

Re-installing and upgrading Stratum-1 @ RAL: Lessons learned

CVMFS Technical Coordination Meeting
8th of December, 2020

Jose Caballero

Context

- Stratum-1 hosts at RAL were updated to SL7.
- During the process, service was also updated to latest version:

2.6.2 ⇒ 2.7.4

- Most probably every Stratum-1 has already done it.
- Still, the lessons learned may be helpful for similar interventions in the future.

Apache: [No] backward compatibility 2.2 → 2.4

- Several issues after the upgrade related to the fact that httpd was updated from 2.2 and 2.4, and previous configuration didn't work anymore.
- Documentation used to fix some of the problems is:

<https://httpd.apache.org/docs/2.4/upgrading.html>

- Issues:
 - Listening port 80 was set twice, in two different config files. That is not allowed anymore.
 - Not valid anymore mixing *options* with +/- signs and *options* with no sign:

```
Options -MultiViews FollowSymLinks -Indexes ⇒ Options -MultiViews +FollowSymLinks -Indexes
```

- syntax for authorization directives has changed:

```
Order allow,deny ⇒ Require all granted  
Allow from all
```

CVMFS configuration to disable GeoIP

WARNING: this is not about whether the GeoIP functionality is relevant or not !!

- After upgrade, CVMFS server refused to work.

```
[root@stratum-1 ~]# cvmfs_server add-replica -p -o root <url> <pub key>
Installing GeoIP Database... CVMFS_GEO_LICENSE_KEY not set
fail
```

- Solved, for the time being, with configuration:

```
[root@stratum-1 ~]# cat /etc/cvmfs/server.local
CVMFS_GEO_DB_FILE=NONE
```

- If GeoIP is not a mandatory requirement for the CVMFS server, maybe the default configuration should be to NOT require it, and let the sys admins to change it when being used ?

Heartbeat

NOTE: this is only relevant for those sites with 2 hosts in a High Availability (HA) setup.

- Red Hat stop support for HA package **heartbeat** on version 6. Not available anymore for RH 7.
- Replacement seems to be **pacemaker**:

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/high_availability_administration/ch-startup-haaa

- For the time being (hopefully!), we keep using **heartbeat** from a yum repo at CERN:

<http://frontier.cern.ch/dist/heartbeat-el7/RPMS/>

Does anyone have experience with pacemaker?

Passwords, secrets, ...

- After re-installing both Stratum-1 hosts, the files not being managed by the Content Management System (usually for security reasons) were gone:
 - ssh public keys to allow hosts to talk to each other.
 - The **awstats** password file:

```
[root@stratum-1 ]# cat /var/log/squid/squidcron.log
...
...
Calling /etc/awstats/run_awstats.sh /var/log/squid/access.log
ERROR: password is missing: /etc/awstats/password-file. Leaving
...
...
```

Passwords, secrets, ...

- After re-installing both Stratum-1 hosts, the files not being managed by the Content Management System (usually for security reasons) were gone:
 - ssh public keys to allow hosts to talk to each other.
 - The **awstats** password file:

```
[root@stratum-1 ]# cat /var/log/squid/squidcron.log
...
...
Calling /etc/awstats/run_awstats.sh /var/log/squid/access.log
ERROR: password is missing: /etc/awstats/password-file. Leaving
...
...
```

Wrong caching policies (another Apache issue)

- Another problem, apparently as a consequence of the non-backward compatibility of apache, was the caching policy for the `.cvmfspublished` file.
- This file contains several pieces of information, like the revision number.
- With the previous configuration in `/etc/httpd/conf.d/cvmfs.conf`, after upgrading apache 2.2 \Rightarrow 2.4, the caches were keeping that file up to 5 days. As a consequence, clients didn't update repos.

```
[root@client ~]# attr -g revision /cvmfs/alice.cern.ch/  
Attribute "revision" had a 4 byte value for /cvmfs/alice.cern.ch/:  
8069  
[root@stratum-1 ~]# cat -v /storage/cvmfs/alice.cern.ch/.cvmfspublish | grep ^S  
S8079
```

```
[root@client ~]# curl -vs http://<url>/cvmfs/alice.cern.ch/.cvmfspublished 2>&1|grep Cache  
< Cache-Control: max-age=259200  
< X-Cache: MISS from stratum01.gridpp.rl.ac.uk
```


Wrong caching policies (another Apache issue)

- Solution
 - change the syntax in apache config file `/etc/httpd/conf.d/cvmfs.conf`
From

```
AddType application/x-cvmfs .cvmfspublished .cvmfswhitelist
```

To

```
<FilesMatch "^\.cvmfs">  
    ForceType application/x-cvmfs  
</FileMatch>
```

- Stop & restart the frontier-squid servers.

Other considerations

- The data is stored in a secondary partition `/dev/sdb1`, so it is not deleted during the re-installation process of the hosts.
- After each host is re-installed, copy the content of `/etc/cvmfs/keys`.
 - Some sort of backup is needed.
 - In the case of a HA setup, each host is the backup for the other one.
- Run command `cvmfs_server add-replica -p -o root <repo> <url> <keys dir>` for each repo.