



EGI CernVM-FS Infrastructure - Status and Plans

CernVM-FS at RAL - Client View

Catalin Condurache, John Kelly

STFC RAL Tier-1

CernVM Users Workshop, RAL, 6-8 June 2016



Science & Technology Facilities Council
Rutherford Appleton Laboratory



GridPP

UK Computing for Particle Physics

Outline

- EGI CernVM-FS infrastructure - Catalin
 - History, status, plans
- CernVM-FS at RAL - client view - John



Bit of History - CernVM-FS at RAL

- Following proven success in experiments software distribution using CernVM-FS to WLCG sites...
- ...RAL Tier-1 started non-LHC Stratum-0 service
 - September 2012
 - '*gridpp.ac.uk*' CernVM-FS domain
- Supported by GridPP UK project
- Initially for UK VOs, then extended





Bit of History - EGI CernVM-FS Task Force

- Expansion to an EGI level infrastructure
 - The European Grid Infrastructure (EGI) enables access to computing resources for European scientists and researchers from all fields of science, from High Energy Physics to Humanities
 - Initiative (Aug 2013) to establish a CernVM-FS infrastructure that allows EGI VOs to use it as a standard method of distribution of their software at grid sites



Bit of History - EGI CernVM-FS Task Force

- Purposes

- To promote the use of CernVM-FS technology among user communities
- To create a network of sites providing CernVM-FS services (Stratum-0, Stratum-1, Squid)
- To encourage cooperation with other organizations (OSG, WLCG, regional grid infrastructures) by cross-replicating repositories for VOs supported by multiple collaborations

