Table of Contents

Installing frontier-squid from source tarball

We recommend installing the frontier-squid package from the frontier-squid rpm, but if you want to use the source tarball, which is more portable, easily relocatable and can be run under any username, the instructions are below.

See the frontier-squid tarball release notes for information on what has changed in each release. Here is what is on this page:

Download and install software

First, create an account with username dbfrontier on your machine. (If for some reason, you can't use the username dbfrontier, any username will work.) Then as user dbfrontier (NOT root) download a tarball into this account.

Unpack and configure the current version of the tarball:

```
$ wget http://frontier.cern.ch/dist/frontier-squid-4.8-2.tar.gz
$ tar -xvzf frontier-squid-4.8-2.tar.gz
$ cd frontier-squid-4.8-2
$ ./configure
```

On prompt enter the directory name where the Squid will be installed. This directory holds the working software, cache, and logs so there should be plenty of space available (unless you relocate the cache and logs as described below). This directory is called "/install_dir" below. It should be on a local disk of the computer you are using and not NFS or AFS mounted. Note that the directory you enter should be an absolute (fully qualified) directory name and not a relative one. You may either re-use a previous install directory or create a new one for each release. Creating a new directory for each release makes it easier to back out to a previous release and ensures a clean installation, but it requires a little extra work to set up (described below).

If the directory you are installing into does not yet contain customize.sh, you will also be prompted for the old installation path. If customize.sh is found in the old installation path, it will be copied into the source directory and the configure step will be finished. If you like, you can avoid the first two questions by passing "--prefix=/install_dir" and "--oldprefix=/oldinstall_dir" parameters to ./configure. If you have not previously installed a release that supports customize.sh, you will be asked a few additional questions about basic configuration parameters.

On prompt enter network/netmask which is allowed to access the Squid.

Examples: 131.154.184.0/255.255.255.0 or 131.154.0.0/255.255.0.0

The script does allow to specify many subnets - just separate them by a blank. Include IPv6 address ranges if you plan to support IPv6 addresses. If you just hit enter, the standard private network addresses 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16, fc00::/7, and fe80::/10 will be allowed. NOTE: The default behavior is to allow any IP address you specify here to use the squid to cache objects coming from any destination address. If you would like a more restrictive policy or other options please see the rpm instructions for restricting the destination.

On prompt enter the amount of cache memory (in MB) the squid should use. This should be at most 1/8 of your hardware memory. 128 MB should be fine, leaving a lot of memory for disk buffering by the OS, because squid performs better for large objects in disk cache buffers than in its own internal memory cache.

On prompt enter the amount of disk space (in MB) the squid should use for a cache. One suggestion is to set this size at 70% of the available space in your disk partition to allow room for the executables, log files, etc. It should be at least 20000, but there's no need to make it more than 100000.

You can double check your responses to the prompts by reading Makefile.conf.inc and edit them there before running **make** if you wish.

On RHEL6-based Linux, a newer version of the compiler is required for squid-4. Before doing any building, do the following:

```
$ sudo yum install devtoolset-2-toolchain scl-utils
$ scl enable devtoolset-2 bash
```

That will put you into a shell where the new compiler is first in the PATH.

Then do:

```
$ make
$ make install
```

After that you should examine `/install_dir/frontier-cache/squid/etc/customize.sh` and make any changes or other customizations you want to. Comments in the default installation of customize.sh give more details on what can be done with it.

Manual control of the server

To do a manual start/stop of the server (as user dbfrontier):

```
$ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh start
```

You can also stop it if you need to:

```
$ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh stop
```

Remember to start your server after you have installed it.

Setup for different install directories with each release

If you choose to use a different install directory for each release, do the following extra things:

- Create a symbolic link at a place you will re-use for each new installation, and use that for the cron job described in the next section and in /etc/init.d/frontier-squid.sh described in the following two sections, so those don't need to be reinstalled for every release.
- Either remember to clean out the old installation's disk cache (in `/install_dir/frontier-cache/squid/var/cache`) and logs (in `/install_dir/frontier-cache/squid/var/logs`) each time or (better) edit customize.sh to set the `cache_log`, `pid_filename`, and `coredump_dir` options and the second parameter of the `cache_dir` option and the first parameter of the `access_log` option to use common directories that you re-use for each new installation. This has an added advantage of not requiring a lot of disk space where you install the software but rather where you choose to put the cache and logs. For example:

```
setoptionparameter("cache_dir", 2, "/data/squid_cache")
setoptionparameter("access_log", 1, "/data/squid_logs/access.log")
setoption("cache_log", "/data/squid_logs/cache.log")
setoption("pid_filename", "/data/squid_logs/squid.pid")
setoption("coredump_dir", "/data/squid_cache")
```

