Cluster synchronization with Csync²

Clifford Wolf, http://www.clifford.at/

May 28, 2006

1 Introduction

Csync² [1] is a tool for asynchronous file synchronization in clusters. Asynchronous file synchronization is good for files which are seldom modified - such as configuration files or application images - but it is not adequate for some other types of data.

For instance a database with continuous write accesses should be synced synchronously in order to ensure the data integrity. But that does not automatically mean that synchronous synchronization is better; it simply is different and there are many cases where asynchronous synchronization is favored over synchronous synchronization. Some pros of asynchronous synchronization are:

1. Most asynchronous synchronization tools (including Csync²) are implemented as single-shot commands which need to be executed each time in order to run one synchronization cycle. Therefore it is possible to test changes on one host before deploying them on the others (and also return to the old state if the changes turn out to be bogus).

2. The synchronization algorithms are much simpler and thus less error-prone.

3. Asynchronous synchronization tools can be (and usually are) implemented as normal user mode programs. Synchronous synchronization tools need to be implemented as operating system extensions. Therefore asynchronous tools are easier to deploy and more portable.

4. It is much easier to build systems which allow setups with many hosts and complex replication rules.

But most asynchronous synchronization tools are pretty primitive and do not even cover a small portion of the issues found in real world environments.

I have developed Csync² because I found none of the existing tools for asynchronous synchronization satisfying. The development of Csync² has been sponsored by LINBIT Information Technologies [2], the company which also sponsors the synchronous block device synchronization toolchain DRBD [3].

Note: I will simply use the term *synchronization* instead of the semi-oxymoron *asynchronous synchronization* in the rest of this paper.

1.1 Csync² features

Most synchronization tools are very simple wrappers for remote-copy tools such as `rsync` or `scp`. These solutions work well in most cases but still leave a big gap for more sophisticated tools such as Csync². The most important features of Csync² are described in the following sections.

1.1.1 Conflict detection

Most of the trivial synchronization tools just copy the newer file over the older one. This can be a very dangerous behavior if the same file has been changed on more than one host. Csync² detects such a situation as a conflict and will not synchronize the file. Those conflicts then need to be resolved manually by the cluster administrator.

It is not considered as a conflict by Csync² when the same change has been performed on two hosts (e.g. because it has already been synchronized with another tool).

It is also possible to let Csync² resolve conflicts automatically for some or all files using one of the pre-defined auto-resolve methods. The available methods are: `none` (the default behavior), `first` (the host on which Csync² is executed first wins), `younger` and `older` (the younger or older file wins), `bigger` and `smaller` (the bigger or smaller file wins), `left` and `right` (the host on the left side or the right side in the host list wins).

The `younger`, `older`, `bigger` and `smaller` methods let the remote side win the conflict if the file

has been removed on the local side.

1.1.2 Replicating file removals

Many synchronization tools can not synchronize file removals because they can not distinguish between the file being removed on one host and being created on the other one. So instead of removing the file on the second host they recreate it on the first one.

$\textsc{Csync}^2$ detects file removals as such and can synchronize them correctly.

1.1.3 Complex setups

Many synchronization tools are strictly designed for two-host-setups. This is an inadequate restriction and so $\textsc{Csync}^2$ can handle any number of hosts.

$\textsc{Csync}^2$ can even handle complex setups where e.g. all hosts in a cluster share the `/etc/hosts` file, but one `/etc/passwd` file is only shared among the members of a small sub-group of hosts and another `/etc/passwd` file is shared among the other hosts in the cluster.

1.1.4 Reacting to updates

In many cases it is not enough to simply synchronize a file between cluster nodes. It also is important to tell the applications using the synchronized file that the underlying file has been changed, e.g. by restarting the application.

$\textsc{Csync}^2$ can be configured to execute arbitrary commands when files matching an arbitrary set of shell patterns are synchronized.

2 The $\textsc{Csync}^2$ algorithm

Many other synchronization tools compare the hosts, try to figure out which host is the most up-to-date one and then synchronize the state from this host to all other hosts. This algorithm can not detect conflicts, can not distinguish between file removals and file creations and therfore it is not used in $\textsc{Csync}^2$.

