

PreflightBuildBot

Preflight mass-checks buildbot

The preflight mass-check buildbot is up and running at <http://buildbot.spamassassin.org:8011/> and also proxied at <http://buildbot.spamassassin.org/preflight/> (although this has buffering issues due to a bug in httpd mod_proxy).

Every time something is checked into SVN, this will wake up and immediately start running mass-checks using that latest code and rules.

The corpus it mass-checks is split in a certain way so that results will be available very quickly – typically in under 10 minutes – with increasing quantities of results becoming available as time elapses.

Progress of the mass-checks are visible on [the Buildbot 'waterfall'](#); as they complete, their results become visible on the [RuleQaApp](#).

The preflight mass-check corpus

This corpus is built from a selection of mail rsync'd up from various people; it's then "smoothed out" into several subsets. These use differing amounts of mail, starting with a small set of mail in the "mc-fast" chunk, and gradually increasing until we get to the largest block in "mc-slower". This division means that early "fast" results can arrive quickly, with less to scan, and as time goes on, more and more of the "slower" slaves complete their mass-checks and upload the results.

The "smoothing" and subset selection happens in mass-check nowadays.

What happens during the preflight buildbot process

[As you can see](#), there are four steps performed by each buildbot slave, as follows:

Update: This performs an 'svn update' to load the latest code.

Configure: runs 'perl Makefile.PL' and 'make' to compile the rules.

Test: the mass-check takes place here. This is usually the time-consuming part.

Configure: a final summarisation step; first off, a 'FAST FREQS REPORT' is output, the [HitFrequencies](#) from the mass-check. Next, the logs from the mass-check are copied to a safe location, and the 'corpus-hourly' script run to generate various reports from them for the [RuleQaApp](#). The URL for viewing the results in the [RuleQaApp](#) is printed prominently.

Uploading corpora

See [UploadedCorpora](#).

Admin: Creating a new buildbot slave to perform mass-checks

I don't see us needing to do this anytime soon, but it's worth recording. Here's the commands that do this, run in the zone.

```
PASSWORD=[randompassword]
NAME=mc-new

sudo mkdir -p /home/bbmass/slaves/$NAME
sudo chown buildbot /home/bbmass/slaves/$NAME
cd /home/bbmass/slaves/$NAME
sudo su buildbot -c "buildbot create-slave --usepty=0 \
    /home/bbmass/slaves/$NAME \
    buildbot.spamassassin.org:9988 $NAME $PASSWORD"

echo $PASSWORD > $HOME/pwd
sudo mv $HOME/pwd /home/buildbot/pwds/$NAME
sudo chown buildbot /home/buildbot/pwds/$NAME
sudo chmod 600 /home/buildbot/pwds/$NAME

sudo vi /home/buildbot/bots/bbmass/master.cfg

    [search for mc-fast and add new lines/entries for $NAME]
    [don't forget the 'scheduler' part!]

sudo vi /etc/init.d/bbmass

    [search for mc-fast and add new lines/entries for $NAME]
```

(history: this was planned at [RulesProjBuildBot](#))