

How to synchronise two ALEPH servers via rsync

by

Bernd Luchner, IDS Basel Bern

@

IGeLU Conference 2011, Haifa

ALEPH is „high maintenance“

Some figures: (IDS Basel Bern, 09/2011)

Sublibraries:	200
BIB-Records:	5'634'862
ADM-Records:	5'333'720
Items:	9'462'321
Holding-Records:	393'896
Loans per year:	1,8M (incl. renewals)
Database Volume:	342GB (incl. 58GB publishing)
Database growth:	18% per year
FTE:	30'000

Suppose you have a training server besides your test and your production servers, and the training server configuration should be in line with your production, and *you don't want to care about.*

Then you should synchronise:

the A-tree,
the U-tree,
Z62, Z66, Z67, Z602

(i.e. the permissions and passwords).

You may *not* synchronise:

the database,
Oracle,
the excludes.

(At least we don't.)

Now, if you don't want and cannot do it manually with tar-files, as described in:

[How to transfer configuration tables and Oracle tables from server A to server B 20+ 20101221.doc](#)

See: Ex Libris Documentation Center, Home > Documentation Center > Ex Libris Documentation Center > Aleph > Support > How To from Support by subject > Miscellaneous

you might try

rsync

What is *rsync*?

rsync is a powerful GNU tool for synchronising files on a server or across servers:

<http://samba.anu.edu.au/rsync/>

Man Page:

<http://samba.anu.edu.au/ftp/rsync/rsync.html>

Now the details. You need to:

1)

set up a script with the rsync command,
exchange authorized keys between the servers,
set up a cronjob or a job list entry.

2)

take care with the excludes,
take care with the job timing.

Script example:

```
#!/usr/bin/bash
# Copy with rsync trees /exlibris/aleph/a20_1 and /exlibris/aleph/u20_1
# to alephschool. List in excludelist files that may NOT be transferred.
# The patterns may NOT contain absolute path names.

# pco, 25.06.2010

RSYNCBIN=/opt/csw/bin/rsync
EXCLUDES=/exlibris/aleph/synca2school/excludes.txt
TARGET=alephschool
TIME="$(date +%Y%m%dT%H%M) "
LOG=/var/log/synca2school/${TIME}-xferlog

if [ ` /usr/ucb/whoami ` != "aleph" ]; then
    echo "This is only to be called as user aleph!"
    exit 1
fi

if ! /usr/sbin/ping $TARGET > /dev/null 2>&1; then
    echo "$TARGET is down!" > $LOG
    exit 1
fi

cd /exlibris/aleph
for i in a20_1 u20_1; do
    $RSYNCBIN --rsync-path=$RSYNCBIN --exclude-from=$EXCLUDES --delete -av
    /exlibris/aleph/$i ${TARGET}:/exlibris/aleph/ >> $LOG 2>&1
done
```

Exclude list example:

```
*.tar
files/que_batch*
scratch/*
print/*
import_files/*
source/*
output/*
.dirlist.html
.dirlist_by_date.html
dead.letter
def_aleph.dat
def_version*
ora_aleph/*
oradiag_aleph/*
tmp/*
log/*
log.1/*
log.2/*
apache_media/*
adam/*
dsv01/tab/tab_publish*
dsv13/*
dsv14/*
dsv51/cash/*
dsv51/edi/*
dsv51/dod/logs/*
dsv51/private*
alephe/apache/*
```

```
alephe/aleph_start*
alephe/aleph_start.private*
alephe/aleph_startup*
alephe/aleph_shutdown*
alephe/www_server.conf
alephe/tab/oai/*
alephe/tab/license*
alephe/tab/license.www_x*
alephe/tab/library_relation*
alephe/tab/job_list*
alephe/tab/tab_base.*
alephe/tab/tab_io_remote*
alephe/tab/tab_library_group*
alephe/tab/tab_suf_lib*
alephe/tab/server_ip_allowed*
alephe/tab/server_info*
alephe/tab/server_info_childs*
local/backup/*
local/cleanup_logs/*
local/system/*
local/statistik/*
local/webstatistik/*
local/AL500*
local/version/Alephcom/tab/library.ini
local/version/Alephcom/tab/searbase.dat
local/version_old/*
local/version_save/*
```

Job timing:

1. run your p_file_04 jobs first (if any),
2. run your servicepack installation utility (if any),
3. then run rsync,
4. run p_file_6 jobs on synced server (if any),
5. restart synced server.

How often do we rsync?
once a week.

Since when?
Since more than a year.

How long does rsync take?
Normally: less than a minute.
After servicepacks: a minute.

For example (after a servicepack):
5'593 log entries (sending and deleting),
sent 332'638'742 bytes received 801'256 bytes
4302451.59 bytes/sec

More questions?

Thanks and good luck!

bernd.luchner@unibas.ch