$\textsc{Csync}^2$ creates a little database with filesystem metadata on each host. This database (`/var/lib/csync2/`*hostname*`.db`) contains a list of the local files under the control of $\textsc{Csync}^2$. The

database also contains information such as the file modification timestamps and file sizes.

This database is used by $\textsc{Csync}^2$ to detect changes by comparison with the local filesystem. The synchronization itself is then performed using the $\textsc{Csync}^2$ protocol (TCP port 30865).

Note that this approach implies that $\textsc{Csync}^2$ can only push changes from the machine on which the changes has been performed to the other machines in the cluster. Running $\textsc{Csync}^2$ on any other machine in the cluster can not detect and so can not synchronize the changes.

Librsync [4] is used for bandwidth-saving file synchronization and SSL is used for encrypting the network traffic. The sqlite library [5] (version 2) is used for managing the $\textsc{Csync}^2$ database files. Authentication is performed using auto-generated pre-shared-keys in combination with the peer IP address and the peer SSL certificate.

3 Setting up $\textsc{Csync}^2$

3.1 Building $\textsc{Csync}^2$ from source

Simply download the latest $\textsc{Csync}^2$ source tar.gz from `http://oss.linbit.com/csync2/`, extract it and run the usual `./configure - make - make install` trio.

$\textsc{Csync}^2$ has a few prerequisites in addition to a C compiler, the standard system libraries and headers and the usual gnu toolchain (`make`, etc):

1. You need librsync, libsqlite (version 2) and libssl installed (including development headers).

2. Bison and flex are needed to build the configuration file parser.

3.2 $\textsc{Csync}^2$ in Linux distributions

As of this writing there are no official Debian, Red-Hat or SuSE packages for $\textsc{Csync}^2$. Gentoo has a $\textsc{Csync}^2$ package, but is has not been updated for a year now. As far as I know, ROCK Linux [6] is the only system with an up-to-date $\textsc{Csync}^2$ package. So I recommend that all users of non-ROCK distributions built the package from source.

The $\textsc{Csync}^2$ source package contains an RPM `.specs` file as well as a `debian/` directory. So it is possible to use `rpmbuild` or `debuild` to build $\textsc{Csync}^2$.

3.3 Post installation

Next you need to create an SSL certificate for the local CSYNC[2] server. Simply running `make cert` in the CSYNC[2] source directory will create and install a self-signed SSL certificate for you. Alternatively, if you have no source around, run the following commands:

```
openssl genrsa \
    -out /etc/csync2_ssl_key.pem 1024
openssl req -new \
    -key /etc/csync2_ssl_key.pem \
    -out /etc/csync2_ssl_cert.csr
openssl x509 -req -days 600 \
    -in /etc/csync2_ssl_cert.csr \
    -signkey /etc/csync2_ssl_key.pem \
    -out /etc/csync2_ssl_cert.pem
```

You have to do that on each host you're running csync2 on. When servers are talking with each other for the first time, they add each other to the database.

The CSYNC[2] TCP port 30865 needs to be added to the `/etc/services` file and inetd needs to be told about CSYNC[2] by adding

```
csync2 stream tcp nowait root \
        /usr/local/sbin/csync2 csync2 -i
```

to `/etc/inetd.conf`.

3.4 Configuration File

Figure 1 shows a simple CSYNC[2] configuration file. The configuration filename is `/etc/csync2.cfg` when no `-C` *configname* option has been passed and `/etc/csync2_configname.cfg` with a `-C` *configname* option.

3.4.1 Synchronization Groups

In the example configuration file you will find the declaration of a synchronization group called `mygroup`. A CSYNC[2] setup can have any number of synchronization groups. Each group has its own list of member hosts and include/exclude rules.

CSYNC[2] automatically ignores all groups which do not contain the local hostname in the host list. This way you can use one big CSYNC[2] configuration file for the entire cluster.

3.4.2 Host Lists

Host lists are specified using the `host` keyword. You can eighter specify the hosts in a whitespace seperated list or use an extra `host` statement for each host.

