



CernVM Program of Work 2025

Jakob Blomer, Laura Promberger, Valentin Völkl



SFT Group Meeting

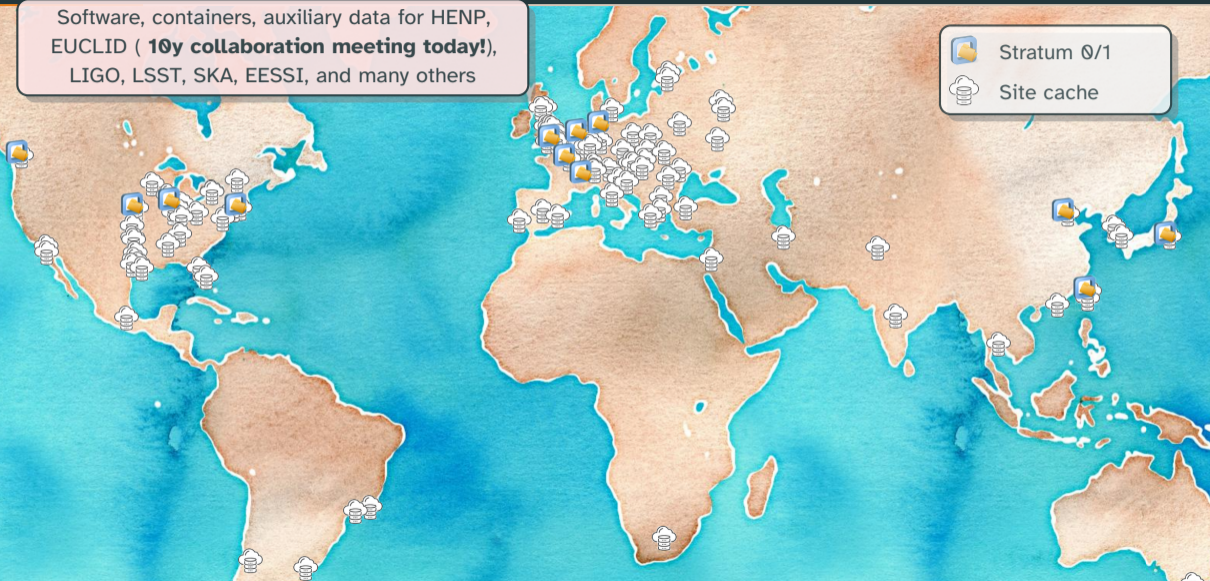
22 January 2025

At a Glance: CernVM-FS Deployment (Grid)



Software, containers, auxiliary data for HENP, EUCLID (**10y collaboration meeting today!**), LIGO, LSST, SKA, EESSI, and many others



 Stratum 0/1
 Site cache



At a Glance: CernVM-FS Deployment (Grid)



Software, containers, auxiliary data for HENP,
EUCLID (**10y collaboration meeting today!**),
LIGO, LSST, SKA, EESSI, and many others

 Stratum 0/1
 Site cache

- ~ 15 Stratum 1s (Europe, North America, Asia)
outage of RAL S-1 for 2nd half of 2024
absorbed by Nikhef
- ~ 5.4 B files in the /cvmfs tree (+210 % in last year)
- ~ 2 PB of data accessible through /cvmfs
out of which ~1PB in *external* files
proven to scale up to 100 PB
- ~ 3550 container images (+60 % in last year)
- ~ 260 repositories (+15 % in last year)



Extras:

- cvmfsexec
- cvmfs-servermon
- github-action-cvmfs
- cvmfs-x509-helper
- repository monitor
- ...

Stand-alone utilities

Preloader

Shrinkwrap

Services (Go)

containerd snapshotter
(now in production)

Container Publishing Tools

Gateway Services

Core Software

Client

Fuse module, libcvmfs,
cache plugins

Server

publisher tools, libcvmfs_server,
Geo-API



Extras:

- cvmfsexec
- cvmfs-servermon
- github-action-cvmfs
- cvmfs-x509-helper
- repository monitor
- ...

