



GridPP

UK Computing for Particle Physics

CVMFS Status

Ian Collier

RAL Tier1

ian.collier@stfc.ac.uk

HEPiX Fall, Vancouver, 27th October 2011

Thanks (and apologies) to Steve Traylen



Research & Technology Support Council
e-Science

- History
- Architecture reminder
- Deployment
- Production Status

- Summer 2010 - Jakob Blomer presented to Virtualisation & Multicore workshop @ CERN - RAL noticed a possible solution to ongoing SW server issues
- 2010 Q3 & Q4 - testing at RAL and other sites for Atlas
- End 2010 - LHCb start testing and using at some Tier 1s
- End 2010 - Stratum 1 Replicas established at RAL, BNL and CERN
- 2010 Q4 - 2011 Q1 HEPiX initiated security audit
- 2011 Q1 - Work starts to move support for service to CERN IT
- Now - Production System

CVMFS client:

- Provides a `/cvmfs/` filesystem area.
- Files served from a web-server.

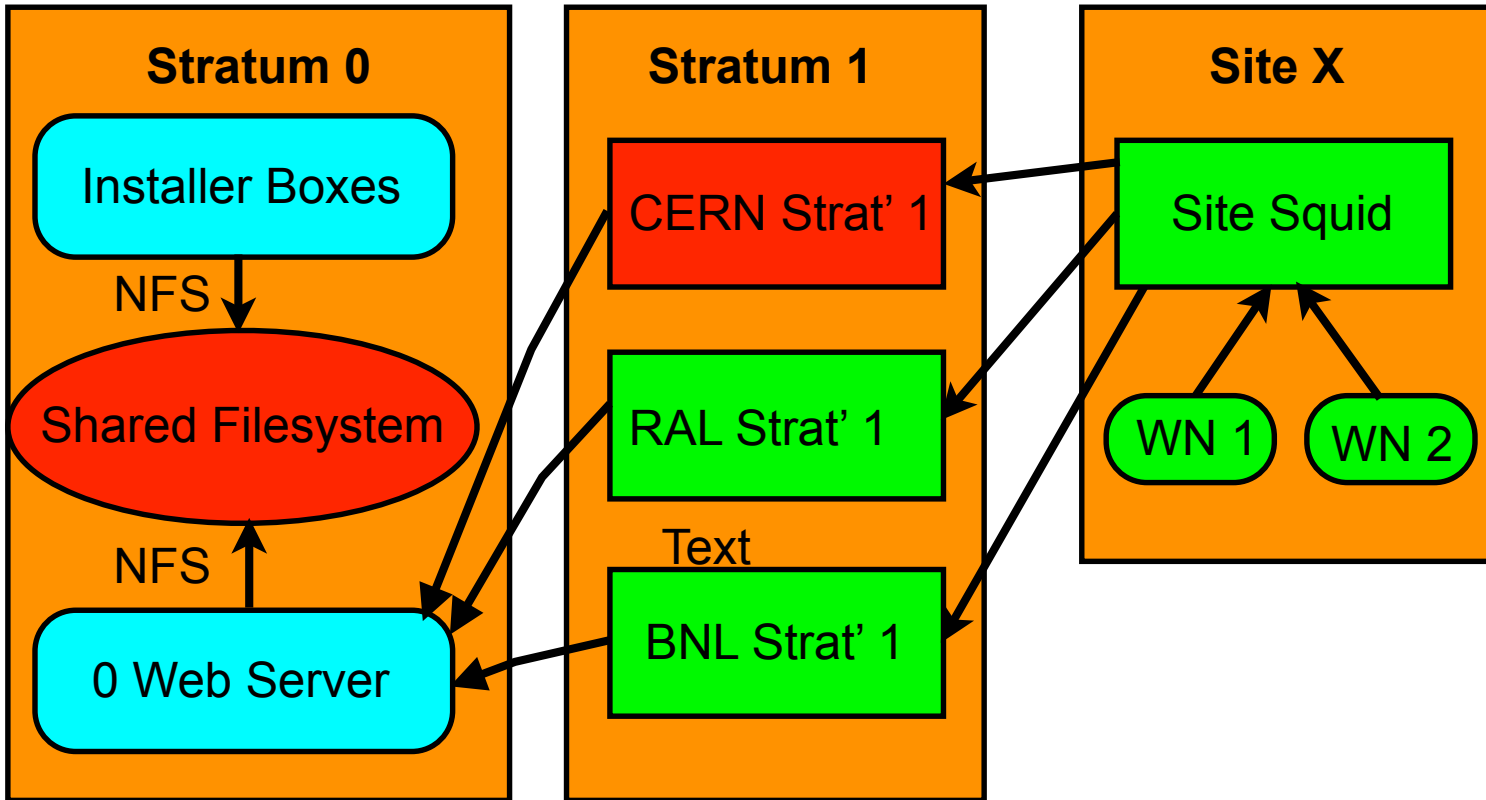
File accesses are intercepted by fuse.

