Mirroring the Virus Database

Luca Gibelli

May 6, 2014

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

1 Introduction

1.1 This doc

The latest version of this document is always available at http://www.clamav.net/doc/mirrors/.
Before going any further, please check that you are reading the latest version.

Japanese sysadmins can find a translated version of this doc at http://www.orange.co.jp/~masaki/clamav/mirror-howto-jp.html (not necessarily up to date).

1.2 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.clamav.net .

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

Visit for the list description and archives.

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

1.3 Virus submission

Whenever you find a new virus which is not detected by ClamAV you should send it to the virusdb team by filling the form at http://www.clamav.net/sendvirus.html. They

will review your submission and update the database so that the whole ClamAV user community can take benefit from it.

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

1.4 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 looks up the TXT record associated with *current.cvd.clamav.net* and extracts the latest database version available from the string returned. If the local database is outdated, freshclam tries to connect to the hostnames listed in freshclam.conf (DatabaseMirror directive). If the first server in the list fails or the latest database is not available on that mirror (e.g. in case there has been a problem sync'ing the mirror), freshclam will sleep for 10 secs and then try again with the next one, and so on.

After freshclam downloads the new database, it sends a notify to clamd (if active) to reload the database.

It is important for the machine running ClamAV to be able to make DNS lookups and 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 30 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 was 100GB/month (May 2003).

On Feb 2004 the traffic on each mirror (11 in total) reached 120GB/month. Thanks to some improvements in freshclam and the increasing number of mirrors (currently 60), the traffic on each mirror was lowered to 40GB/month (Aug 2004). That makes about 2.5TByte/month of global traffic.

Our users are encouraged to add the following directives to their freshclam.conf :

DatabaseMirror db.XY.clamav.net

DatabaseMirror db.local.clamav.net

where XY stands for the country the server lives in [1]

Each db.XY.clamav.net record points to the mirrors available in that country[2] or, in case there are none, the continent.

If freshclam can't connect to db.XY.clamav.net, it will fail back on db.local.clamav.net, which **attempts** to redirect the user to the closest pool of mirrors by looking up its ip source address in GeoIP database (http://www.maxmind.com/app/geoip_country).[3] We are aware that looking up the ip source address is not an accurate method to find the user location from a network topology point of view. We accept the risk.

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[4]
- Unlimited traffic
- At least 50MB of web space
- Support for our *push-mirroring* system
- The mirror has to be available to all ClamAV users. We DO NOT support private mirrors.
- ssh 2 (read on)

We also appreciate (but do not require) having shell access to the server hosting the mirror. FTP access is no longer accepted.
The virusdb team will use the account *only* to update the virus database.

2.3 How to become a mirror

Before setting up a mirror contact *luca -at- clamav.net*!

You have to follow these steps:

[1] a full list is available at http://www.iana.org/cctld/cctld-whois.htm

[2] For a complete list of the mirrors available in each country visit http://www.clamav.net/mirrors.html

[3] See:
http://www.iana.org/assignments/ipv4-address-space
http://ip-to-country.webhosting.info/
http://ftp.apnic.net/stats/apnic/
http://www.ripe.net/db/erx/erx-ip/
Some of the information were contributed by Walter Hop (from transip.nl).

[4] Traffic is bursty, that's why we request such a large pipe

1. Set up a virtual host for
   http://database.clamav.net, http://db.*.clamav.net and http://clamav.your-domain.tld
   Note there is an asterisk in the second hostname. A literal asterisk.
   Do not replace it with your country code.
   If you are using name based virtual hosts[5] 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 john@clamav.foo.com
   DocumentRoot /home/users/clamavdb/public_html
   ServerName database.clamav.net
   ServerAlias db..clamav.net*
   ServerAlias clamav.foo.com
   </VirtualHost>
   If you are not using Apache and you cannot create wildcard vhosts, you must use IP based virtual hosts!
   Please note that an http redirect (e.g. RedirectPermanent) is not enough! freshclam can't handle redirects yet.
   If you are running Apache 2.x, please use the following directive for proper logging:
   LogFormat "%h %l %u %t \"%r\" %>s %O \"%{Referer}i\" \"%{User-Agent}i\"" combinedrealsize
   CustomLog /path/to/clamav-access.log combinedrealsize
   See the "Statistics" paragraph for more info.

