



CernVM-FS - *the* Tool for Software Distribution

EGI CVMFS Task Force Status

Catalin Condurache

STFC RAL Tier-1

EGI Community Forum, Helsinki, 20 May 2014



Science & Technology Facilities Council
Rutherford Appleton Laboratory



GridPP

UK Computing for Particle Physics



Outline

- Introduction
- Brief history
- CernVM-FS deployment at RAL
- EGI CVMFS Task Force
- EGI CVMFS Infrastructure
- CVMFS for non-LHC - status and plans





Introduction - CernVM-FS ?

- Read-only, globally distributed filesystem, optimized to access experiment software and conditions data
- Built using standard technologies (fuse, sqlite, http, squid and caches)
- Files and directories are hosted on standard web servers and mounted in universal */cvmfs* namespace
- Software needs one single installation, then it is available at any site with CVMFS client installed and configured
- Repository meta-data and files are downloaded on demand and locally cached





Brief History - CernVM-FS

- Last 4 years it transformed distribution of experiment software and conditions data to WLCG sites
 - No need of
 - local installation jobs
 - local SW tags
 - (heavy loaded) file servers
 - It relies on a robust decentralised network of repositories replicas
- In parallel, growing number of repositories and CVMFS servers around world serving other user communities
- 18 Months ago RAL Tier-1 offered non-LHC Stratum-0 service





CernVM-FS Deployment at RAL

- Supported by GridPP UK project
- Initially (Dec'12) for UK VOs, then extended to international small VOs (2+ NGIs)
- Local installation jobs used to automatically publish the Stratum-0 - now obsolete
- Replaced by web interface used by SGMs to upload, unpack tarballs and publish
- Recently, GSI interface added to transfer and process tarballs and publish



CernVM-FS Deployment at RAL

CVMFS Stratum-0 Uploader

You are logged in as: /C=UK/O=eScience/OU=CLRC/L=RAL/CN=ian.collier

/cvmfs/gridpp.egi.eu

project [/cvmfs/gridpp.egi.eu](#)
mounting point [/cvmfs](#)

[Upload](#) [Back](#)

uploaded	deployed	error	deleted
name	uploaded at		
uploader-0.0.3.tar.bz2	02 Oct 2013, 10:17		
gipsy.tgz	01 Oct 2013, 09:44		
michal.tar	01 Oct 2013, 09:43		

project content

- + - [deploy](#) [mkdir](#)
- sl6/
 - gipsy/
 - tmp/
 - linux64/
 - 2 files
 - ggitovtk.o
 - grpckg1.inc
 - sys/
 - 64 files

CernVM-FS Deployment at RAL

- At the moment 11 repos are hosted and published at RAL
- Separate Stratum-1 service for non-LHC VOs
 - initially shared the LHC Stratum-1



EGI CVMFS Task Force

- EGI initiative to establish a CVMFS infrastructure that allows EGI VOs to use it as a standard method of distribution of their SW at grid sites
- Kick-off meeting - August 2013
- Members: reps from NGIs and other collaborations, VO SGMs, sysadmins at sites
- Regular meetings
- https://wiki.egi.eu/wiki/CVMFS_Task_Force