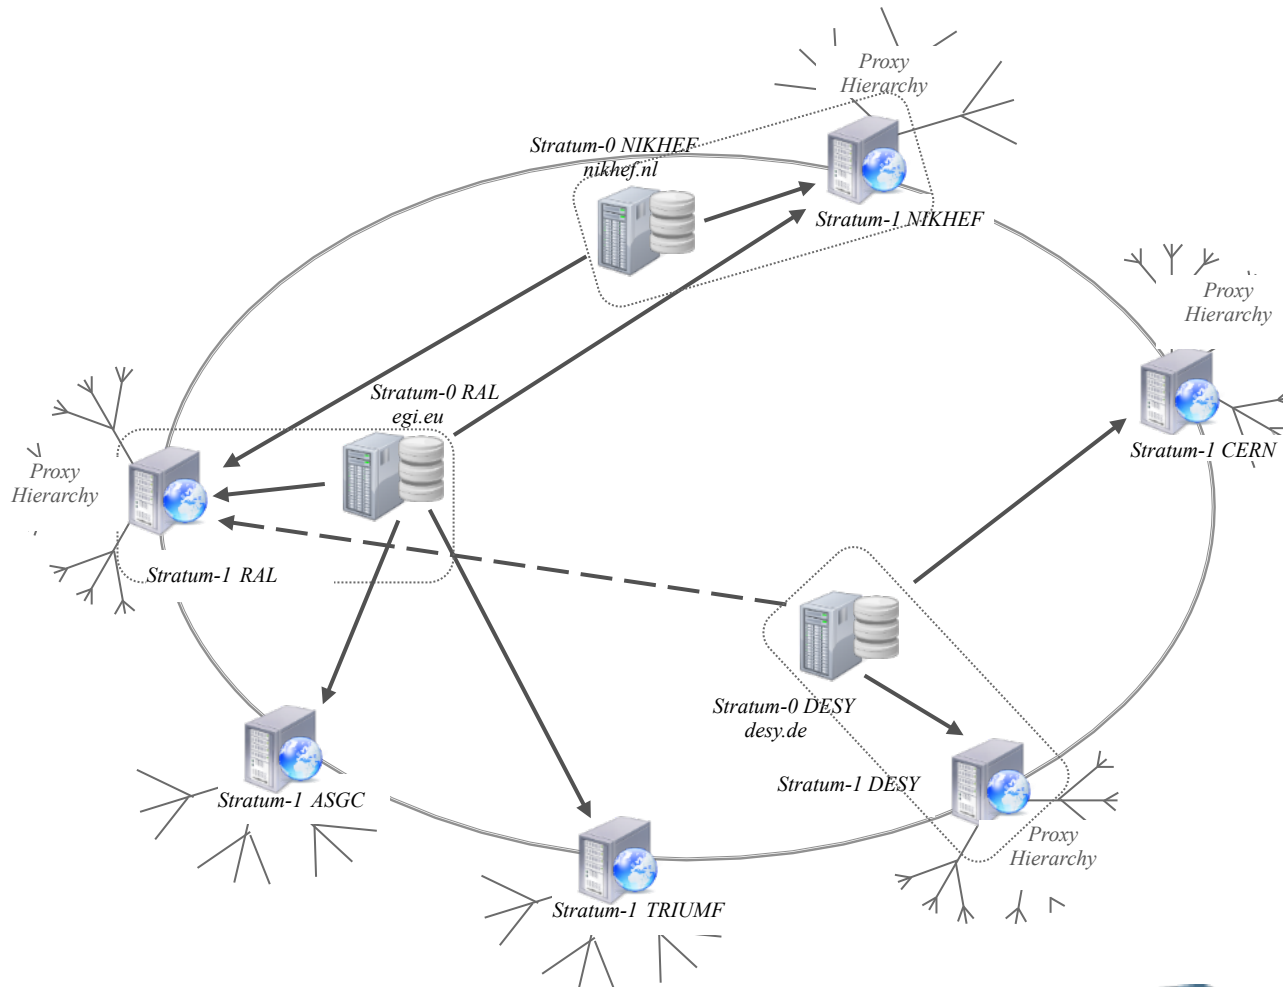


EGI CernVM-FS Infrastructure Status

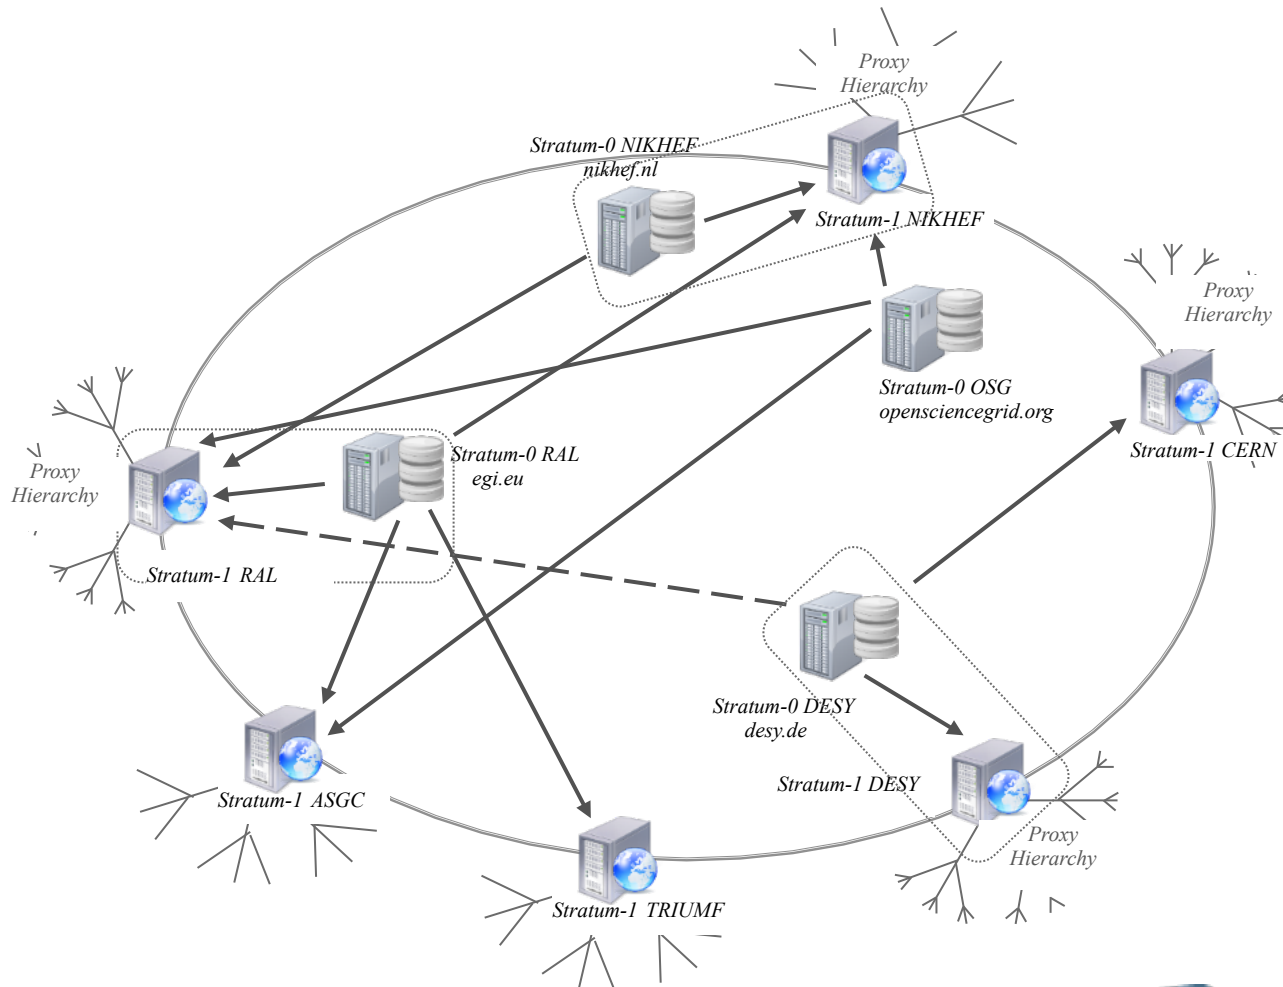
- ‘*egi.eu*’ new space name for repositories
 - ‘*gridpp.ac.uk*’ used for UK regional VO repos (5)
- 30 software repos currently hosted and published at RAL - 900+ GB
 - HEP and non-HEP
 - Stratum-1s at RAL, NIKHEF, TRIUMF, ASGC
 - 5 repos at the time of EGI CernVM-FS TF kick-off
 - BIG change in nearly three years!