The hostnames used here must be the local hostnames of the cluster nodes. That means you must use exactly the same string as printed out by the `hostname` command. Otherwise csync2 would be unable to associate the hostnames in the configuration file with the cluster nodes.

The `-N` *hostname* command line option can be used to set the local hostname used by CSYNC[2] to a different value than the one provided by the `hostname` command. This may be e.g. useful for environments where the local hostnames are automatically set by a DHCP server and because of that change often.

Sometimes it is desired that a host is receiving CSYNC[2] connections on an IP address which is not the IP address its DNS entry resolves to, e.g. when a crossover cable is used to directly connect the hosts or an extra synchronization network should be used. In this cases the syntax *hostname@interfacename* has to be used for the `host` records (see `host4` in the example config file).

Sometimes a host shall only receive updates from other hosts in the synchronization group but shall not be allowed to send updates to the other hosts. Such hosts (so-called *slave hosts*) must be specified in brackets, such as `host3` in the example config file.

3.4.3 Pre-Shared-Keys

Authentication is performed using the IP addresses and pre-shared-keys in CSYNC[2]. Each synchronization group in the config file must have exactly one `key` record specifying the file containing the pre-shared-key for this group. It is recommended to use a separate key for each synchronization group and only place a key file on those hosts which actually are members in the corresponding synchronization group.

The key files can be generated with `csync2 -k` *filename*.

```
group mygroup                          # A synchronization group (see 3.4.1)
{
        host host1 host2 (host3);       # host list (see 3.4.2)
        host host4@host4-eth2;

        key /etc/csync2.key_mygroup;    # pre-shared-key (see 3.4.3)

        include /etc/apache;            # include/exclude patterns (see 3.4.4)
        include %homedir%/bob;
        exclude %homedir%/bob/temp;
        exclude *~ .*;

        action                          # an action section (see 3.4.5)
        {
                pattern /etc/apache/httpd.conf;
                pattern /etc/apache/sites-available/*;
                exec "/usr/sbin/apache2ctl graceful";
                logfile "/var/log/csync2_action.log";
                do-local;
        }

        backup-directory /var/backups/csync2;
        backup-generations 3;           # backup old files (see 3.4.11)

        auto none;                      # auto resolving mode (see 3.4.6)
}

prefix homedir                          # a prefix declaration (see 3.4.7)
{
        on host[12]: /export/users;
        on *:        /home;
}
```

Figure 1: Example CSYNC^2 configuration file

3.4.4 Include/Exclude Patterns

The `include` and `exclude` patterns are used to specify which files should be synced in the synchronization group.

There are two kinds of patterns: pathname patterns which start with a slash character (or a prefix such as the `%homedir%` in the example; prefixes are explained in a later section) and basename patterns which do not.

The last matching pattern for each of both categories is chosen. If both categories match, the file will be synchronized.

The pathname patterns are matched against the beginning of the filename. So they must either match the full absolute filename or must match a directory in the path to the file. The file will not be synchronized if no matching `include` or `exclude` pathname pattern is found (i.e. the default pathname pattern is an exclude pattern).

The basename patterns are matched against the base filename without the path. So they can e.g. be used to include or exclude files by their filename extensions. The default basename pattern is an include pattern.

In our example config file that means that all files from `/etc/apache` and `%homedir%/bob` are

```
csync2 -cr /
if csync2 -M; then
        echo "!!"
        echo "!! There are unsynced changes! Type 'yes' if you still want to"
        echo "!! exit (or press crtl-c) and anything else if you want to start"
        echo "!! a new login shell instead."
        echo "!!"
        if read -p "Do you really want to logout? " in &&
           [ ".$in" != ".yes" ]; then
                exec bash --login
        fi
fi
```

Figure 2: The `csync2_locheck.sh` script

synced, except the dot files, files with a tilde character at the end of the filename, and files from `%homedir%/bob/temp`.

3.4.5 Actions

Each synchronization group may have any number of `action` sections. These `action` sections are used to specify shell commands which should be executed after a file is synchronized that matches any of the specified patterns.