- Metadata operations (`ls`, `cd`, ...) work on an sqlite database - downloaded from web-server.
- File operations (`cat`, `open`, ...) trigger download of file from web server.

Lot's of caching:

- On batch workers across job executions.
- On intermediate squid servers.

Transparent to user/jobs

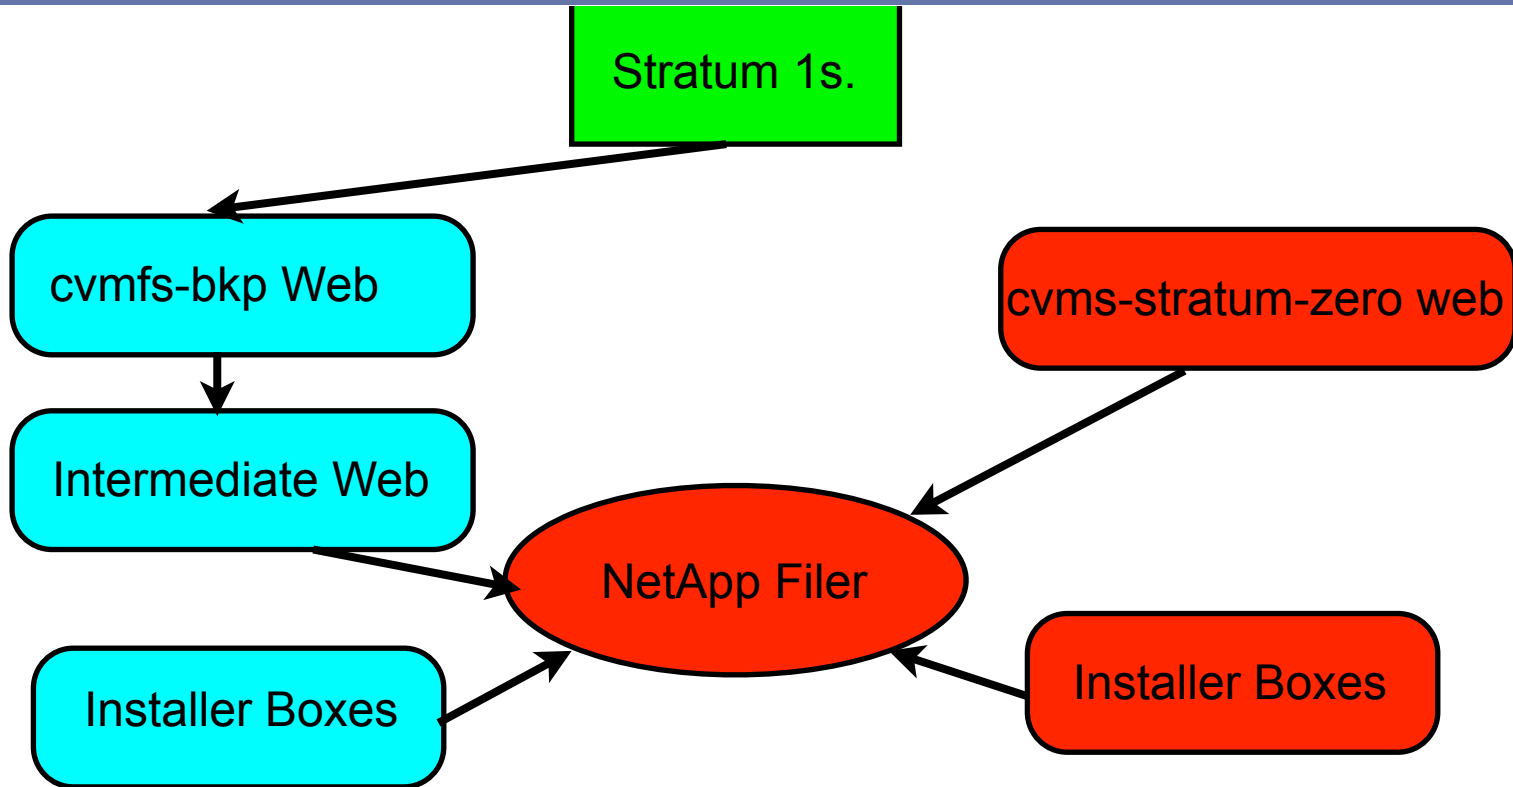


Location - CERN IT, CERN PH, Other

Stratum 1s fully redundant and production.

- No more are needed, nothing to be done.

Stratum 0 migration from PH to IT - next slide



Locations: **CERN IT**, **CERN PH**, **Other**

- IT installers installing on non-production file path.
- IT cvmfs-stratum-zero ready for production now.
- Intermediate web server will be dropped



stratum-zero and IT installer boxes.

- Migration of these is decoupled.

stratum-zero web can be migrated tomorrow.

- Small coordination with stratum one providers but transparent to all.

IT installer boxes, 4 now in place.

- Currently undergoing internal testing
- Testing with experiment installers in progress.
- Eventual switch will be quick:
 - half a day (at most) of downtime for installers.
 - There is no requirement to reinstall all existing software.

All data is now hosted on IT provided NetApp.

- This has improved reliability greatly.

Two parts:

– File Serving:

- Once cvmfs-stratum-zero is source for stratum ones the full chain file serving is in IT.

– File Installation:

- Service won't be improved much if at all by moving nodes from PH to IT.
- Same resilience. A box dies, reinstall next working day.

Central service is as production as it is going to be.

(Next 2 Slides provided by Predrag Buncic)

CernVM-FS was developed in context of an R&D Activity by PH/SFT (Software Development for Experiments) group

- It provides an ideal mechanism for SFT to deliver software to clients (LHC experiments, ...)

Responsibilities for future support and maintenance of software need to be agreed (in-house product)

