

Antispam Evaluation Guide

White Paper

Table of Contents

1 Testing antispam products within an organization: 10 steps	3
2 What is spam?.....	4
3 What is a detection rate?.....	4
4 What is a false positive rate?	4
5 Quickly evaluating the effectiveness of a spam filter	5
6 Comparing different spam filters.....	5
7 Comprehensive spam filter testing (for one or more filters).....	6
8 Common errors in evaluating spam filters	8
9 Common user errors with Kaspersky Anti-Spam.....	10

Users often run into difficulties evaluating spam filter systems when selecting an antispam product or when using an antispam product. It may seem like an easy task, but experience has shown that those evaluating spam filters can actually run into a number of stumbling blocks.

We at Kaspersky Lab recommend familiarizing yourself with this document before you begin testing and evaluating an antispam product. The criteria listed below should apply to all contemporary antispam solutions.

We hope that we can help you to make the right choice!

1 Testing antispam products within an organization: 10 steps

The following steps should be followed in order to ensure that you have selected the antispam solution which best meets your needs.

1. Make your own clear definition of spam and determine which parameters you plan to use to assess the spam filter(s) (detection rate, false positive rate, etc.).
2. Identify the products that you will be comparing.
3. Install the products on a test board (an email server) using settings which are as close as possible to one another across all products.
4. Make sure the products are set up to receive updates in line with vendor recommendations.
5. Forward email traffic to the test server.
6. Use email traffic taken from at least 20 accounts.
7. Conduct testing for at least two weeks.
8. Become familiar with the errors that are commonly made when evaluating spam filters (see below) and how to avoid them.
9. When you have completed the testing process, tally up the number of spam messages that were detected and the number that slipped through the filter, which should help determine the detection rate for each product (the number of detected spam emails compared to the total number of all spam messages received).
10. Tally up the number of false positives and calculate the false positive rate for each product (the number of messages erroneously identified as spam compared to the number of all legitimate messages). Remember that the delivery of email that you need is more important than a few spam messages that get past the filter.

These points are discussed in more detail below.

2 What is spam?

Before evaluating an antispam product, it is important to understand exactly what spam is.

Spam is unsolicited, mass, anonymous electronic correspondence. It is important that these three factors are included in the definition. For example, a mass mailing that a user subscribed to in the past (and from which s/he can opt to unsubscribe) is not spam. Mass mailings that are unsolicited and that are sent manually from a known address (for example, a sales manager) are also not spam.

When evaluating the spam filters used in different antispam products, it is important to remember that not every unwanted email is actually spam.

3 What is a detection rate?

The detection rate (or the percentage of spam that is rerouted by a spam filter) is the only objective way to determine the effectiveness of a spam filter. *The detection rate is the number of emails defined as spam by the spam filter compared to the number of all spam messages in email traffic over a certain period of time.* In other words, if your personal inbox receives 100 spam messages in 24 hours, and 10 made it past the spam filter, then your antispam's detection rate is 90%. If a company receives 100,000 spam messages over the course of a week, and 95,000 were blocked by the product, then the detection rate is 95%.

It is not possible to automatically calculate the detection rate, since you first need to know the number of all spam messages received on your server (which can be done manually). This is why determining the detection rate can be rather difficult, although it is important to do.

These days, a spam filter is considered effective if it has a detection rate of 90% or higher. However, the detection rate is not the only criterion for evaluating antispam solutions. Users should also consider the false positive rate.

4 What is a false positive rate?

A false positive occurs when a regular message is mistakenly identified as spam (as a result of which, the user does not receive the message). The false positive rate is the number of false positives divided by the number of all wanted correspondence in email traffic over a set period of time.

Just like the detection rate, the false positive rate must be calculated manually.

False positives are tricky, and just when you think you have done what you can to rule them out, they sneak up on you anyway. False positives depend greatly on the quality of the spam filter. If the product demonstrates a strong detection rate, but is always "making mistakes" on 1-2% of normal emails, this outweighs its strengths and can mean trouble for a company. A few extra spam messages a day is a small price to pay, but the failure to receive an email from your partner or client can have *very serious consequences*.

Some manufacturers resort to the following tactic: they say the false positive rate is zero (which is highly unlikely) and then, in the fine print, they specify that it is zero only when a spam folder has been designated. This is where all detected spam is saved, which is a useful function, especially in terms of false positives. However, if you force a user to check this folder daily for false positives, it means that the spam filter is not doing its job and that the user must do most of this work manually.

An acceptable false positive rate is 0.001% or less (one false positive per every 100,000 messages). The best solutions have filters that demonstrate these numbers as they are recorded (not in marketing materials, etc.).

5 Quickly evaluating the effectiveness of a spam filter

It is difficult to evaluate a spam filter's detection rate quickly, especially in large organizations or for large email systems, mostly because it involves a great deal manual calculation.

