Model Email Service

An email service must be reliable, convenient, and above all secure. Email recipients rely on their email services to protect them from spam, viruses, forgery, identity theft, and even more serious crimes. We have gathered what we think are the best open-source programs to provide a secure setup. There are three basic components to this security – blocking a Denial of Service (DoS) attack, checking the identity and reputation of an unknown sender, and filtering messages from senders whose reputation is unknown.

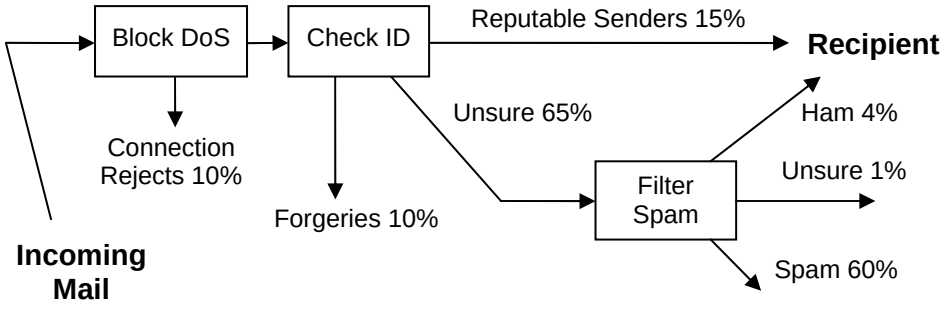

Default Settings

Check ID: Reputable Senders have less than 1 spam in: 100 emails

Filter Thresholds:

– IP Blacklist: - - - - - Moderate

– Spam if greater than: 75

– Ham if less than: 50

The Block DoS gate uses a blacklist of IP addresses that are currently being heavily abused. There are many IP blacklist services available. Each is a tradeoff between blocking addresses that are used for spamming and not blocking legitimate mail. We have chosen a list on the conservative end, since it is intended to block only the most voluminous sources of spam, not provide a complete anti-spam solution. Reaction time is important also, since DoS attacks usually start very suddenly.

The Check ID gate determines the identity and reputation of the sender. A query to the Registry of Public Email Senders provides information on most legitimate senders, including what methods the sender offers to authenticate their Identity, and ratings of that sender by various Rating Services. Rejection of forgeries is based on the policy of the ID owner, using the ID owner's authentication records. Thus no rejection of legitimate mail (except a very few at the Block DoS stage) will occur based on the Receiver's policy.

Acceptance of all mail from Reputable Senders is based on a threshold set by each Recipient. Spam haters can set the threshold high, and very few senders will qualify. Recipeints that need the utmost reliability in delivery of mail addressed to them will set a lower threshold, and tolerate a higher level of spam. Any mail from senders that don't authenticate, or that don't have an acceptable reputation, will go to the Spam Filter.

The Spam Filter uses a variety of methods to sort the remaining mail into three categories. These methods may include more aggressive IP blacklists, heuristic rules that identify common characteristics of spam, and statistical analysis of the message content. Recipient options control which blacklists are selected, and what thresholds are set for the three categories.

Modern spam filters do a good job of classifying most messages as clearly spam or clearly ham, so there are usually few messages in the Unsure category. The default is to accept the message if its spam score is less than 50, and send it to the spam bucket if greater than 75. This sets a wide margin for false rejects. Most recipients will reduce this margin after they gain some confidence that they are not seeing many false rejects with a score as high as 75.

A Recipient who prefers to use the spam filter that comes with his own email program, can set the ham threshold to 100 and bypass the shared filter above. A filter with individual training is likely to be more accurate, particularly for recipients who have "spammy" words in their normal email.

DMQ 01/07/06