2. Create an account with login "clamavdb" and give it write access to the virtual host's DocumentRoot.
   You may want to disable password authentication for this account and change the password to something obscure.
   The "clamavdb" user's shell must be /bin/sh or /bin/bash . Otherwise the user won't be able to run the command associated with the ssh public key[6].

[5] You can check whether the mirror setup is correct or not, simply by adding a line like this:
your-server-ip database.clamav.net
to the /etc/hosts on your client machine. Then visit http://database.clamav.net and see if you can download files from your mirror's directory.

[6] Take a look at the content of *"authorized_keys_noshell"*: the only command which can be executed by the owner of the corresponding ssh private key is ~/bin/clam-clientsync. We will only be able to trigger the execution of that script and nothing else!

3. Download the following files:
   clam-clientsync.conf
   clam-clientsync
   authorized_keys_shell
   authorized_keys_noshell
   authorized_keys_shell.sig
   authorized_keys_noshell.sig
   from http://www.clamav.net/doc/mirrors/

4. Verify the signature using:
   $ gpg –verify authorized_keys_noshell.sig authorized_keys_noshell
   $ gpg –verify authorized_keys_shell.sig authorized_keys_shell
   My PGP public key is available on most keyservers and on ClamAV web site. It can eventually be verified by telephone. Contact me by email first.

5. If you don't want to give us shell access, copy *authorized_keys_noshell* to *~clamavdb/.ssh/authorized_keys*:
   $ cp authorized_keys_noshell ~/.ssh/authorized_keys
   If you want to give us shell access, use *authorized_keys_shell* instead:
   $ cp authorized_keys_shell ~clamavdb/.ssh/authorized_keys
   $ chmod go-w ~clamavdb
   $ chmod 700 ~clamavdb/.ssh
   $ chmod 600 ~clamavdb/.ssh/authorized_keys

6. Copy clam-clientsync to ~clamavdb/bin/
   Copy clam-clientsync.conf to ~clamavdb/etc/
   chmod 600 ~clamavdb/etc/clam-clientsync.conf
   chmod 755 ~clamavdb/bin/clam-clientsync
   Everything must be owned by user clamavdb.
   The clam-clientsync requires the "lockfile" program, which is part of the *procmail* package. Before going any further, please check that "lockfile" is available.

7. Send the server's details (ip address, country, virtual host aliases, available bandwidth and sysadmin's full name and email address) to *luca at clamav.net* .

8. Edit ~clamavdb/etc/clam-clientsync.conf . If your DocumentRoot (see paragraph 1) is */home/users/clamavdb/public_html* , your login is *foo* and your password *guessme*, then your clam-clientsync.conf will look like this:

However, shell access is really appreciated. If you are willingly to give us shell access, use *authorized_keys_shell* instead which contains Luca Gibelli and Tomasz Papszun ssh public keys too.

TO=/home/users/clamavdb/public_html
RSYNC_USER=foo
RSYNC_PASSWORD=guessme
EXCLUDE="–exclude local_*"

9. Reconfigure your packet filter to allow incoming connections on port 22/tcp and outgoing connections to ports 873/tcp and 873/udp.
   You can furtherly restrict access to these ports by only allowing connections from/to the following IP addresses:
   194.109.142.194, 64.18.103.6, 194.242.226.43 .
   rsync.clamav.net is a round robin record which points to our master mirror servers.
   Any changes to this record will be announced on the clamav-mirrors mailing-list.

10. You are welcome to put your company logo on the mirror home page. Just copy it to the DocumentRoot and rename it to "local_logo.png". The index.html is unique for every mirror. Please note that any file in the DocumentRoot whose name doesn't match "local_*" will be deleted at every mirror sync.

