

CernVM-FS

Site Deployment

Ian Collier - RAL Tier 1
Atlas Software Workshop
6th April 2011

`ian.collier@stfc.ac.uk`

Principles

- Instead of NFS server and mounts you have cvmfs client mounting virtual filesystem using fuse and downloading files as required by http
- For a site you can (and in fact you should) accelerate this - as well as reducing external network traffic and load on the central cvmfs repositories - by using squid cache(s) at your site
- We save network traffic, software server load, and the experiments have just one place to install software. Squids are easier to resource and maintain than performant NFS servers
- **NOTE The following applies to cvmfs 0.2.61**

What you need

- A squid cache (ideally two or more for resilience rather than load)

Configured (at least) to accept traffic from your site to one or more cvmfs repository servers

You could use existing frontier-squids

- Some worker nodes
- Some rpms installed on your WNs:

`cvmfs` (currently 0.2.61-1)

`cvmfs-init-scripts` (currently version 1.0.8-2)

`cvmfs-keys` (currently version 1.1-1)

`cvmfs-auto-setup` (optional - aimed at T3s)

`fuse` (should be at latest version from OS)

`fuse-libs` (should be at latest version from OS)

`autofs` (should be at latest version from OS)

What you need to configure

- The repositories required at your site (atlas, cms, lhcb, etc.)
- The source repository URL(s) to use

`http://cernvmfs.gridpp.rl.ac.uk/opt/@org@`

`http://cvmfs-stratum-one.cern.ch/opt/@org@`

BNL is coming

Ideally set one primary and at least one secondary - failover is built in to the client

- The size of the cache(s) on your WNs
- The **cvmfs** user account
- You need an entry in the autofs master map **/etc/auto.master** (**/cvmfs /etc/auto.cvmfs**) and also **/etc/fuse.conf** (**user_allow_other**)
- **VO_<voname>_SW_DIR** - i.e:
VO_LHCB_SW_DIR, VO_ATLAS_SW_DIR etc.

What you need to configure

There is also `/usr/bin/cvmfs_config` which can be run manually or be triggered by the `cvmfs-auto-config` rpm (aimed at T3s), since we carry out system config with Quattor I don't have direct experience.

`cvmfs_config`

Common configuration tasks for CernVM-FS

```
Usage: /usr/bin/cvmfs_config {setup  
[nouser] [nocfgmod] [noservice] [nostart] |  
chksetup | showconfig [<repository>]}
```

What you need to configure

Depending on how your existing site software areas are set up, you may need some other steps.

At the moment, for Atlas at RAL, we have some steps to 'simulate' the NFS served area, with some links to it. This would best be worked between the site and those who currently install software to local servers.

For LHCB, for example this was not required - as of this morning they are successfully testing using just cmvfs.

Clearly the medium term goal would be to have perhaps one script to set up the site environment.

Configuration files

`/etc/cvmfs`

```
[cvmfs]# ls -R *
```

```
config.sh  default.conf  default.local
```

config.d:

alice.cern.ch.conf

atlas.cern.ch.conf

atlas.cern.ch.local

atlas-condb.cern.ch.conf

atlas-condb.cern.ch.local

boss.cern.ch.conf

cms.cern.ch.conf

cms.cern.ch.local

grid.cern.ch.conf

hepsoft.cern.ch.conf

hone.cern.ch.conf

lcd.cern.ch.conf

lhcb.cern.ch.conf

lhcb.cern.ch.local

na61.cern.ch.conf

sft.cern.ch.conf

domain.d:

cern.ch.conf **cern.ch.local**

keys:

cern.ch.pub

profile.d:

The ones in red are added by the site (RAL in this case)

Specific changes

`/etc/cvmfs/default.local`

```
CVMFS_USER=cvmfs
CVMFS_NFILES=32768
#CVMFS_DEBUGLOG=/tmp/cvmfs.log
CVMFS_SYSLOG_LEVEL=0
CVMFS_REPOSITORIES=atlas,atlas-condb,lhcb,cms
CVMFS_CACHE_BASE=/pool/cache/cvmfs2/
CVMFS_QUOTA_LIMIT=5000
CVMFS_HTTP_PROXY="http://lcg0679.gridpp.rl.ac.uk:
3128;http://lcg0617.gridpp.rl.ac.uk:3128+http://
lcg0617.gridpp.rl.ac.uk:3128;http://
lcg0679.gridpp.rl.ac.uk"
```

For `CVMFS_HTTP_PROXY`, the client randomly picks the first or second couplet, and binds to the first squid in that couplet, failing over to the second if required. Debug line can be uncommented as required.

