Mirroring the Virus Database

Luca Gibelli

6th November 2003

Some guidelines for people interested in contributing to the distribution of ClamAV virus database.

1 Introduction

1.1 Who is responsible for the virus database

The virusdb team take care of reviewing virus signatures, checking for new viruses in the wild and committing changes to the virus database file.

The updates are released quite often (usually no less than three times a week). If you want to be notified whenever the virus database is updated subscribe to clamav-virusdb *at* lists.sourceforge.net .

Every time the virusdb team updates the database, the ChangeLog will be posted to the mailing-list.

Visit `http://lists.sourceforge.net/mailman/listinfo/clamav-virusdb` for the list description and archives.

If you need to contact the virusdb team please write to: virus-team *at* clamav.net

1.2 Virus submission

Whenever you find a new virus which is not detected by ClamAV you should send it to the virusdb team. They will review your submission and update the database so that the whole ClamAV user community can take benefit from it.

There are two ways to submit a virus sample:

- fill the form at `http://clamav.sourceforge.net/cgi-bin/sendvirus.cgi` (preferred method)
- send it by email to the following address: virus *at* clamav.net

Never send virus samples to ClamAV mailing-lists or developers addresses.

1.3 Getting a copy of the latest virus database

The most important factor for an antivirus's efficiency is to be up to date. ClamAV comes with a tool to update the virus database automatically: its name is *freshclam*.

freshclam reads a list of hostnames from *mirrors.txt* and tries to connect to them. If the first server in the list fails, freshclam will sleep for 10 secs and then try again with the next one, and so on.

If freshclam finds a new database, it downloads it and then sends a notify to clamd (if active) to reload the database.

It is important for the machine running ClamAV to be able to connect to port 80 of external hosts on Internet either directly or through a proxy. There are known problems with some transparent proxies caching what they shouldn't cache. If you should run into this kind of problem, please check your proxy configuration before reporting a bug.

2 Mirroring the database

2.1 The need for mirrors

To prevent the spread of worms it is essential to check for updates frequently. ClamAV users often configure freshclam with a check interval of 10 minutes.

With an exponentially growing number of ClamAV users, the servers hosting the virus database files get easily overloaded.

Without mirrors, the traffic on our main site reached 100GB/month. So if you are going to set up a new mirror, you can expect a traffic of about 100GB divided by the number of already existing mirrors+1.

freshclam downloads the database from `http://database.clamav.net/`.

database.clamav.net is a round robin record that tries to equally balance the traffic between all the database mirrors. The round robin record allows us to alter the mirrors list in real-time, thus if a mirror stops working or ceases to get updates it can be removed immediately from the list without any intervent on the user side.

2.2 Requirements to become a mirror

We need fast reliable mirrors. Servers eligible for becoming mirrors have to provide:

- At least a 10Mbit/s link to the Internet
- Unlimited traffic
- At least 50MB of web space
- Ability to let the virusdb team update the web space via scp/ftp

2.3 How to become a mirror

You have to follow these steps:

1. Set up a virtual host for `http://database.clamav.net/`. The database files will be placed at `http://database.clamav.net/database/`
   Additional server aliases for the above virtual host can be added too, but are not required. If you are using name based virtual hosts see `http://httpd.apache.org/docs/mod/core.html#serveralias` for more information.
   Here is an example for a typical setup:

   <VirtualHost 10.1.2.3>
   ServerAdmin webmaster@clamav.foo.com
   DocumentRoot /home/users/clamavdb/public_html
   ServerName clamav.foo.com
   ServerAlias database.clamav.net
   </VirtualHost>

2. Create an account with login "clamavdb" and give it write access to the virtual host's DocumentRoot.
   SCP is the preferred transferring method. You can find more information on creating a scponly account at:
   `http://www.sublimation.org/scponly/`
   FTP is also accepted.
   The virusdb team will use this account *only* to update the virus database.

3. Send the server details (account info, ip address, country, virtual host aliases, available bandwidth and sysadmin's email address) to nervoso *at* users.sourceforge.net

.
You can encrypt the message with this PGP public key:
`http://www.nervous.it/lucagibelli.asc`.

4. Subscribe to clamav-mirrors *at* lists.sourceforge.net: see
`http://lists.sourceforge.net/mailman/listinfo/clamav-mirrors`
for more info.
Subscribe requests have to be approved.

If your server meets the above requirements you'll receive a subscribe confirmation soon after and your server will be added to the round robin record.

2.4 Admin's duty

- Scheduled downtimes should be announced on the clamav-mirrors mailing-list in advance.
- IP address changes should be notified in advance too.
- It is essential to be able to contact the sysadmin responsible for the mirror server and get a quick response. Whenever a problem with a mirror occurs we need to immediately find out its cause and act consequently.

3 Mirror status

Every mirror is continously monitored to ensure that every ClamAV user gets the latest virus database.

Every six hours we attempt to upload a file called *timestamp* on every mirror. Every hour we choose a random mirror and check that *timestamp* is fresh. If the file is one day old or unavailable, the mirror if marked as "old" and ClamAV team receive a warning. If the situation persists for two days, the mirror is temporarily removed from the list.

You can view the current status of every ClamAV database mirror at `http://clamav.sourceforge.net/mirrors.html`.