The sys admin's daily grind: SA-Update

NEW ORDER

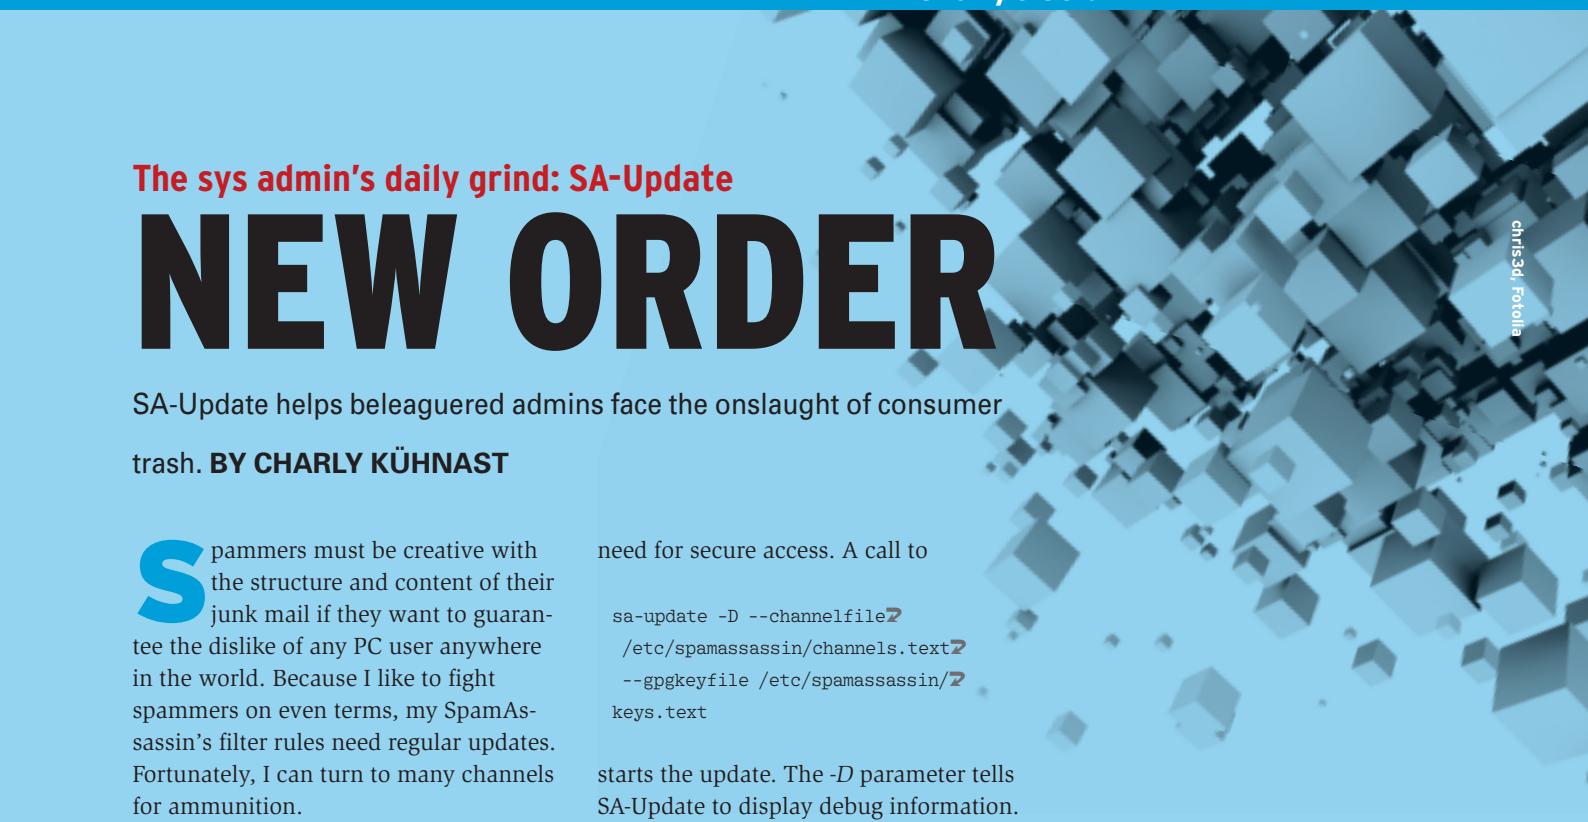

SA-Update helps beleaguered admins face the onslaught of consumer trash. **BY CHARLY KÜHNAST**