EGI CernVM-FS Infrastructure Status

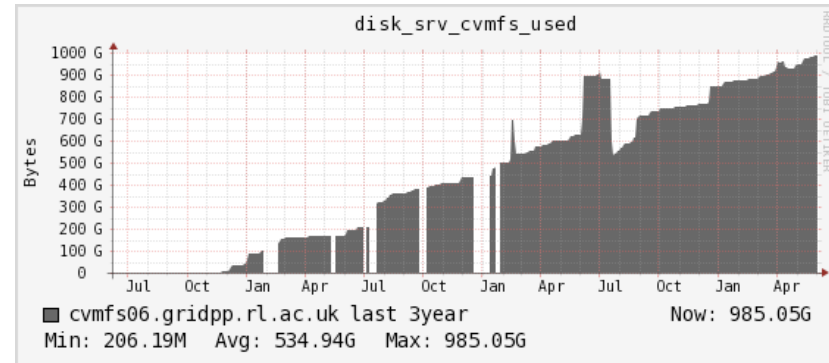


EGI + OSG CernVM-FS Infrastructure Status



EGI CernVM-FS Infrastructure Status

- Stratum-0 at RAL
 - VM, 24GB RAM, 1.2TB, cvmfs-server v2.2.0-0.131.78...
 - Plan to be replaced with an existing Stratum-1 node (see next slides) and v2.2.3+
 - 30 repos
 - 2.73 Mio files
 - 8.57 Mio on uploader
 - 980 GB
 - Between 36KB and 3MB average filesize per repo