Stand-alone utilities

- ~150 kLOC, mostly C++, then Go, Bash
- 10+ supported platforms
- ~1/3 of the code base are tests
- 40-50 bugs / year reported, diagnosed, and fixed
- Additional 40-50 improvement tickets / year, 2024: 3 labeled as new features

In line with experiments request to put
"Priority on bug fixes, client and Stratum 0 speedup rather than new features"

Services (Go)

containerd snapshotter
(now in production)

Container Publishing Tools

Gateway Services

Core Software

Server

libcvmfs_server,

API

CernVM Team




		2024	2025
Jakob Blomer	Staff	25 %	10 %
Valentin Völkl	Staff	100 %	100 %
Laura Promberger	Fellow	100 %	50 %
Amal Thundiyil	Tech	70 %	—
Siem Vermeulen	Tech	—	50 %
<i>TBS</i>	Tech	—	50 %
<i>TBS</i>	Fellow	—	60 %
	FTE	~2.05	~2.5

Note: 1 technical student externally funded by Jump Trading

Significant code contributors:

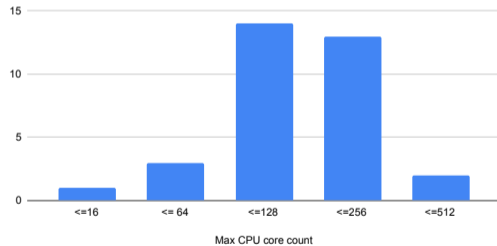
Matt Harvey (Jump Trading), Dave Dykstra (FNAL), Yuriy Belikov (Ukrainian/GSoC student)

Review of 2024






What is the maximum CPU core count on machines at your sites?

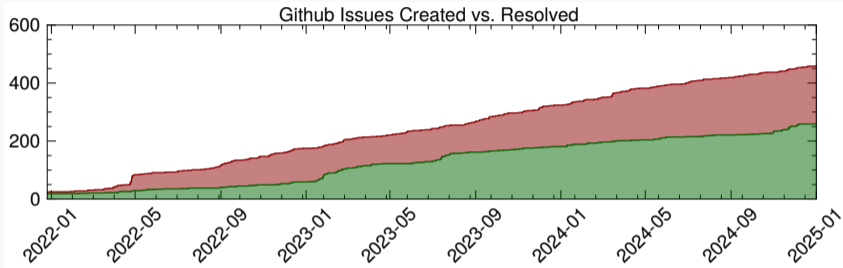


Which version of CVMFS is currently deployed at your site?

[Copy chart](#)

34 responses





Support load in 2025:

- ~ 350 posts on mailing lists and forum
- Occasional tickets in SNOW, GGUS

Tickets and PRs in 2024:

- 84 open issues, 51 closed
- 32 open pull requests

Growing backlog!

- Bugfix sprint / hackathon planned for 2025
- Still relevant open issues in JIRA



Highlights

- Client resilience improvements for extreme conditions
 - Out of memory recovery
 - Fix of rare races when root file catalogs are reloaded
- Container integration: production release and benchmarking of snapshotter
- Prototype on MacOS using FUSE-T library
- Metalink support (for Stashcache)
- Infrastructure migrations and updates
- CernVM Workshop @ CERN with >45 participants with speakers from libFUSE, Jump Trading, Varnish, EESSI
- Dissemination: ▶ CHEP'24 ▶ HEPiX Spring

Unfinished Tasks

- Gateway functionality and performance improvements (*ongoing, slightly behind schedule*)
- Prefetching prototype
- Production release of FUSE-T on MacOS
 - More issues than anticipated, but likely possible
- Update of container unpacker

CernVM Appliance 2025



1. Ready to use platform for HEP application stacks
2. Reference platform for **long-term data preservation**

2025 Plan of Work **est 3 FTW [VV]**

- Upgrade of CernVM 4
 - MicroVM model still has many users (BaBar, Boinc, ATLAS HLT) for whom CernVM 5 is a hard sell
- Deprecate `cernvm.cern.ch/config`
- Support and advice for software preservation and open data efforts
 - Investigate possibility of using CernVM 5 - better suited in this context!