The `exec` statement is used to specify the command which should be executed. Note that if multiple files matching the pattern are synced in one run, this command will only be executed once. The special token `%%` in the command string is substituted with the list of files which triggered the command execution.

The output of the command is appended to the specified logfile, or to `/dev/null` if the `logfile` statement is omitted.

Usually the action is only triggered on the targed hosts, not on the host on which the file modification has been detected in the first place. The `do-local` statement can be used to change this behavior and let Csync[2] also execute the command on the host from which the modification originated.

3.4.6 Conflict Auto-resolving

The `auto` statement is used to specify the conflict auto-resolving mechanism for this synchronization group. The default value is `auto none`.

See section 1.1.1 for a list of possible values for this setting.

3.4.7 Prefix Declarations

Prefixes (such as the `%homedir%` prefix in the example configuration file) can be used to synchronize directories which are named differently on the cluster nodes. In the example configuration file the directory for the user home directories is `/export/users` on the hosts `host1` and `host2` and `/home` on the other hosts.

The prefix value must be an absolute path name and must not contain any wildcard characters.

3.4.8 The `nossl` statement

Usually all Csync[2] network communication is encrypted using SSL. This can be changed with the `nossl` statement. This statement may only occur in the root context (not in a `group` or `prefix` section) and has two parameters. The first one is a shell pattern matching the source DNS name for the TCP connection and the second one is a shell pattern matching the destination DNS name.

So if e.g. a secure synchronization network is used between some hosts and all the interface DNS names end with `-sync`, a simple

```
nossl *-sync *-sync;
```

will disable the encryption overhead on the synchronization network. All other traffic will stay SSL encrypted.

3.4.9 The `config` statement

The `config` statement is nothing more then an include statement and can be used to include other config files. This can be used to modularize the configuration file.

3.4.10 The `ignore` statement

The `ignore` statement can be used to tell Csync[2] to not check and not sync the file user-id, the file group-id and/or the file permissions. The statement is only valid in the root context and accepts the parameters `uid`, `gid` and `mod` to turn off handling of user-ids, group-ids and file permissions.

3.4.11 Backing up

Csync[2] can back up the files it modifies. This may be useful for scenarios where one is afraid of accidentally syncing files in the wrong direction.

The `backup-directory` statement is used to tell Csync[2] in which directory it should create the backup files and the `backup-generations` statement is used to tell Csync[2] how many old versions of the files should be kept in the backup directory.

The files in the backup directory are named like the file they back up, with all slashes substituted by underscores and a generation counter appended. Note that only the file content, not the metadata such as ownership and permissions are backed up.

Per default Csync[2] does not back up the files it modifies. The default value for `backup-generations` is 3.

3.5 Activating the Logout Check

The Csync[2] sources contain a little script called `csync2_locheck.sh` (Figure 2).

If you copy that script into your `~/.bash_logout` script (or include it using the `source` shell command), the shell will not let you log out if there are any unsynced changes.

4 Database Schema

Figure 3 shows the Csync[2] database schema. The database can be accessed using the `sqlite` command line shell. All string values are URL encoded in the database.

The `file` table contains a list of all local files under Csync[2] control, the `checktxt` attribute contains a special string with information about file type, size, modification time and more. It looks like this:

```
v1:mtime=1103471832:mode=33152:
uid=1001:gid=111:type=reg:size=301
```

This `checktxt` attribute is used to check if a file has been changed on the local host.

If a local change has been detected, the entry in the `file` table is updated and entries in the `dirty` table are created for all peer hosts which should be updated. This way the information that a host should be updated does not get lost, even if the host in question is unreachable right now. The `force` attribute is set to `0` by default and to `1` when the cluster administrator marks one side as the right one in a synchronization conflict.

The `hint` table is usually not used. In large setups this table can be filled by a daemon listening on the inotify API. It is possible to tell Csync[2] to not check all files it is responsible for but only those which have entries in the `hint` table. However, the Linux syscall API is so fast that this only makes sense for really huge setups.

