What nobody told you about how to interpret Spam filter scores

MailMonitor
Email Deliverability Simplified

Table Of Content

It’s time to put on our referee hats and take a look at those Spam filter scores.

First of all, let’s talk about the term: Spam Filter Scores. To put it simple, this score is determined by the number of emails that land either in your inbox compared to the number of emails that land in the spam folder. The higher the score is, the larger is the number of spam emails.

In this article we’ll interpret the Spam Filter Score of these three companies: Spam Assassin, Symantec, Barracuda and Brightmail.

Spam Assassin

Let’s start with the professional killer of spam, Spam Assassin, the popular filter from the Apache Software Foundation.

The scores that Spam Assassin provides are pretty straight forward. There are only two options: if the score is anything above the value of 5, then the software indicates “spam” and if the score is anything below the value of 5, then there are no spam messages identified.

Of course, there is an option to set a higher or lower value, depending on your own personal restriction preferences but just to be on the safe side, it’s indicated to maintain the score under 5.

The software produces a header reveals a lot on information and details about the score it provided.

Let’s take a look at it and analyze the text.

```
X-Spam-Checker-Version: SpamAssassin 3.2.5 (2008-06-
10) on JetWeb
X-Spam-Level:
X-Spam-Status: No, score=-0.4 required=5.0 tests=ALL_
TRUSTED,AWL,DKIM_SIGNED,
DKIM_VERIFIED,HTML_MESSAGE,URIBL_BLACK
autolearn=disabled
version=3.2.5
X-Spam-Report:
* -1.4 ALL_TRUSTED Passed through trusted hosts only
via SMTP
* -0.0 DKIM_VERIFIED Domain Keys Identified Mail:
signature passes
* verification
* 0.0 DKIM_SIGNED Domain Keys Identified Mail:
message has a signature
* 0.0 HTML_MESSAGE BODY: HTML included in message
* 2.0 URIBL_BLACK Contains an URL listed in the URIBL
blacklist
* [URIs: websitehere.com]
* -0.9 AWL AWL: From: address is in the auto white-
list
```

This may look confusing but there's nothing much to it at all.

1. The first line indicated the version of the software.

2. The second line would display a * for each point the email gets. But in this case, it didn't detect anything wrong.

3. The third line tell you directly if the message is spam or not and gives you the exact score.

4. The fourth one indicates the broken-down test that the software did on the text.

Now that we know how to read the score provided by Spam Assassin, let's move on to someone else.

Symantec

The score provided by **Symantec** (full name Symantec Messaging Gateway) is calculated a little bit different than **Spam Assassin**. **Symantec** calculates each email and giving it a score from 1 to 100.

In an email receives a score between 90 and 100 it is definitely spam. You can also set a range between 25 and 89, the messages falling in the personalized value, will be considered spam.

The only time a message is not considered spam is if it scores below the threshold.

Pretty easy right? But the main advice that the company gives is to not adjust the spam threshold until you have had decent exposure to the filtering patters. And even then, the value you modify would only slightly be lowered, 1 point to up to 5 points each week.

The third entry on our list is devouring spam left and right. Let's talk about Barracuda.

Barracuda

Ready to hook some spams? Well, if a message gets a value over 10, then, you know it, it's definitely spam. Best would be to aim for a score below 3.5 .

Based on this score, the software can tag, quarantine, block or allow (for outbound) the scanned message.

To provide a score and give a precise answer, Barracuda has to go through a number of steps, but they're all worth it.

1. Checking the IP Block list just to make sure the sender IP is not already suspicious.

2. Checking for viruses.

3. Double check for viruses. (you know a filter is good when it's checking twice.)

4. Check to see is the text contains any words or codes that the used specifically marked them as spam.

5. Checking "Spam Fingerprint". This is a fancy way if saying that it checks whether or not the email has been fingerprinted before. When a message has been marked as spam, the information is sent to the Barracuda Central.

6. Intention Analysis.

7. Bayesian Spam Analysis.

8. Spam rules-Based Scoring.
The third entry on our list is devouring spam left

Just like **Spam Assassin**, **Barracuda** also created a header that looks like this:

```
X-Barracuda-Start-Time: 1332864901
X-Barracuda-URL: http://172.26.14.249:8000/url-mod/address.com
X-Barracuda-Bayes: SPAM GLOBAL 1.0000 1.0000 4.3430
X-Barracuda-Spam-Score: 2.03
X-Barracuda-Spam-Status: No, SCORE=2.03 using global scores of TAG_
LEVEL=3.0 QUARANTINE_LEVEL=5.0 KILL_LEVEL=7.0 tests=BSF_SC0_SA_TO_
FROM_DOMAIN_MATCH, BSF_SC7_SA578_CH, DATE_IN_PAST_12_24, DATE_IN_
PAST_12_24_2, HEAD_LONG, HTML_MESSAGE, LONG_TERM_PRICE
X-Barracuda-Spam-Report: Code version 3.2, rules version 3.2.2.92409
pts RULE_NAME description
---- ---------------------- ----------------------
0.50 HEAD_LONG Message headers are very long
0.01 DATE_IN_PAST_12_24 Date: is 12 to 24 hours before Received: date
0.21 LONG_TERM_PRICE BODY: LONG_TERM_PRICE 0.00 HTML_MESSAGE BODY:
HTML included in message
0.50 BSF_SC7_SA578_CH Custom Rule SA578_CH
0.01 BSF_SC0_SA_TO_FROM_DOMAIN_MATCH Sender Domain Matches Recipient
Domain
0.80 DATE_IN_PAST_12_24_2 DATE_IN_PAST_12_24_2
X-SA-Exim-Connect-IP: 12.237.60.52
X-Spam-Checker-Version: SpamAssassin 3.3.1 (2010-03-16) on server-4
X-Spam-Level: -0.5
X-Spam-Status: No, score=-0.5 required=5.7 tests=BAYES_00,DATE_IN_
PAST_12_24, HTML_MESSAGE,LONG_TERM_PRICE,L_BILLS,L_TAX1,T_LOTS_OF_
MONEY autolearn=no version=3.3.1
```

Let’s talk about what it means, starting from the first line to the last, in that order.

First of all, you’ll be shown the Spam score.
Next, the header indicated whether or not the message has been indicated as Spam or not, along with other information like the score, and tests.

You’ll also see the threshold that was set up and the action taken on the specific email.

Last but not least, let’s brighten up our day and talk about how the scoring works for **Brightmail**.

Brightmail

Similar to the other filters, Brightmail calculates the score by giving a certain text a value ranging from 1 to 100 but even a score of 3 could raise suspicions.

The software gives you the option to play around with the numbers, depending on how aggressive you aim to be. You can define a range of scores from 25 to 89, meaning that the email is “suspected spam”. Also, just to be safe, you should only modify this range when you gain certain traffic experience and only then gradually lowering it by as much as 5 points per week.

The user can also help the software perform better by enabling certain features like:

1. **Uniform Resource Identifiers (URI) reporting.**

This sends a report that contains URIs that are scanned as spam. Ultimately all this information is used to develop new filters.

2. Probe accounts.

By forwarding unusual email accounts or invalid addresses, the software uses them to attract spammers. Finally, these spam messages are used to create better filters.

These were the four titans of filtering spam and making the online world a better and peaceful place.

Going through all these headers and numbers or letters, we found out it’s not that difficult to understand a score and now we have a better idea on how to protect ourselves from spammers.

See You Next Time!

MailMonitor
Email Deliverability Simplified

@mailmonitor
www.mailmonitor.com