

J3 Web Services Spam Filtering in SmarterMail

Spam filtering in SmarterMail uses a number of different technologies. Below is an explanation of each of those technologies, followed by an explanation of how they all work together.

Spam Checks and Weights

For each spam check, SmarterMail uses the configured numerical weight for a particular check to tally up the overall spam score for a particular message.

Blacklists – Several blacklists of IP addresses that are known spammers exist for checking the IP address of the incoming mail server. If the IP of the incoming mail server is found on the list, that list's configured spam weight is added to the spam score of the message. J3 uses four different black lists:

- CBL, cbl.abuseat.org, weight: 5
- PSBL, psbl.surriel.com, weight: 5
- SpamCop, bl.spamcop.net, weight: 10
- Spamhaus, sbl-xbl.spamhaus.org, weight: 10

Reverse DNS – When you type in “www.google.com” in your web browser, a special server known as a Domain Name Service (DNS) server tells your computer what IP address to go to. Comparatively, per Internet IP address guidelines, each IP address should have a corresponding domain name associated with it. This IP to domain name mapping is known as Reverse DNS. Many mail servers used for spamming do not follow this Reverse DNS guideline. If an incoming message's sending server does not have a Reverse DNS entry, the message is assigned a weight of 30.

Sender Policy Framework (SPF) – A relatively new technology, SPF attempts to combat email address spoofing. The owner of a domain puts a special entry in the DNS zone for their domain that tells others what IP addresses are authorized to send mail for that domain. So if an incoming message says it's from joesmith@hotmail.com, but it is coming from an IP address that is not an authorized mail server for the domain hotmail.com, then the receiving mail server knows that the message's from address is spoofed. Not all domains have adopted this technology, so it is not as effective as it could be.

- SPF Pass, or no SPF record – If the SPF check passes or the domain has not defined an SPF record, the message is assigned a weight of 0
- SPF Error – If the SPF record for the sending domain exists but is not configured properly, the message is assigned a weight of 5
- SPF Soft Fail – A soft fail means that the IP address of the sending server is probably not authorized to send for that domain, but the results of the check were not conclusive. Weight: 20
- SPF Fail – The IP address of the sending server is definitely not authorized for the domain that the from address claims to be from. Weight: 30

Bayesian Filtering – This method of statistically analyzing the content of a message, looking for a certain percentage of typical spam words relative to the length of the entire message, was introduced in SmarterMail 3. Historically, it has not been very accurate, allowing a lot of obvious spam to get through. However, when it does tag a message as spam, it almost always is, so J3 uses this setting as one more line of defense. Weight: 10

SpamAssassin – Arguably one of the best, most accurate spam filters on the market. SpamAssassin gives each message a spam weight based on the content of the message and how likely that content is to be spam. Weight: SA produces a scale from 0 on up. There isn't a set top limit, and some messages have reached as high as 50 with this one setting.

SMTP Blocking

SmarterMail uses SPF, Reverse DNS, and the blacklists for blocking an incoming message from being accepted if its combined spam weight from these checks is over the block threshold. Currently, the block threshold is set to 50, meaning if the combined weight of these checks is 50 or more, the message is refused and an error is sent to the sender. This level was chosen to prevent any single check from triggering a block. Based on the weights assigned to these individual checks, if the combined weight is 50 or more, the certainty that the message is spam is virtually 100%.

Greylisting

This is a complex, yet very powerful spam blocking tool. The basic premise is that it is not cost effective for spammers to retry messages to addresses that return a temporary error, while nearly all legitimate mail servers are set to retry delivery for a message that receives temporary error. The following terms are used in the description of greylisting:

Message Triplet: The combination of the from address, to address, and sending server's IP address for a given message

Block Period: The amount of time messages with unknown triplets are told to wait before retrying delivery (currently set to 9 minutes)

Pass Period: The amount of time a mail server has to retry a message before having to start the process over (currently set to 12 hours)

Record Expiration: Upon the successful retry of a message in the Pass Period, the amount of time that subsequent messages with the same triplet are passed unhindered (currently set to 45 days)

Here's how the process works:

1. SmarterMail records the from address, to address, and sending server's IP address for each incoming message. This information is known as the Message Triplet. For each unknown triplet, the sending message is told to wait the period of time specified in the Block Period setting.
2. Once the Block Period has expired, the triplet is held for the Pass Period. If it is seen again, meaning the sending mail server retries delivery of the message, the message is accepted for immediate delivery.
3. Upon the successful delivery of a message with a specific triplet, any incoming message received with the same triplet during the Record Expiration period will be accepted for immediate delivery. Additionally, subsequent messages with the same triplet cause the Record Expiration period for that triplet to be reset to its initial value.

Putting it all together

New messages flow through the following path as they are delivered to SmarterMail:

SMTP Block Check → Greylist Check → Spam Check → Message Delivery

SMTP Block Check and Greylist Check are explained above. The Spam Check phase runs the message through each of the checks listed above in the Spam Checks and Weights section. The weights for each check are then tallied to give the message its final spam weight. At the Message Delivery phase, the message is processed according to its spam weight. SmarterMail defines three spam probability levels and automated actions that can be taken at each level. SmarterMail defines a Low probability, Medium Probability, and High Probability of being spam, and these probabilities are set to spam weights of 10, 20, and 30 respectively. By default, anything with a spam weight of 10 or higher is sent to the recipient's Junk Email folder on the server, accessible through the SmarterMail web interface at <http://mail.j3web.com>.

Changing the Defaults

Some of the settings described above can be changed from their default values.

SMTP Blocking – this setting is only at the server level and is the same for everyone. In the unlikely event that legitimate email is blocked because of this setting, the block threshold will be increased to prevent it from happening again.

Greylisting – Users can disable greylisting for their account. After logging in to the SmarterMail web interface at <http://mail.j3web.com>, click on Settings → My Settings. On the Mailbox tab, check “Bypass Greylisting”, then click “Save”.

Spam Check – Specific checks can be disabled on a per-domain basis. Requests to have certain checks disabled for a particular domain can be made by emailing support [at] j3web.com.

Message Delivery – Users can change the default actions for handling Low, Medium, and High spam. After logging in to the SmarterMail web interface at <http://mail.j3web.com>, click on Settings → My Spam Filtering. On the Options tab, click “Override spam settings for this account”. Then click the Actions tab. From there users can change the default actions. ****WARNING** If you choose to delete a message automatically as the result of it being classified as spam, there is no way to recover it. Do not use this setting unless you are willing to accept the risk that a legitimate message may be deleted.**

If you have any questions regarding the information here, contact J3 at support [at] j3web.com.