The `action` table is used for scheduling actions. Usually this table is empty after Csync[2] has been terminated. However, it is possible that Csync[2] gets interrupted in the middle of the synchronization process. In this case the records in the `action` table are processed when Csync[2] is executed the next time.

The `x509_cert` table is used to cache the SSL cetrificates used by the other hosts in the csync2 cluster (like the SSH `known_hosts` file).

5 Running Csync[2]

Simply calling `csync2` without any additional arguments prints out a help message (Figure 4). A more detailed description of the most important usage scenarios is given in the next sections.

5.1 Just synchronizing the files

The command `csync2 -x` (or `csync2 -xv`) checks for local changes and tries to synchronize them to the other hosts. The option `-d` (dry-run) can be

```
CREATE TABLE file (
        filename, checktxt,
        UNIQUE ( filename ) ON CONFLICT REPLACE
);

CREATE TABLE dirty (
        filename, force, myname, peername,
        UNIQUE ( filename, peername ) ON CONFLICT IGNORE
);

CREATE TABLE hint (
        filename, recursive,
        UNIQUE ( filename, recursive ) ON CONFLICT IGNORE
);

CREATE TABLE action (
        filename, command, logfile,
        UNIQUE ( filename, command ) ON CONFLICT IGNORE
);

CREATE TABLE x509_cert (
        peername, certdata,
        UNIQUE ( peername ) ON CONFLICT IGNORE
);
```

Figure 3: The CSYNC[2] database schema

used to do everything but the actual synchronization.

When you start CSYNC[2] the first time it compares its empty database with the filesystem and sees that all files just have been created. It then will try to synchronize the files. If the file is not present on the remote hosts it will simply be copied to the other host. There also is no problem if the file is already present on the remote host and has the same content. But if the file already exists on the remote host and has a different content, you have your first conflict.

5.2 Resolving a conflict

When two or more hosts in a CSYNC[2] synchronization group have detected changes for the same file we run into a conflict: CSYNC[2] can not know which version is the right one (unless an auto-resolving method has been specified in the configuration file). The cluster administrator needs to tell CSYNC[2] which version is the correct one. This can be done

using CSYNC[2] -f, e.g.:

```
# csync2 -x
While syncing file /etc/hosts:
ERROR from peer apollo:
    File is also marked dirty here!
Finished with 1 errors.

# csync2 -f /etc/hosts
# csync2 -xv
Connecting to host apollo (PLAIN) ...
Updating /etc/hosts on apollo ...
Finished with 0 errors.
```

5.3 Checking without syncing

It is also possible to just check the local filesystem without doing any connections to remote hosts: csync2 -cr / (the -r modifier tells CSYNC[2] to do a recursive check).

csync2 -c without any additional parameters checks all files listed in the hints table.

```
csync2 1.26 - cluster synchronization tool, 2nd generation
LINBIT Information Technologies GmbH <http://www.linbit.com>
Copyright (C) 2004, 2005  Clifford Wolf <clifford@clifford.at>
This program is free software under the terms of the GNU GPL.

Usage: csync2 [-v..] [-C config-name] \
              [-D database-dir] [-N hostname] [-p port] ..

With file parameters:
        -h [-r] file..          Add (recursive) hints for check to db
        -c [-r] file..          Check files and maybe add to dirty db
        -u [-d] [-r] file..     Updates files if listed in dirty db
        -f file..               Force this file in sync (resolve conflict)
        -m file..               Mark files in database as dirty

Simple mode:
        -x [-d] [[-r] file..]   Run checks for all given files and update
                                remote hosts.

Without file parameters:
        -c      Check all hints in db and eventually mark files as dirty
        -u [-d] Update (transfer dirty files to peers and mark as clear)

        -H      List all pending hints from status db
        -L      List all file-entries from status db
        -M      List all dirty files from status db

        -S myname peername      List file-entries from status db for this
                                synchronization pair.

        -T              Test if everything is in sync with all peers.

        -T filename     Test if this file is in sync with all peers.

        -T myname peername      Test if this synchronization pair is in sync.

        -T myname peer file     Test only this file in this sync pair.

        -TT     As -T, but print the unified diffs.

        The modes -H, -L, -M and -S return 2 if the requested db is empty.
        The mode -T returns 2 if both hosts are in sync.

        -i      Run in inetd server mode.
        -ii     Run in stand-alone server mode.
        -iii    Run in stand-alone server mode (one connect only).

        -R      Remove files from database which do not match config entries.

Modifiers:
        -r      Recursive operation over subdirectories
        -d      Dry-run on all remote update operations

        -B      Do not block everything into big SQL transactions. This
                slows down csync2 but allows multiple csync2 processes to
                access the database at the same time. Use e.g. when slow
                lines are used or huge files are transferred.

        -A      Open database in asynchronous mode. This will cause data
                corruption if the operating system crashes or the computer
                loses power.

        -I      Init-run. Use with care and read the documentation first!
                You usually do not need this option unless you are
                initializing groups with really large file lists.

        -X      Also add removals to dirty db when doing a -TI run.
        -U      Don't mark all other peers as dirty when doing a -TI run.

        -G Group1,Group2,Group3,...
                Only use this groups from config-file.

        -P peer1,peer1,...
                Only update this peers (still mark all as dirty).

        -F      Add new entries to dirty database with force flag set.

        -t      Print timestamps to debug output (e.g. for profiling).
Creating key file:
        csync2 -k filename

Csync2 will refuse to do anything when a /etc/csync2.lock file is found.
```

