

How Scores Are Assigned

How are the scores assigned?

The scores are assigned using a neural network trained with error back propagation ([Perceptron](#)). Both systems attempt to optimize the efficiency of the rules that are run in terms of minimizing the number of false positives and false negatives.

You can find all of the currently active rules and their scores in the Subversion repository under [/trunk/rules](#) or by downloading the latest published set using the sa-update tool.

You can help this system by providing statistics on your mail spool via [NightlyMassCheck](#) and [RescoreMassCheck](#).

Confusing scores

Scores for "learn" rules (example the various BAYES_?? rules) are scored using the same method. This can produce scores which seem incorrect (example BAYES_80 with a higher score than BAYES_99). This is due to the fact that rules are not related to one another, they're separate rules have separate scores.

Messages with high probability from a "learn" rule will most likely match other rules. This lets the score generation system lower the "learn" rule score preventing false positives. The message still is recognized as spam due to the sum of all rule scores.

Some DNS blacklist rules are distributed with scores of 0. These generally request or require payment are disabled by default. Feel free to enable the lookups, if you've paid for them.

A score of 0 will stop a rule from being run.

In version 2.x, the scores are assigned using a genetic algorithm (GA).