Set up cron jobs

As user dbfrontier, set up cron jobs to rotate the logs, with crontab entries like this:

```
7 7 * * * /install_dir/frontier-cache/utils/cron/daily.sh </dev/null 2<&1
8,23,38,53 * * * * /install_dir/frontier-cache/utils/cron/hourly.sh </dev/null 2<&1
```

You could get the above crontab by doing (with the appropriate value of install_dir)

```
$ crontab /install_dir/frontier-cache/utils/cron/crontab.dat
```

You can change the hour and minutes as you like, but leave the first hourly.sh to be one minute after daily.sh, and avoid multiples of 5 for the minute because it can interfere with the monitoring probes which happen every 5 minutes. The hourly.sh script will rotate the logs if access.log goes over a given size, default 5GB when compression is used (which is the default) or 1GB when compression is not used. You can change that value by setting the environment variable SQUID_MAX_ACCESS_LOG to a different number of bytes. You can also append an M for megabytes or G for gigabytes. For example for 20GB you can use:

```
8,23,38,53 * * * * SQUID_MAX_ACCESS_LOG=20G /install_dir/frontier-cache/utils/cron/hourly.sh
```

To help with choosing the max log size, see the corresponding rpm instructions, except that if you choose to disable compression you can set `SQUID_COMPRESS_LOGS=false` in the cron command line. If disk space for the logs is a concern see the section on the Access log growth issue below.

As root, set up start at boot time

This is the only step to be done as root.

```
# cp /install_dir/frontier-cache/utils/init.d/frontier-squid.sh /etc/init.d
# /sbin/chkconfig --add frontier-squid.sh
```

Testing your installation

Do the same test as in rpm testing instructions.

Monitoring

Enable monitoring the same way as in the rpm monitoring instructions.

Some Useful Commands

`$ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh` with any parameter or no parameter will recreate squid.conf after changing customize.sh

`$ /install_dir/frontier-cache/squid/sbin/squid -k parse` will just read squid.conf to see if it makes sense

`$ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh reload` sends a HUP signal and has squid reread squid.conf

`$ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh status` checks if squid is running

`$ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh restart` stops squid and starts squid without clearing the cache

`$ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh cleancache` deletes and recreates the cache, like a start does, but without starting squid

`$ /install_dir/frontier-cache/squid/bin/squidclient mgr:info` outputs operational information about your squid

Access log growth issue

With many active clients, it is still possible for the squid access.log to grow to unmanageable size. The squid will crash if it runs out of available diskspace. There are a couple ways to avoid this problem:

1) Make sure that you have the hourly.sh cron job enabled as described in the Set up cron jobs section above to rotate the log when it grows over a size you choose.

2) The other possibility is to disable writing to access.log by putting the following in `/install_dir/frontier-cache/squid/etc/customize.sh`:

```
    setoption("access_log", "none")
```

and then do

```
    $ /install_dir/frontier-cache/utils/bin/fn-local-squid.sh reload
```

to update squid.conf and load it if the squid is already running, otherwise just use `start` instead of `reload`.

The squid installation script has the access log turned on by default. It is recommended that a new installation be installed with it on, the functioning of the squid verified by reading the access log, then if disk space is limited, turn the access log off when the squid is in production. Even if you do turn the access log off, you should still run the daily.sh script once per day to rotate the other logs.

Filedescriptors

At some installations with a very large number of worker nodes it may be possible to see error messages about running out of filedescriptors in your cache.log. It is easy to avoid this problem:

1) As root, add the following line to `/etc/security/limits.conf`

```
    * - nofile 16384
```

2) Reboot the machine.

You can check your file descriptor limit and usage by doing:

```
    $ /install_dir/frontier-cache/squid/bin/squidclient mgr:info
```

Other issues

See also the rpm common issues which also apply to the tarball.

Responsible: DaveDykstra

This topic: Frontier > InstallSquidTarball
Topic revision: r42 - 2019-08-29 - DaveDykstra