With this in mind, one can use an email account that receives *only* spam. Many organizations have an account like this, such as info@domain, webmaster@domain, etc. Once the filter settings are specified for this account and the user has been able to evaluate the number of spam messages that are successfully filtered, then s/he can make an *approximate* calculation of the detection rate.

Let's say that the webmaster@domain email receives an average of 99.9% spam. Over a test period of one week, the account received 1,000 emails, 920 of which were filtered out as spam. That means that the approximate detection rate (since we are not calculating the exact detection rate and we do not know how many "real" email messages are in the mailbox) is 92%.

Always keep the following in mind:

- When using this kind of approach, the false positives are not added up and are not considered.
- The nature of spam changes with time, and a testing period of less than two weeks will give very rough approximations.

The nature of spam may depend on a specific email account, and the profile of the specific email account chosen for testing purposes may differ from the profile of most email accounts on an email system.

6 Comparing different spam filters

Despite all of the drawbacks to the approach listed above, it can be used in order to gain a basic idea to compare detection rates (*but not false positive rates*) between different antispam products.

In order to do this, you will need to install several spam filters for the comparison. Each filter should be subject to a test email traffic flow, which includes 99.9% spam. It is important that the same emails all go through each filter. Then you can compare what made it past each filter.

All of the drawbacks described above apply to this approach as well.

Similar fast methods may be used to gain a rough estimate of the situation in order to make a decision without having to undergo a lengthy full-scale testing process described below.

7 Comprehensive spam filter testing (for one or more filters)

If you are ready to dedicate the time and the effort (which is worth it) in order to obtain an adequate, precise evaluation of the quality of your antispam solution(s), follow the steps below to test the performance of one or more spam filters.

Filtering email traffic in real time

Often, when testing a product, a system engineer will subject it to a set collection of spam messages (for example, a spam-heavy inbox with emails collected over the past year) and review the number of emails that were caught by the filter. A quick-and-dirty approach to this kind of testing will only result in an inaccurate assessment of the product's quality.

Each vendor strives to ensure that their antispam solution will identify all of today's spam as it is sent. Thus catching spam that was sent a year ago and that will never be sent again seems like a pretty useless approach. That is why the results of testing on a collection of spam are very unpredictable and is substantially different from testing spam filters in real time, and later, when fully operational.

In order to achieve reliable results, spam filter testing must be carried out in real time *only*, using real emails as they are sent and delivered.

Test for at least two weeks

The contents of spam change over time. A specific spam campaign may last anywhere from 1-2 weeks and can influence the content of spam on the Internet. Different filters do better with some types of spam, while they perform more poorly with others. Correspondingly, a spam filter's performance may vary from day to day, depending on the type of incoming spam.

In order to obtain an objective evaluation, you should test the product(s) for at least two weeks.

Use at least 20 email accounts for testing

The contents of spam messages depend on the email address to which they are sent. An address such as info@domain may see all kinds of different types of spam, while individual employee addresses may see more specialized types of spam – it all depends on how, when and which spammer databases get which addresses. Likewise, the detection rate for each account may differ.

In order to assess the entire email system as a whole, we recommend testing your antispam solution using at least 20 different email accounts simultaneously.

Messages should be sent and received without any errors

Testing should use messages that are sent from different addresses to the server hosting the product. However, you must also make sure that the correspondence has been set up in such a way that emails arrive in the inbox without any errors (conversion of text, subject headings, etc.). This is especially relevant when it comes to emails that are initially received by a Microsoft Exchange Server.

This is a very important factor, since errors can both increase or lower test results.

Equal testing conditions

Settings for all of the products undergoing a comparative test should be as consistent as possible across all solutions:

- Each product should filter the same email traffic.

- If two or more filters use the same method(s) for identifying spam (such as DNSBL), this option should be selected or unselected for all of the products.

The frequency of product updates should also be set so that it is in line with vendor recommendations.

8 Common errors in evaluating spam filters

“I get a lot of spam in my inbox.”

As strange as it may seem, the number of spam messages that end up in a user's inbox over the course of one day does not really tell you how efficient the spam filter is.

For example, let's say a user receives 10 spam messages and is unhappy about it. However, if s/he examines the situation more closely, s/he will find that 100 emails were actually filtered out that day. That means that the detection rate is $100/(100 + 10) = 91\%$. What would s/he have thought if all of that spam has been delivered to his/her inbox?

A spam filter should block 100% of all spam

Many people think that if you pay for protection from spam, then you should not ever see spam again. Unfortunately, that is impossible, and there is no solution capable of blocking 100% of all spam.

It is important that a spam filter does not result in false positives, which would mean the loss of important correspondence, which means a user is forced to be much more careful about their incoming email. That is, it is much better to receive one spam message than to block any wanted or necessary correspondence.

