

# BayesExpiration

The Bayes database stores up to a certain number of tokens, configured via `bayes_expiry_max_db_size` in `local.cf` (default: 150000 tokens).

Each token has an access time which records when it last contributed to a classification or appeared in a learned email. A mixture of obsolete (often ephemeral) tokens and the most-infrequently seen tokens are occasionally purged, according to a schedule and algorithm explained in the [sa-learn documentation](#).

Thus, even if you force an expiry run every month, it doesn't mean that you only have a month of data; the most important tokens never get purged.

To view the access time of the oldest token in the database: `date -r {{sa-learn --dump magic | grep "oldest atime" | cut -f 3 -w}}`

[partially adapted from a [post](#) by RW to the spamassassin-users mailing list]