



CernVM-FS – Beyond LHC Computing

Catalin CONDURACHE, Ian COLLIER
STFC Rutherford Appleton Laboratory, UK Tier-1 Center, Didcot, UK

In the last three years the CernVM FileSystem (CernVM-FS) has transformed the distribution of experiment software to WLCG sites. Now established and proven to work at scale, CernVM-FS is beginning to perform a similar role for non-LHC computing. The deployment of CernVM-FS service at RAL Tier-1 is presented, as well as the proposed development of a network

Stratum-0 and Stratum-1 replicas somewhat modeled upon the infrastructure developed to support WLCG computing. A case study of one non-LHC Virtual Organization (VO) is also included, describing their use of the CernVM-FS Stratum-0 service, along with a web interface intended to be used as a tool to upload software at Stratum-0 sites.

What is CernVM-FS?

- Read-only, globally distributed filesystem optimized for software distribution, originally developed for the CERN Virtual Machine project.
- The data integrity and validity are ensured by the digitally signed file catalog and access authentication for software server updates (done by Software Grid Manager or other privileged member of the VO).
- Built using standard technologies (fuse, sqlite, http, squid proxies and caches).
- It removes the need for local installation jobs and conventional NFS servers at sites and its use standardizes the computing environment across the Grid.
- Once the signed catalog has been downloaded and mounted using fuse, metadata operations require no further network access. Together with the file based de-duplication this makes CernVM-FS very efficient in terms of disk usage and network traffic.
- The software needs one single installation and then is available at any site with CernVM-FS client installed

CernVM-FS History at RAL Tier-1


- Summer 2010 - RAL was the first Tier-1 centre to test CernVM-FS at scale and worked towards getting it accepted and deployed within WLCG.
- February 2011 – first global CernVM-FS Stratum-1 replica for LHC VOs in operation outside CERN.
- December 2012 – first CernVM-FS Stratum-0 services deployed for mice and na62 VOs supported by the GridPP UK Project.
- January – June 2013 – CernVM-FS Stratum-0 service extended to international small VOs (hone, phys.vo.ibergrid.eu, enmr.eu).
- Ongoing involvement in verification of CernVM-FS client software at scale before it is released into production.
- Now entering the 'EGI service' phase with more Virtual Organizations supported (hyperk.org, biomed, vlemed, auger, glast.org), offering also a separate Stratum-1 service for non-LHC VOs only.

CernVM-FS Task Force


- A working group setup by European Grid Infrastructure (EGI) to establish a CernVM-FS infrastructure that allows EGI VOs to use it as a standard method of distribution of their software across the EGI computing resources by:
 - promoting the use of CernVM-FS technology by VOs
 - creating a network of sites providing CernVM-FS services (Stratum-0, Stratum-1, squid)
- The Task Force also promotes cooperation with other organizations, such as Open Science Grid (OSG) and WLCG, on monitoring tools for CernVM-FS and by cross-replicating repositories for VOs supported by multiple collaborations.

CernVM-FS 'relaxed' Topology

The proposed model for the EGI community



CernVM-FS WLCG deployment



CernVM-FS EGI deployment

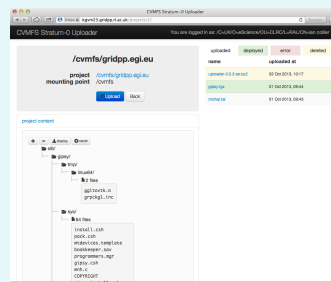
enmr.eu CernVM-FS Stratum-0 at RAL Tier-1

- WeNMR is a Virtual Research Community (VRC) supported by EGI which aims to bring together complementary research teams in Structural Biology and Life Sciences into a VRC at a worldwide level and provide them with a platform integrating and streamlining the computational approaches necessary for NMR and SAXS data analysis and structural modelling. Its Virtual Organization – enmr.eu – is currently the largest in the area of Life Sciences with over 570 registered users and a steady growth rate.
- enmr.eu has an active presence across the Grid at more than 25 sites which makes the distribution of new software releases time consuming.
- The deployment and configuration of the CernVM-FS Stratum-0 repository for enmr.eu at RAL Tier-1 has reduced the effort required at supporting sites maintaining multiple NFS software areas.
- Installation jobs run by the enmr.eu/Role=Manager at RAL Tier-1 batch farm upload and configure new software releases under /cvmfs/wenmr.egi.eu directory.
- Work is ongoing to consolidate the enmr.eu Stratum-0 and to move toward a CernVM-FS only environment.

- The software is installed by LHC VOs at Stratum-0 hosted at CERN and replicated to Stratum-1 hosted by WLCG Tier-1 sites
- CernVM-FS clients connect to one of the Stratum-1 services (via local squid caches)
- Transparent failover to other Stratum-1 service in case of connection problems

- Stratum-0s are disjoint and represent the source repositories where software is installed by the EGI VOs
- Stratum-0 and Stratum-1 can be geographically co-located, or not
- Stratum-1 can replicate a whole Stratum-0 (→), or can partially replicate (---→) – the 'relaxed' model
- A distributed network of Stratum-1s makes CernVM-FS software distribution more resilient

CernVM-FS Stratum-0 Web Frontend



- Web application for CernVM-FS Stratum-0 uploads used as an alternative to installation jobs or 'power users'.
- Developed by a student on an Erasmus Programme placement at RAL-Tier 1 UK.
- Users can upload tarballs and unpack them within the /cvmfs/<repo_name> 'space', followed by synchronization with the real CernVM-FS Stratum-0 repository.
- Features
 - Authenticates with X509 certificates (managed by a web server)
 - Further authentication mechanisms can be added