CernVM File System 2025



1. **Integration in HPC sites**
2. **Technology exploration of VARNISH as a proxy server**
3. **Carry-over tasks: ZSTD compression, file bundles**
4. **Stratum-1 tooling efficiency**
5. **Many-core performance engineering**



- Documentation
 - Switch documentation from RST to markdown + mkdocs
 - Restructure documentation
 - Conform to Debian/EPEL packaging guidelines (man pages ...)
- Infrastructure and Testing
 - Finish cc7 infrastructure migrations
 - Move to new cloudtest setup, using Jenkins Openstack plugin, fix flaky tests
- Code renovation
 - Switch to C++11 and remove backported stdlib functionality
 - Make Codebase conform to clang-format and clang-tidy
 - Virtualize cloud test suites, increase coverage of normal suite, add stress test suite [[summer student proposal](#)]



- File-bundle prefetching est 8 FTW [VV]
- Investigate new libfuse features est 8 FTW [all]
- Investigation of VARNISH est 8 FTW [all]
- Investigate smarter cache cleanup, skipping open files est 4 FTW [VV]
- Production release of parallel download est 4 FTW [LP]
- Use of libnettle for hashing primitives est 4 FTW [JB]



- Resilient cvmfs_config bugreport when mountpoints unresponsive **est 2 FTW [LP]**
- Improved macOS support **est 6 FTW [VV]**
 - macOS support for FUSE-T, and FSKit evolution
- Extension of CVMFS_INFO_HEADER, including review of capabilities **est 4 FTW [DD]**
- Use shrinkwrap with diffs of named snapshots (requested by CMS) **est 4 FTW [VV]**
- Investigate Rucio-FS-like functionality [[summer student proposal](#)]
- Offline mode for developers (stretch goal)



- MacFUSE library stable, but requires kernel extensions enabled
 - More difficult to do with each update, currently 3 reboots
- FUSE-T is an alternative implementation using an NFS-server
 - Less mature, and less performant, but without installation procedure
 - can be used on github actions shared runners
- GSoC project in 2024 evaluated FUSE-T
 - finding few non-critical peculiarities, fixing one critical bug (in FUSE-T)
 - a major blocker (wrong directory listings) remains
 - but homebrew infrastructure, m1 builds in place
- FSKit will likely shake up filesystem APIs again, requiring further attention



- Partial replication for Stratum-1s (EESSI usecase) **est 4 FTW** [VV]
- Priority queue for lease requests **est 4 FTW** [all]
- Interruptible garbage collection **est 4 FTW** [VV]
- Investigate better package upgrade mechanisms **est 4 FTW** [VV]
- High-availability / multiple gateways **est 4 FTW** [SV, VV]
- UI improvements for cvmfs_server **est 4 FTW** [SV, VV]
- Elimination of legacy bulk hashes in cern.ch repositories **est 4 FTW** [VV]
- Bugfix: CVMFS_AUTO_TAG_TIMESPAN to be set from publishers connected to gateway **est 1 FTW** [VV]
- Investigate ML Model distribution on CVMFS [GSoC student proposal]



- Centos 7 (as soon as all experiments have migrated)
 - Will facilitate move to newer C++ standards, slimmer codebase
- x509 authenticated repositories
 - LIGO, the only user, will stop their usage of authenticated data repos in 2025
- Javascript code in current cvmfs-monitor website
 - Opportunity to remove number of programming languages in the project; rewrite in golang
 - How important is this monitoring?

Community Interaction



- Developers and operators meet in a monthly coordination call (no changes for 2025)
- Regular alignment with IT-ST
- established contacts with experiment representatives in charge of operations

- Interaction with external users and industry (Jump Trading, EESSI, EUCLID, LIGO, etc.)
- Conferences on the radar: CHEP, OTF, XRootD workshop, HEPiX, Kubecon, SC

- Preparation of the CernVM Workshop 2026

Discussion

(and a link to the live document for PoW completion: [Link](#))
