



The CernVM File System


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

Jakob Blomer, Predrag Buncic, Artem Harutyunyan

ATLAS Software Week, April 2011

Software Distribution with CernVM-FS

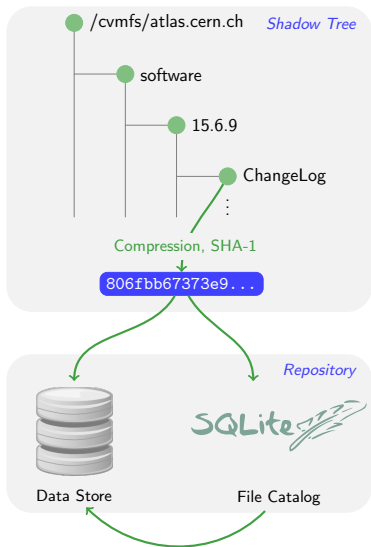
Principle: Virtual software installation by means of an HTTP File System



Essential Properties:

- Distribution of read-only binaries
- Files and file meta data are downloaded on demand and locally cached
- Self-contained (e.g. `/cvmfs/atlas.cern.ch`), does not interfere with base system

Content Addressable Storage



⇒ Immutable Files,
trivial to check for corruption

Data Store

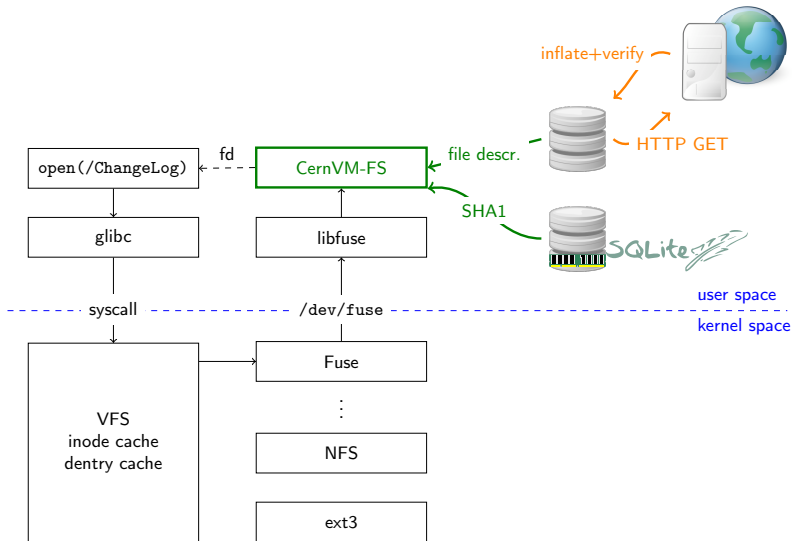
- Compressed Chunks (Files)
- Eliminates Duplicates
- Never Deletes


File Catalog

- Directory Structure
- Symlinks
- SHA1 of Regular Files
- Digitally Signed
- Time to Live
- Nested Catalogs

Components

Client





Mirror servers: Web servers listening on port 80, 8000

Proxy servers: Local load-balanced Squid forward proxy (SL5 Squid)

Fuse Module

- Namespace:
/cvmfs/<repository>
e. g. /cvmfs/atlas.cern.ch
- One process per repository + watchdog
- One mount per repository
- Exclusive cache directory per repository
- Cache directory LRU managed
 - Each file < quota/2
 - Loaded metadata < quota/2
 - Cleanup from quota to quota/2


Mount scripts

- Setup environment (number of file descriptors, access rights, ...)
- Used by autofs on /cvmfs
- Used by system mount

```
mount -t cvmfs atlas.cern.ch  
/cvmfs/atlas.cern.ch
```


Diagnostics

- `service cvmfs probe`
- `cvmfs_config chksetup`
- `cvmfs_config showconfig`
- `cvmfs_fsck`
- `cvmfs-talk`, connect to running instance



Proxies SL5 Squid, Load-Balancing + Fail-Over
e. g. `CVMFS_HTTP_PROXY="A|B|C"`
as of next revision: semantics like Frontier/Squid

Mirrors Fail-Over, set order manually in
`/etc/cvmfs/domain.d/cern.ch.local`
Mirrors at CERN, RAL, BNL operational by the end of the
month

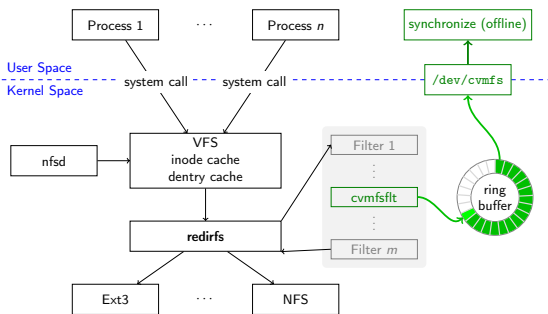


Proxies SL5 Squid, Load-Balancing + Fail-Over
e. g. `CVMFS_HTTP_PROXY="A|B|C"`
as of next revision: semantics like Frontier/Squid

Mirrors Fail-Over, set order manually in
`/etc/cvmfs/domain.d/cern.ch.local`
Mirrors at CERN, RAL, BNL operational by the end of the
month

CernVM-FS Backend Storage

- *Release manager machine* acts as "template".
- Maintains shadow tree and CernVM-FS repository



Publishing:

- 1 Install and test new software
- 2 Synchronize changes with repository
- 3 Sign file catalogs
- 4 Create new file system snapshot

Steps 2–4: Taken care of by
CernVM-FS server tools

Known users: CernVM, ATLAS U.S. Tier3s, RAL, PIC, NIKHEF, CERN, Wuppertal, QMUL, Munich, Lancaster, Dortmund, JINR, ...

- `/cvmfs/atlas.cern.ch` — Production Software
5 release managers, 24 releases SLC4 + 31 releases SLC5
590 GB, 11 Million files, 16 Million entries (shadow)
85 GB and 1.5 Million files (repository)
- `/cvmfs/atlas-condb.cern.ch` — ATLAS Condition Flat Files
Release manager machine hosted by CERN IT
Automatic update several times a day
30 GB, 110 000 files, 7 000 directories, 3 000 symlinks (shadow tree)
30 GB, 70 000 files (repository)
Only fraction of all conditions data
- `/cvmfs/atlas-nightlies.cern.ch` — ATLAS Nightlies
to be done

Mid of April Next CernVM-FS revision

End of April Mirror servers operational

End of May CernVM-FS 2.0 Stable

June/July ATLAS Nightlies available (not mirrored)

Second half of 2011 Migration of release manager machines and backend storage to CERN/IT

Some 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>