Fighting spam for fun and profit

the long road to SpamAssassin 4.0

Giovanni Bechis
<gbechis@apache.org>

Fosdem 2019, Brussels

SNB Information Technology & Web Solutions

SA as a framework

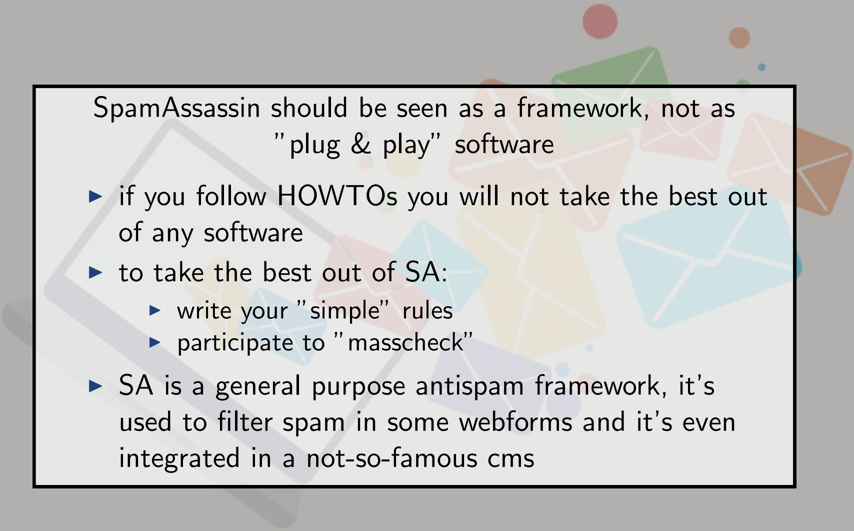

SpamAssassin should be seen as a framework, not as ”plug & play” software

- if you follow HOWTOs you will not take the best out of any software
- to take the best out of SA:
  - write your ”simple” rules
  - participate to ”masscheck”
- SA is a general purpose antispam framework, it’s used to filter spam in some webforms and it’s even integrated in a not-so-famous cms

What's Masscheck ?

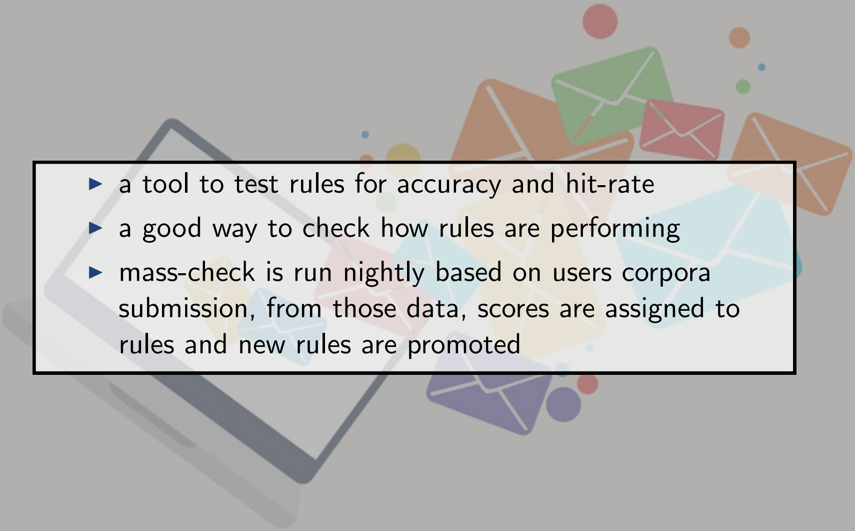

- a tool to test rules for accuracy and hit-rate
- a good way to check how rules are performing
- mass-check is run nightly based on users corpora submission, from those data, scores are assigned to rules and new rules are promoted

Checking how rules are performing

RuleQA:

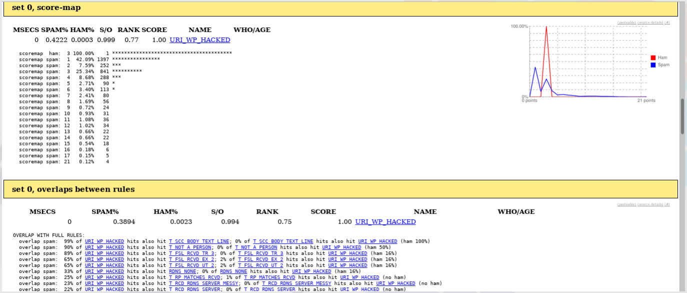

set 0, score-map

(pasteable) (source details) (#)