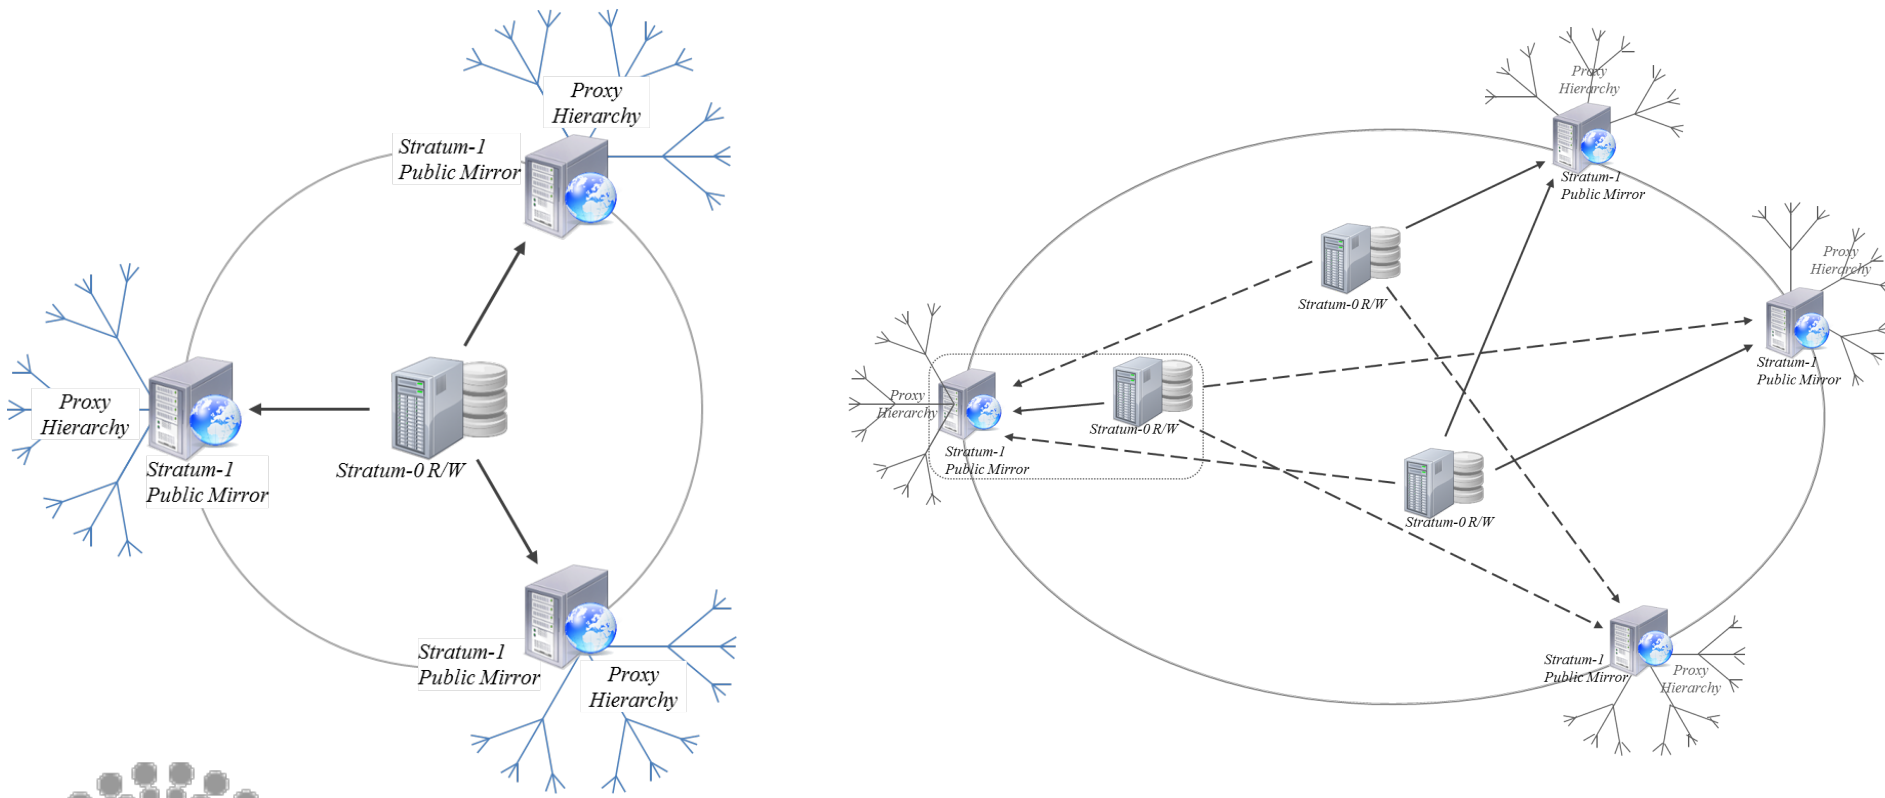
EGI CVMFS Task Force

- It promotes the use of CVMFS technology by user communities
 - “Developments for/by LHC experiments are/should be re-used by others” - CHEP’13 Track 3A Summary
 - Good example: CVMFS is *the* tool for software distribution
 - Bad example: ...
- It creates a network of sites providing CVMFS services (Stratum-0, Stratum-1, squid)
- It encourages cooperation with other organizations (OSG, WLCG, regional grid infrastructures) by cross-replicating repositories for VOs supported by multiple collaborations