Specific changes

```
/etc/domain.d/cern.ch.local
```

```
CVMFS_SERVER_URL=http://  
cernvmfs.gridpp.rl.ac.uk/opt/  
@org@,http://cvmfs-stratum-one.cern.ch/  
opt/@org@
```

Here we tell cvmfs that for repositories hosted originally at CERN (ie all the WLCG experiment software, first look to the replica at RAL, and fail over to CERN. We can add BNL when it becomes available.

But as other cvmfs servers become available (for nightlies, non lhc VOs etc.) this mechanism will allow multiple sources to be accessed from the same WN.

Specific changes

```
/etc/cvmfs/config.d/atlas.cern.ch.local
```

```
CVMFS_QUOTA_LIMIT=10000
```

Here we tell cvmfs that for Atlas, the repository quota limit is 10GB instead of the 5GB in **/etc/cvmfs/default.local**, and so on for the other repositories used if required.

Tools

```
cvmfs_talk
```

```
cvmfs_config chksetup
```

```
cvmfs_config showconfig <repository>.cern.ch
```

```
service cvmfs status
```

```
service cvmfs restartclean
```

```
service cvmfs restartautofs
```

cvmfs_talk allows us to 'interrogate' the caches

cvmfs_config shows/verifies the configuration

the service commands allow granular stopping/starting/restarting of components

Tools - examples

```
cvmfs-talk cache size
```

```
Current cache size is 93MB (98116906 Bytes), pinned: 93MB  
(98115584 Bytes)
```

```
Current cache size is 21MB (22661463 Bytes), pinned: 21MB  
(22660096 Bytes)
```

```
Current cache size is 10MB (10551590 Bytes), pinned: 10MB  
(10550272 Bytes)
```

```
Current cache size is 0MB (19746 Bytes), pinned: 0MB  
(18432 Bytes)
```

Tools - examples

```
service cvmfs status
```

```
automount (pid 2946) is running...
```

```
CernVM-FS mounted on /cvmfs/atlas.cern.ch with pid 32325
```

```
CernVM-FS not mounted on /cvmfs/atlas-condb.cern.ch
```

```
CernVM-FS mounted on /cvmfs/lhcb.cern.ch with pid 32385
```

```
CernVM-FS not mounted on /cvmfs/cms.cern.ch
```

Tools - examples

```
cvmfs_config showconfig atlas.cern.ch
```

```
CVMFS_USER=cvmfs
CVMFS_NFILES=32768
CVMFS_CACHE_BASE=/pool/cache/cvmfs2/
CVMFS_CACHE_DIR=/pool/cache/cvmfs2//atlas.cern.ch
CVMFS_MOUNT_DIR=/cvmfs
CVMFS_QUOTA_LIMIT=20000
CVMFS_SERVER_URL=http://cernvmfs.gridpp.rl.ac.uk/opt/atlas
CVMFS_OPTIONS=allow_other,entry_timeout=60,attr_timeout=60,timeout=60,negative_timeout=60,catalog_timeout=60
CVMFS_DEBUGLOG=
CVMFS_HTTP_PROXY=http://lcg0679.gridpp.rl.ac.uk:3128;http://lcg0617.gridpp.rl.ac.uk:3128+http://lcg0617.gridpp.rl.ac.uk:3128;http://lcg0679.gridpp.rl.ac.uk
CERNVM_CDN_HOST=
CERNVM_GRID_UI_VERSION=
CVMFS_SYSLOG_LEVEL=
CVMFS_TRACEFILE=
CVMFS_DEFAULT_DOMAIN=cern.ch
CVMFS_PUBLIC_KEY=/etc/cvmfs/keys/cern.ch.pub
CVMFS_FORCE_SIGNING=
CVMFS_STRICT_MOUNT=
```

Volunteers

RAL has a full mirror, updated hourly, of the CVMFS repositories at CERN - as described elsewhere.

This is visible externally and we are looking for a small number of sites to test configuring their worker nodes and squids to use the RAL mirror.

Also, the Atlas conditions database repository - atlas-conddb - is now available, but has not been used much. It is different in nature from the software repositories. Some experiments are needed to see what effect it has on squid and local (to WN) cache resources.

Please contact me if you are interested.

Links

Download <http://cernvm.cern.ch/portal/downloads>

Yum <http://cvmrepo.web.cern.ch/cvmrepo/yum>

Mailing list cvmfs-talk@cern.ch

News <http://twitter.com/cvmfs>

Bug Tracker <https://savannah.cern.ch/bugs/?group=cernvm>

RAL CVMFS Pages:

http://www.gridpp.ac.uk/wiki/RAL_Tier1_CVMFS