Figure 4: The Csync[2] help message

The command `csync2 -M` can be used to print the list of files marked dirty and therfore scheduled for synchronization.

5.4   Comparing the hosts

The `csync2 -T` command can be used to compare the local database with the database of the remote hosts. Note that this command compares the databases and not the filesystems - so make sure that the databases are up-to-date on all hosts before running `csync2 -T` and run `csync2 -cr /` if you are unsure.

The output of `csync2 -T` is a table with 4 columns:

1. The type of the found difference: `X` means that the file exists on both hosts but is different, `L` that the file is only present on the local host and `R` that the file is only present on the remote host.

2. The local interface DNS name (usually just the local hostname).

3. The remote interface DNS name (usually just the remote hostname).

4. The filename.

The `csync2 -TT` *filename* command can be used for displaying unified diffs between a local file and the remote hosts.

5.5   Bootstrapping large setups

The `-I` option is a nice tool for bootstrapping larger Csync[2] installations on slower networks. In such scenarios one usually wants to initially replicate the data using a more efficient way and then use Csync[2] to synchronize the changes on a regular basis.

The problem here is that when you start Csync² the first time it detects a lot of newly created files and wants to synchronize them, just to find out that they are already in sync with the peers.

The `-I` option modifies the behavior of `-c` so it only updates the `file` table but does not create entries in the `dirty` table. So you can simply use `csync2 -cIr /` to initially create the Csync² database on the cluster nodes when you know for sure that the hosts are already in sync.

The `-I` option may also be used with `-T` to add the detected differences to the dirty table and so induce Csync² to synchronize the local status of the files in question to the remote host.

Usually `-TI` does only schedule local files which do exist to the dirty database. That means that it does not induce Csync² to remove a file on a remote host if it does not exist on the local host. That behavior can be changed using the `-X` option.

The files scheduled to be synced by `-TI` are usually scheduled to be synced to all peers, not just the one peer which has been used in the `-TI` run. This behavior can be changed using the `-U` option.

5.6 Cleaning up the database

It can happen that old data is left over in the Csync² database after a configuration change (e.g. files and hosts which are not referred anymore by the configuration file). Running `csync2 -R` cleans up such old entries in the Csync² database.

5.7 Multiple Configurations

Sometimes a higher abstracion level than simply having different synchronization groups is needed. For such cases it is possible to use multiple configuration files (and databases) side by side.

The additional configurations must have a unique name. The configuration file is then named `/etc/csync2_myname.cfg` and the database is named `/var/lib/csync2/hostname_myname.db`. Accordingly Csync² must be called with the `-C` *myname* option.

But there is no need for multiple Csync² daemons. The Csync² protocol allows the client to tell the server which configuration should be used for the current TCP connection.