MSECS	SPAM%	HAM%	S/O	RANK	SCORE	NAME	WHO/AGE
0	0.4222	0.0003	0.999	0.77	1.00	URI_WP_HACKED	

```
  scoremap  ham:  3 100.00%    1 ****************************************
  scoremap spam:  1  42.09% 1397 ****************
  scoremap spam:  2   7.59%  252 ***
  scoremap spam:  3  25.34%  841 **********
  scoremap spam:  4   8.68%  288 ***
  scoremap spam:  5   2.71%   90 *
  scoremap spam:  6   3.40%  113 *
  scoremap spam:  7   2.41%   80
  scoremap spam:  8   1.69%   56
  scoremap spam:  9   0.72%   24
  scoremap spam: 10   0.93%   31
  scoremap spam: 11   1.08%   36
  scoremap spam: 12   1.02%   34
  scoremap spam: 13   0.66%   22
  scoremap spam: 14   0.66%   22
  scoremap spam: 15   0.54%   18
  scoremap spam: 16   0.18%    6
  scoremap spam: 17   0.15%    5
  scoremap spam: 21   0.12%    4
```

set 0, overlaps between rules

(pasteable) (source details) (#)

MSECS	SPAM%	HAM%	S/O	RANK	SCORE	NAME	WHO/AGE
0	0.3894	0.0023	0.994	0.75	1.00	URI_WP_HACKED	

```
OVERLAP WITH FULL RULES:
  overlap spam:  99% of URI_WP_HACKED hits also hit T_SCC_BODY_TEXT_LINE; 0% of T_SCC_BODY_TEXT_LINE hits also hit URI_WP_HACKED (ham 100%)
  overlap spam:  90% of URI_WP_HACKED hits also hit T_NOT_A_PERSON; 0% of T_NOT_A_PERSON hits also hit URI_WP_HACKED (ham 50%)
  overlap spam:  89% of URI_WP_HACKED hits also hit T_FSL_RCVD_TR_3; 0% of T_FSL_RCVD_TR_3 hits also hit URI_WP_HACKED (ham 16%)
  overlap spam:  65% of URI_WP_HACKED hits also hit T_FSL_RCVD_EX_2; 2% of T_FSL_RCVD_EX_2 hits also hit URI_WP_HACKED (ham 16%)
  overlap spam:  65% of URI_WP_HACKED hits also hit T_FSL_RCVD_UT_2; 2% of T_FSL_RCVD_UT_2 hits also hit URI_WP_HACKED (ham 16%)
  overlap spam:  33% of URI_WP_HACKED hits also hit RDNS_NONE; 0% of RDNS_NONE hits also hit URI_WP_HACKED (ham 16%)
  overlap spam:  25% of URI_WP_HACKED hits also hit T_RP_MATCHES_RCVD; 1% of T_RP_MATCHES_RCVD hits also hit URI_WP_HACKED (no ham)
  overlap spam:  23% of URI_WP_HACKED hits also hit T_RCD_RDNS_SERVER_MESSY; 0% of T_RCD_RDNS_SERVER_MESSY hits also hit URI_WP_HACKED (no ham)
  overlap spam:  22% of URI_WP_HACKED hits also hit T_RCD_RDNS_SERVER; 0% of T_RCD_RDNS_SERVER hits also hit URI_WP_HACKED (no ham)
```

- score assigned to messages that has been hit by a rule
- ham/spam hit by a single rule
- rules that overlaps on a particular rule

What have SpamAssassin done in 3 years and a half ?