EGI CVMFS Infrastructure

- The model follows the CVMFS for WLCG topology



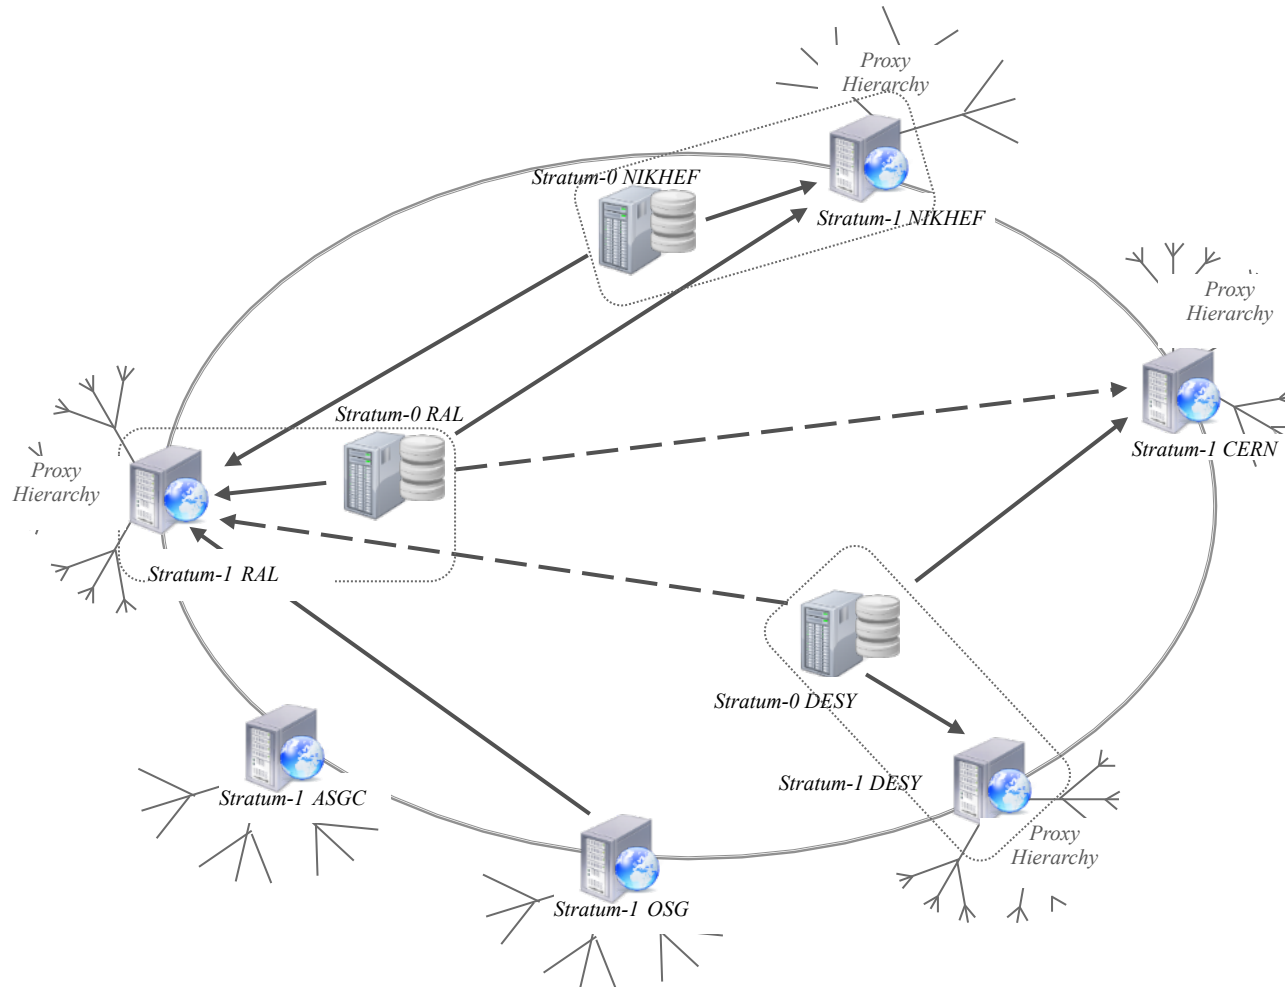


EGI CVMFS Infrastructure

- *'relaxed'* topology
- Stratum-0 are disjoint and represent the source repositories where SW is installed by VOs
- Stratum-0, Stratum-1 can be geographically co-located or not
- Stratum-1 can replicate an entire Stratum-0 or can only partially replicate - the *'relaxed'* model
- It partially makes use of the existent hierarchy of proxy servers used for LHC software distribution



EGI CVMFS Infrastructure



EGI CVMFS Infrastructure

- 17 SW repositories currently hosted (HEP and non-HEP)
 - 5 repos at the time of kick-off
 - BIG change in one year!
- Main infrastructure
 - Enough for new VO's (SW repos) on medium term
 - Stratum-0 - SW repositories
 - RAL - 160GB
 - DESY - 20 GB
 - NIKHEF - 12 GB

EGI CVMFS Infrastructure

- Stratum-1 - repository replica

	gridpp.ac.uk	desy.de	amc.nl
RAL	😊	😊	😊
NIKHEF	😊	😞	😊
DESY	😞	😊	😞
CERN	😊	😊	😞
OSG	😊	😞	😞
ASGC	😊		

EGI CVMFS Deployment Issues

- Experiments
 - Are lacking manpower to progress more
 - Technology and its advantages are understood, just they do not have time
 - Subgroups within same VO are using subsets of SW at specific sites
 - Coordination needed to consolidate SW under single CVMFS space
 - Sometimes VOs wait too long for sites to help them with initial tests before rolling out more widely



EGI CVMFS Deployment Plans

- Stay in touch with VO SGMs
 - Software distribution consolidation
 - 2 more experiments to switch to CVMFS shortly
- WLCG resources to be used where possible
 - Italian sites ready to support EGI VOs - need formal request from VOs
 - French sites almost ready
 - Regional Squid servers to be open to EGI VOs as well (NGIs)





EGI CVMFS Deployment Plans

- NGIs could approach sites and nicely ask them to install / configure CVMFS clients for small VOs
 - Once technology made available (or config done), experiments might be less reluctant to try it
- Cross-replication in Asia (also Latin America?) might also give a push to worldwide VOs (enmr.eu, biomed) to try and consolidate their SW distribution





EGI CVMFS Deployment Plans

- CernVM-FS Working Group
 - CERN, OSG developers, CVMFS managers
 - Involved in the effort to build a more flexible worldwide CVMFS infrastructure
 - Key and configuration distribution
 - Automatic ordering of Stratum-1s
 - Monitoring





Thank you!

(more) Questions?