6 Performance

In most cases Csync² is used for syncing just some (up to a few hundred) system configuration files. In these cases all Csync² calls are processed in less than one second, even on slow hardware. So a performance analysis is not interesting for these cases but only for setups where a huge amount of files is synced, e.g. when syncing entire application images with Csync².

A well-founded performance analysis which would allow meaningful comparisons with other synchronization tools would be beyond the scope of this paper. So here are just some quick and dirty numbers from a production 2-node cluster (2.40GHz dual-Xeon, 7200 RPM ATA HD, 1 GB Ram). The machines had an average load of 0.3 (web and mail) during my tests..

I have about 128.000 files (1.7 GB) of Linux kernel sources and object files on an ext3 filesystem under Csync² control on the machines.

Checking for changes (`csync2 -cr /`) took 13.7 seconds wall clock time, 9.1 seconds in user mode and 4.1 seconds in kernel mode. The remaining 0.5 seconds were spent in other processes.

Recreating the local database without adding the files to dirty table (`csync2 -cIr` after removing the database file) took 28.5 seconds (18.6 sec user mode and 2.6 sec kernel mode).

Comparing the Csync² databases of both hosts (`csync2 -T`) took 3 seconds wall clock time.

Running `csync2 -u` after adding all 128.000 files took 10 minutes wall clock time. That means that Csync² tried to sync all 128.000 files and then recognized that the remote side had already the most up-to-date version of the file after comparing the checksums.

All numbers are the average values of 10 iterations.

7 Security Notes

As statet earlier, authentication is performed using the peer IP address and a pre-shared-key. The traffic is SSL encrypted and the SSL certificate of the peer is checked when there has been already an SSL connection to that peer in the past (i.e. the peer certificate is already cached in the database).

All DNS names used in the Csync^2 configuration file (the `host` records) should be resolvable via the `/etc/hosts` file to guard against DNS spoofing attacks.

Depending on the list of files being managed by Csync^2, an intruder on one of the cluster nodes can also modify the files under Csync^2 control on the other cluster nodes and so might also gain access on them. However, an intruder can not modify any other files on the other hosts because Csync^2 checks on the receiving side if all updates are OK according to the configuration file.

For sure, an intruder would be able to work around this security checks when Csync^2 is also used to sync the Csync^2 configuration files.

Csync^2 only syncs the standard UNIX permissions (uid, gid and file mode). ACLs, Linux ext2fs/ext3fs attributes and other extended filesystem permissions are neither synced nor flushed (e.g. if they are set automatically when the file is created).

8   Alternatives

Csync^2 is not the only file synchronization tool. Some of the other free software file synchronization tools are:

8.1   Rsync

Rsync [7] is a tool for fast incremental file transfers, but is not a synchronization tool in the context of this paper. Actually Csync^2 is using the rsync algorithm for file transfers. A variety of synchronization tools have been written on top of rsync. Most of them are tiny shell scripts.

8.2   Unison

Unison [8] is using an algorithm similar to the one used by Csync^2, but is limited to two-host setups. Its focus is on interactive syncs (there even are graphical user interfaces) and it is targeting on syncing home directories between a laptop and a workstation. Unison is pretty intuitive to use, among other things because of its limitations.

8.3   Version Control Systems

Version control systems such as Subversion [9] can also be used to synchronize configuration files or application images. The advantage of version control systems is that they can do three way merges and preserve the entire history of a repository. The disadvantage is that they are much slower and require more disk space than plain synchronization tools.

9   References

[1] Csync^2
http://oss.linbit.com/csync2/

[2] LINBIT Information Technologies
http://www.linbit.com/

[3] DRBD
http://www.drbd.org/

[4] Librsync
http://librsync.sourceforge.net/

[5] SQLite
http://www.sqlite.org/

[6] ROCK Linux
http://www.rocklinux.org/

[7] Rsync
http://samba.anu.edu.au/rsync/

[8] Unison
http://www.cis.upenn.edu/~bcpierce/unison/

[9] Subversion
http://subversion.tigris.org/