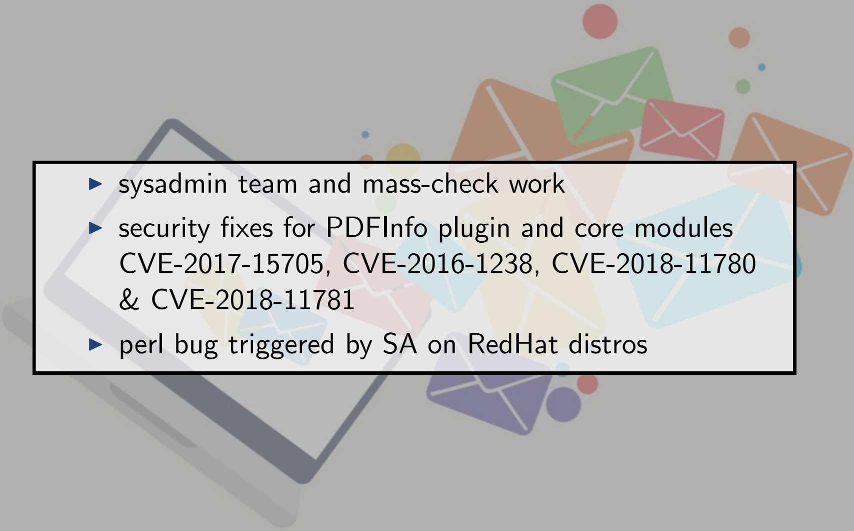

- sysadmin team and mass-check work
- security fixes for PDFInfo plugin and core modules CVE-2017-15705, CVE-2016-1238, CVE-2018-11780 & CVE-2018-11781
- perl bug triggered by SA on RedHat distros

What have SpamAssassin done in 3 years and a half ?

Assorted improvements:

- faster startup code and free(3) fixes for spamc(1)
- SSLv3 support removed from spamc(1)
- freemail antiforge improvements
- added possibility to score based on continents in geo-aware plugins
- improvements in URILocalBL plugin
- TxRep file descriptor leak fixes
- better check for http[s] mismatch plugin
- regression tests switched to Test::More

What have SpamAssassin done in 3 years and a half ?

HashBL plugin

The HashBL plugin is the interface to The Email Blocklist (EBL).

The EBL is intended to filter spam that is sent from IP addresses and domains that cannot be blocked without causing significant numbers of false positives.

What have SpamAssassin done in 3 years and a half ?

HashBL plugin

Checking Bitcoin scams with HashBL.

Using HashBL plugin (in 4.x release) or using an external plugin you can check if a Bitcoin address has been used for fraudolent purposes or not by asking via a dns query to $bitcoinaddress.bl.btcblack.it

What have SpamAssassin done in 3 years and a half ?

GeoIP2 support

Starting on 04/01/2018 Maxmind legacy geoip databases have been discontinued.

GeoIP2 support has been added to RelayCountry and URILocalBL plugins. In addiction RelayCountry supports also IP::Country::DB_File as an option.

What have SpamAssassin done in 3 years and a half ?

Anti phishing plugin

A new anti phishing plugin has been developed, it searches phishing uri in a database downloaded from PhishTank or from OpenPhish. More antiphishing databases will be added soon.

What have SpamAssassin done in 3 years and a half ?

Resource limits plugin

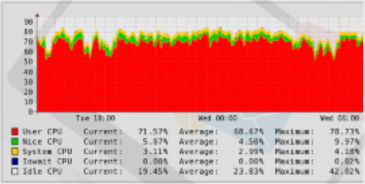

A new plugin that uses BSD::Resource perl module to assure your spamd child processes do not exceed specified CPU or memory limit.

What have SpamAssassin done in 3 years and a half ?

"From Name" spoof plugin

A new plugin that perform various tests to detect spoof attempts using the From header name section.

```
From:  "safeaddress@paypal.com"
       <hacked@hacked.eu>
```

International channels

International channels

Channels are a set of signed rules, they are important and very effective because standard rules are mostly base on english emails.

New features in SpamAssassin 4.0

"Ole Macro" detection plugin

A plugin have been developed to check if an email contains an Office attachment with a macro, it tries to detect if the attached macro is malicious or not.

New features in SpamAssassin 4.0

"Authentication-Results" parser plugin

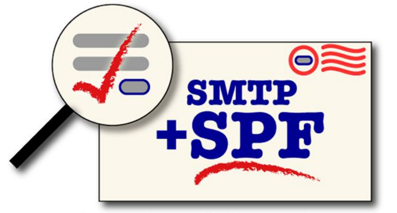

A plugin have been developed to check "Authentication-Results" header fields, it can supply the results obtained to other plugins, to avoid repeating checks that have been performed already.

KAM.cf rules: respond faster to spam

KAM.cf rules

MAPLE LEAF
READY TO EAT · PRÊT À MANGER
KAM
LUNCHEON MEAT / PAIN DE VIANDE
HIGH IN PROTEIN TENEUR ÉLEVÉE EN PROTÉINES
340 g

KAM.cf is a set of "additional/unofficial" rules developed to respond faster to spam, standard rules takes some days to be deployed due to masscheck. They are very effective but there could be "very few" false positives.

New features in SpamAssassin 4.0

SpamAssassin 4.0 and future releases

- full utf-8 support
- GeoDB module for a better geolocalization support
- better TxRep handling

Questions ?