phishing. Can we alert the receiver for every quarantined mail that was heading to their inbox? Realistically, no, because the malware traffic can reach really high levels. Add to this that almost all of it is spoofed and you risk to get into situations in which innocent people are blamed for spreading malware, when their only fault would be the fact that their address is known by spammers. So, when a mail is quarantined, neither the sender, nor the recipients are usually aware of it. If through some other means one or both of them find out about the missing mail, typically the receiver will have to alert his/her administrator of the mishap in order to gain access to the quarantined mail.

- **Tagging**, often applied to spam messages, means the message will usually end up in a sub-folder of the recipient's mailbox, typically named "Spam". He/she may or may not check that folder for false positives, but as no filter is perfect, sooner or later some legitimate mail will end up in "Spam". When that happens, neither the sender, nor the recipient will be aware of it. If through some other means, one or both of them find out about the missing mail, the receiver will usually have to dig through its spam folder to find the legitimate mail. This may be quick if he/she knows the exact coordinates of that mail (sender, date, subject), or may be a daunting task if the "Spam" folder is rather large and the data is very vague (eg. "Should have received a mail from a South American company with some financial info").
- **Refusing** at the SMTP level would almost instantly alert the recipient, because the sender's mail server has to generate a Non-Delivery Report. As you may see in the above SMTP conversation that we have used as an example, the AXIGEN server does not accept the mail; it rejects it at SMTP level, which means that the other mail server, the one that delivers the mail in the name of the sender, will have to immediately generate a Non-Delivery Report (NDR) back to the sender.

Keep in mind that the vast majority of malware or spam messages is traffic for which the only accountable information is the origin of the connection (the IP); everything else may be spoofed: be it the sender envelope, the **From** header etc. Spoofing leads to a lot of confusion among users, so minimizing the mail traffic that enters your mail systems goes a long way towards minimizing the impact of spoofing.

□ Final thoughts

This was just one simple example, in which the decision to reject a mail was taken by the Milter filter by evaluating the results of the antivirus scanner. But a Milter filter may take such decisions based on a multitude of checks: antivirus filter, spam filters, blacklists, spoofing checks, ehlo checks, grey-listing, spam traps etc. It can also be used to integrate the AXIGEN Mail Server with the [MIMEDefang](#) email filtering tool to help you reduce the number of spam messages processed by the server.

To find out more about this innovative approach on email filtering, [contact us](#) at any time.

Additionally, if you'd like to find out more about how the AXIGEN Mail Server supports business growth by combining core messaging features (SMTP, POP, IMAP) with AJAX WebMail access, advanced groupware, integrated security tools and full Web-based control panel where administrators can manage and configure the system, sign up for our free live webinar organized in partnership with The Radicati Group: ["Linux Messaging for Large Scale Enterprise and Service Provider Environments"](#).

AXIGEN Copyright © 2009 Gecad Technologies SA [AXIGEN]. All rights reserved.

This material or parts of the information contained herein cannot be reproduced in any form or by any means without the prior written permission of AXIGEN. The product and the documentation that comes with the product are protected by AXIGEN copyright. [AXIGEN \(http://www.axigen.com\)](http://www.axigen.com) reserves the right to revise and modify its products and documentation according to its own necessities, as well as this document content. This material describes a status, as it was in the moment this material was written and may not correctly describe the latest developments. This material is provided on an "as-is" basis.