11. Subscribe to clamav-mirrors *at* lists.clamav.net: see
    http://lists.clamav.net/mailman/listinfo/clamav-mirrors for more info.
    Subscribe requests have to be approved. We will approve your subscription request only *after* reviewing your server's info.

When everything is done, your server's IP address will be added either to your country's dns record (db.XY.clamav.net) or one of the round robin record (db.<continent>.clamav.net) and your company will be listed on our mirrors list page.

2.4 Statistics

Although it's not required, we really appreciate if you can make access statistics of your mirror available to us. They should be available at http://your-mirror-host-name/local_stats/ and they **must** be protected with login and password. You should use the same login and password you are using in your ~clamavdb/etc/clam-clientsync.conf file.

If possible, please tell your statistics generator to ignore requests made by the "ClamAV-MirrorCheck" agent.

If you are using Webalizer, you can add the following directive to your conf. file:
HideAgent ClamAV-MirrorCheck
If you are using AWStats, you can add this one instead:
SkipUserAgents="ClamAV-MirrorCheck"

Refer to your stats generator's manual for more info.

Important note for Apache2 users:

As stated in the Apache documentation from http://httpd.apache.org/docs/2.0/mod/mod_log_config.html:

Note that in httpd 2.0, unlike 1.3, the %b and %B format strings do not represent the number of bytes sent to the client, but simply the size in bytes of the HTTP response (which will differ, for instance, if the connection is aborted, or if SSL is used). The %O format provided by mod_logio will log the actual number of bytes sent over the network.

2.5 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.
- Changes in the ssh host public key of the mirror host should be announced on the clamav-mirrors mailing-list.
- 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 Notes for sigmakers

New sigmakers should send their ssh2 public key to *luca at clamav.net* . Their public key will be added to rsyncX.clamav.net authorized_keys (after being verified).

Sigmakers can upload a new database to either rsync1.clamav.net or rsync2.clamav.net using a (scp|sftp|rsync)-only account.

The new database won't be available to other people immediately. First, sigmakers have to notify the rsyncX.clamav.net server that a new database is available.

Here is the step-by-step procedure to release a new database version and propagate it around the world:

1. Assume your ssh private key is ~/.ssh/id_rsa and you've just built a new daily.cvd. Assume you want to use rsync1.clamav.net
2. In order to upload the new database, you have to run:
   $ rsync -tcz –stats –progress -e 'ssh -i ~/.ssh/id_rsa' daily.cvd clamupload@rsync1.clamav.net:public_html/

3. Next, you need to notify rsync1.clamav.net that a new database is available:
   $ ssh rsync1.clamav.net -i ~/.ssh/id_rsa -l clamavdb sleep 1

4. rsync1.clamav.net will verify the digital signature of the newly uploaded database using *sigtool -i*. If it finds an error, it will refuse to distribute the database to other mirrors.

5. rsync1.clamav.net will copy the previously uploaded database to its rsync shared directory.

6. rsync1.clamav.net will notify every mirror that a new database is available

7. Every mirror will rsync its copy of the database from *rsync1.clamav.net::clamavdb* (only mirrors can access the rsync server at rsync1.clamav.net, it's password protected)

As a fallback, every three hours, either rsync1.clamav.net or rsync2.clamav.net force an update on every mirror.

If rsync1 can't reach rsync2 or viceversa, the automatic update doesn't take place. This is done to avoid propagating an old database.

To avoid conflicts, sigmakers should use rsync1 by default and if it fails, switch to rsync2. Whenever a sigmaker uses rsync2, he should announce it on the clamav-team mailing-list so that every other sigmaker uses rsync2 too, until the issues with rsync1 are over.

4 Mirror status

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

Every three hours we 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 the 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://www.clamav.net/mirrors.html .

Please note that this page doesn't reflect how *often* the database is propagated to mirrors. It just shows the trend of mirrors availability.