EGI CernVM-FS Infrastructure Status

- 30 repositories on RAL Stratum-0
 - Some more active / less active
 - pheno.egi.eu - 465 releases
 - t2k.egi.eu - 152 releases
 - phys-ibergrid.egi.eu - 5 releases
 - glast.egi.eu - 4 releases
 - Some big / not so big
 - chipster.egi.eu - 313GB
 - biomed.egi.eu - 177GB
 - t2k.egi.eu - 59GB
 - phys-ibergrid.egi.eu - 69MB
 - ligo.egi.eu - 2.9GB
 - wenmr.egi.eu - 3.8GB

EGI CernVM-FS Infrastructure Status

- Stratum-1 at RAL - *cvmfs-egi.gridpp.rl.ac.uk*
 - 2-node HA cluster, metal boxes, 32GB, 12TB
 - Integrated with WLCG Stratum-1 (*cernvmfs* endpoint)
 - It replicates 58 repos
 - '*egi.eu*', '*cern.ch*', '*nikhef.nl*',
 - '*desy.de*', '*gridpp.ac.uk*',
 - '*opensciencegrid.org*' domains
 - Very soon to be replaced with new pair of servers (64GB, 55TB, RAID6) and v2.2.3



Over 10.5 TB



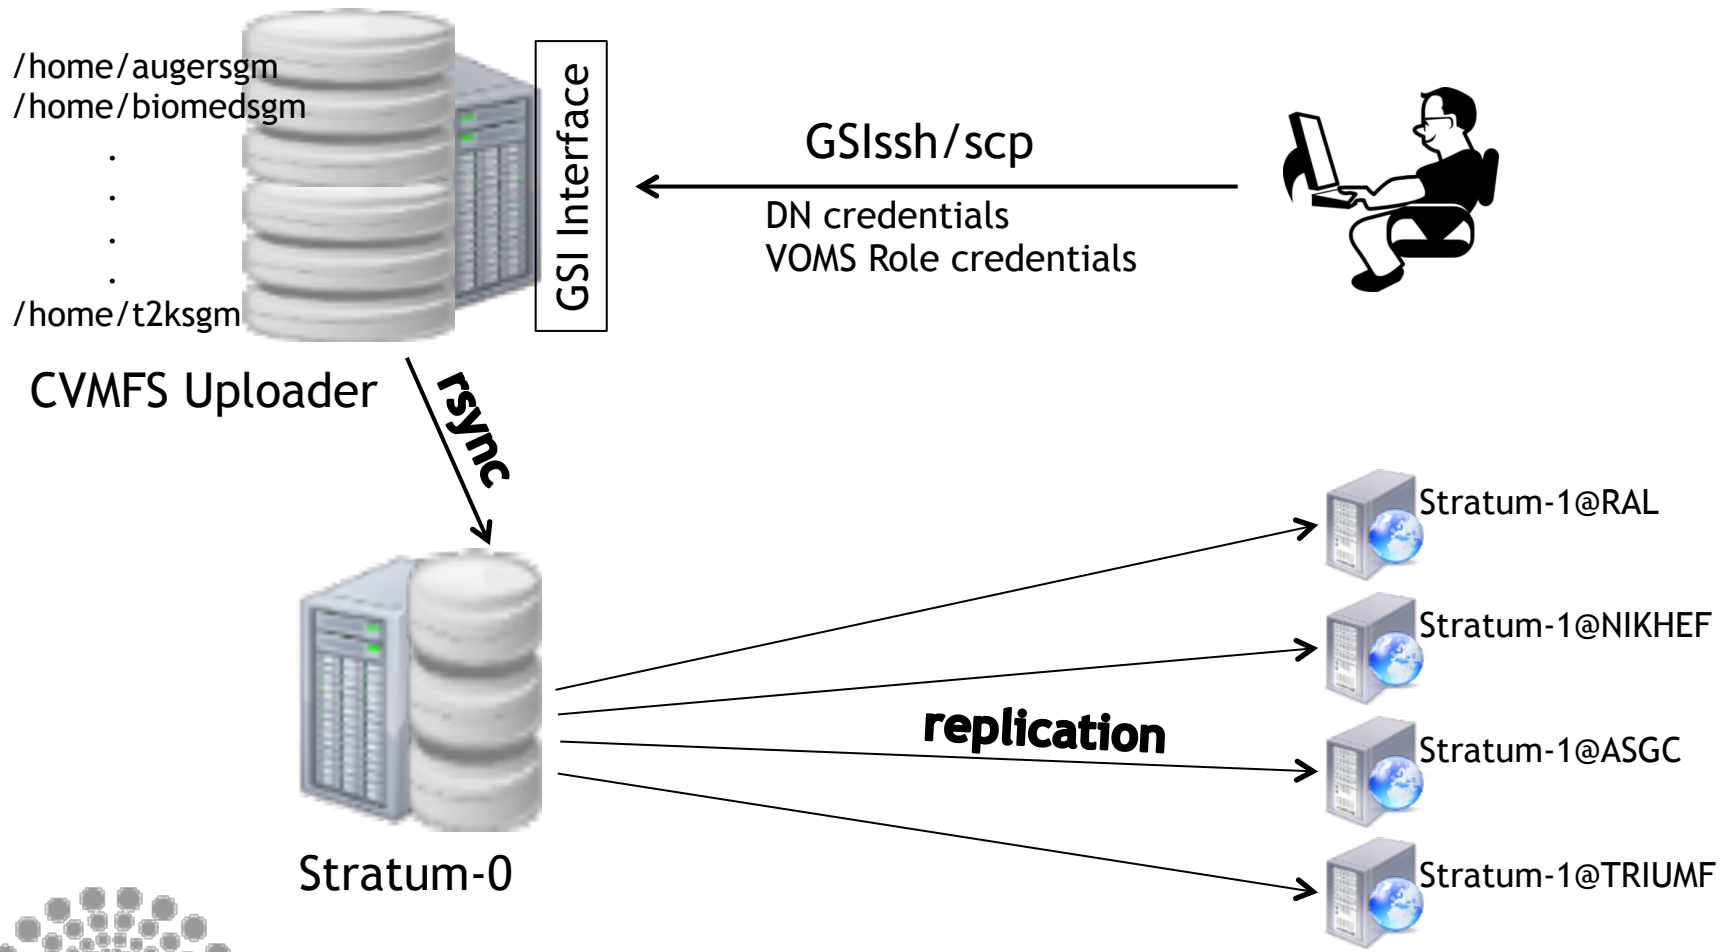


EGI CernVM-FS Infrastructure - RAL

- CVMFS uploader - old-ish disk server, 12GB RAM, 36TB disk
- Squid machines (*cvmfs-squid* alias) at RAL are shared for Frontier and LHC/non-LHC CernVM-FS access
- EGI Early Adopter site
 - Staged Rollout process for server and client releases



Software Installation Mechanism at RAL





EGI CernVM-FS Infrastructure - Plans

- Operations Level Agreement for Stratum-0@RAL
 - Between STFC and EGI.eu
 - Provisioning, daily running and availability of service
 - Service to be advertised through the EGI Service Catalogue
- Use S3 interface to CEPH backend storage for Stratum-1@RAL
 - Initial tests successful
 - Waiting now for 'production' quality (S3 + CEPH storage)





EGI CernVM-FS Infrastructure - More Plans

- Squid network/access for FedCloud
 - Discussion with FedCloud people re: setting up a network of squids
 - Or a mechanism to discover the closest available squid integrated into the client configuration
- Integration of Africa-Arabia repositories within egi.eu name space
 - Long outstanding task
 - Maybe configurations config-egi.egi.eu?





Acknowledgments

- Many, many thanks to CernVM[-FS] team (Dave, Rene, Jakob) for their invaluable support
- Also thanks to *cvmfs-*@cern.ch* members





Thank you and over to John!