Because of these idiosyncrasies, no spam filter can identify 100% of all spam. It is more important to have a relatively high detection rate and a zero false positive rate.

“We tested using our spam collection.”

As explained above, testing the detection rate of spam solutions using spam that was initially sent a year ago and which will never be sent again is an exercise in futility. This is why no serious antispam solution is developed using old spam. An antispam product must block spam that is being sent today, right this very minute. Testing with a collection of old spam will provide inaccurate results.

“The percentage of spam is just 50%”

Users often do not know exactly what the detection rate is and assume it is the percentage of spam that is found in the inbox.

For example, let's say that a user receives 100 emails – both spam and wanted correspondence. The spam filter caught 50% (or 60%, or 40%, and so on) of all of the spam messages. But, this number does not have anything to do with the quality of the spam filter's performance, since we do not know how many spam messages were in those 100 emails, and how many of those emails were wanted correspondence.

In order to determine just how well a product performs, you will have to apply the testing methods described above.

“Only system administrators understand how antispam products work.”

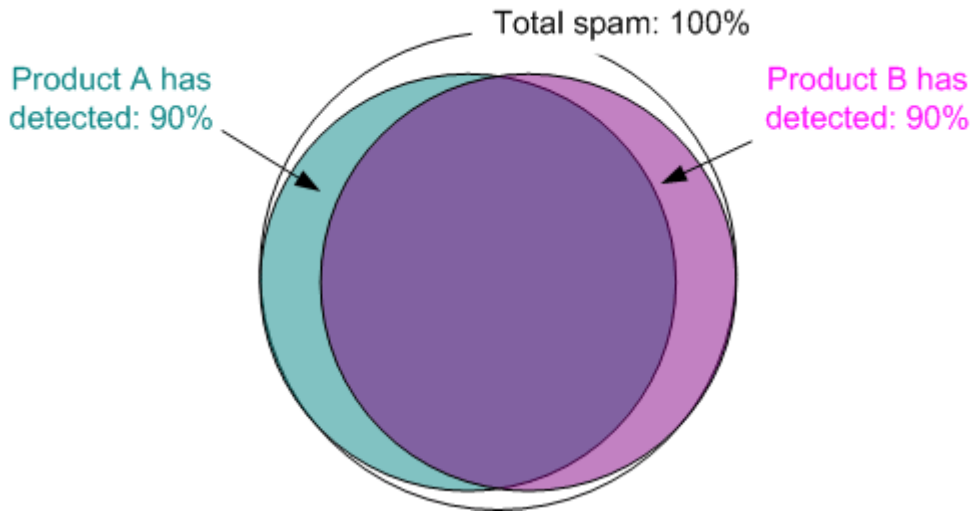
Many antispam solutions use a marking system for emails that have been checked. A message that is identified as spam is also marked as such, and they are delivered to the user with that identifier. The user's email client can be set to automatically file those messages in a separate spam folder.

If you install this kind of product in an organization, but do not tell users all of the information, such as what the “SPAM” identifier means, why it appears, and how to change your settings to automatically file these messages, etc, the results for most users will be disappointing.

“Filter A catches only 10% of what gets through filter B”

Sometimes, when two or more products undergo comparative testing, they are installed one after the other. The email that goes through the first filter is then filtered by the second filter, and so on. As a result, the second product filters what got past the first. This is an unreliable method, because the second filter's performance is

then based on a very small portion of all email traffic (the products ought to be tested using the same traffic – see the section on Equal Testing Conditions). Testing antis spam products that are installed one on top of the other will give you nothing but inaccurate results.



Different products use different methods, and even if they work in the same way – for example, they both have a 90% detection rate – that 90% will be different for each product (see the illustration above).

9 Common user errors with Kaspersky Anti-Spam

Scheduling updates for less than once every 20 minutes

The more often a product is updated, the faster it will react to new spam. If you change the settings to schedule updates as recommended by Kaspersky Lab, then you will be using the most up-to-date algorithms, spam analysis rules, and signatures.

Turning off UDS technology

Kaspersky Anti-Spam 3.0 uses UDS technology, which allows real-time reaction to spam messages. It is extremely important to make sure that this feature is active, since it filters out the fastest spam mailings and the types of spam that are the most difficult to filter.

Turning off the RBL filter

Using the blacklists of spammers' IP addresses can help improve the solution's detection rate.

Using an outdated version of the product

Spammers are constantly improving upon their previous methods and are always coming up with new tricks to get around spam filters. You can ensure that you have maximum protection against spam by installing the latest version of Kaspersky Anti-Spam, since it uses the most up-to-date methods and algorithms to identify spam and counters the latest tricks used to get around spam filters.