Spammers must be creative with the structure and content of their junk mail if they want to guarantee the dislike of any PC user anywhere in the world. Because I like to fight spammers on even terms, my SpamAssassin's filter rules need regular updates. Fortunately, I can turn to many channels for ammunition.

SA-Update

The tool that retrieves the updates and copies them to the right spot goes by the name of SA-Update [1].

A GPG key prevents various manipulation techniques such as DNS spoofing. To rejuvenate the default channel, *updates.spamassassin.org*, I first need the matching public key:

```
wget http://spamassassin.➥
apache.org/updates/GPG.KEY
gpg --import GPG.KEY
sa-update --import GPG.KEY
```

Create Files

Next, I create two files in the SpamAssassin folder. One of them, *channels.text*, lists the update channels. The second, *keys.text*, holds the GPG key IDs that I need for secure access. A call to

```
sa-update -D --channelfile➥
 /etc/spamassassin/channels.text➥
 --gpgkeyfile /etc/spamassassin/➥
keys.text
```

starts the update. The *-D* parameter tells SA-Update to display debug information. Without this parameter – SA-Update is as taciturn as Charles Bronson's character Harmonica [2] – there is no such thing as your average verbose mode.

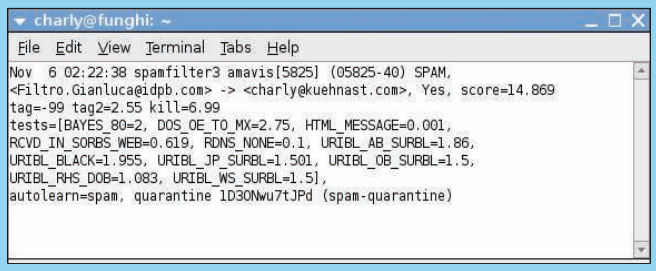

Figure 1: A quick inspection of the mail log reveals what SpamAssassin didn't like about any given message and which ruleset it used.

Filters, Please!

The return value gives an easy method of checking for a successful update. A return value of *0* means that SA-Update has added new filter rules. A value of *1* means that the ruleset was already up to date. A value of *4* or more indicates an error, and that means I need to check the debug output more closely.

To improve the spam detection rate, I like to add channels such as OpenProtect [3] or Daryl O'Shea [4]. A useful overview of the rules of the SpamAssassin Rules Emporium (SARE) are available online [5], and the default ruleset is explained in detail [6]. The filter rule short forms appear in the mail logs; thus, you can tell at a glance what SpamAssassin doesn't like about a message and which ruleset the tool used (Figure 1).

The most important question is, "Is it worthwhile?" Definitely! My spam filter's detection rates benefit considerably by extending the rulesets. Still, I like to keep an eye on the logfiles: The danger of false positives grows with each new filter rule you add. ■

THE AUTHOR

Charly Kühnast is a Unix operating system administrator at the Data Center in Moers, Germany. His tasks include firewall and DMZ security and availability. He divides his leisure time into hot, wet, and eastern sectors, where he enjoys cooking, fresh water aquariums, and learning Japanese, respectively.

INFO

[1] SA-Update: *http://wiki.apache.org/spamassassin/RuleUpdates*

[2] "Once Upon a Time in the West" ("C'era una Volta il West"), 1968, *http://www.imdb.com/title/tt0064116/*

[3] OpenProtect: *http://saupdates.openprotect.com*

[4] Daryl O'Shea: *http://daryl.dostech.ca/sa-update/sare/sare-sa-update-howto.txt*

[5] SARE: *http://www.rulesemporium.com/rules.htm*

[6] Default ruleset: *http://spamassassin.apache.org/tests_3_2_x.html*