SFT has been very active in support of this first exploitation phase

- bug fixes; 13 releases so far in 2011
- responding to questions and solving installation issues; ~900 mails in 2011

Requests for new features.

In pipeline:

- ATLAS: support for diskless servers, sharing local cache between repositories, exporting CVMFS via NFS
- LHCb : faster turn-over for publishing updates (<1hour)
- CMS : Mac OS X support
- Small projects: simplified tool set to create repositories
- Grid Sites : monitoring, security, encrypted repositories (Nordic Tier 1), extended documentation (RAL), archival of old software releases (CERN)

Natural for SFT to continue providing software maintenance in immediate future (next 2 years)

- Main goal will be to ensure stability and performance

One solution can never work everywhere however:

- CVMFS does not alter the one node installation method. i.e it does not break anything that exists.

Sites with SUSE and Ubuntu are running CVMFS clients.

- MacOS - code exists, not released.

Scenarios that are still difficult:

- Diskless batch workers.
- WNs with software provision strictly via NFS or similar.
 - Both will require client development.

CVMFS publication of experiment software is production ready.

- CERN IT under WLCG has committed to run central services, long term in equivalent fashion to LFC, VOMS,
- Service is now as stable as it is likely to ever be.
 - IT will do no better than PH is doing for remaining items.
- Final service layout awaiting IT installers to be fully operational.

Support in place - GGUS, CERN Service NOW.

- Items for CVMFS developers will be relayed just as for FTS, VOMS, apache, ..
- Existing lines to CVMFS developers will persist.

CVMFS discussion list.

- cvmfs-talk@cern.ch

CVMFS Releases, including release notes & configuration examples.

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

CERN-IT CVMFS Central Services

- A new homepage for central CVMFS service
- <https://twiki.cern.ch/twiki/bin/view/CvmFS/>