CernVM Program of Work 2020

J Blomer for the CernVM Team

SFT Meeting
20 January 2020

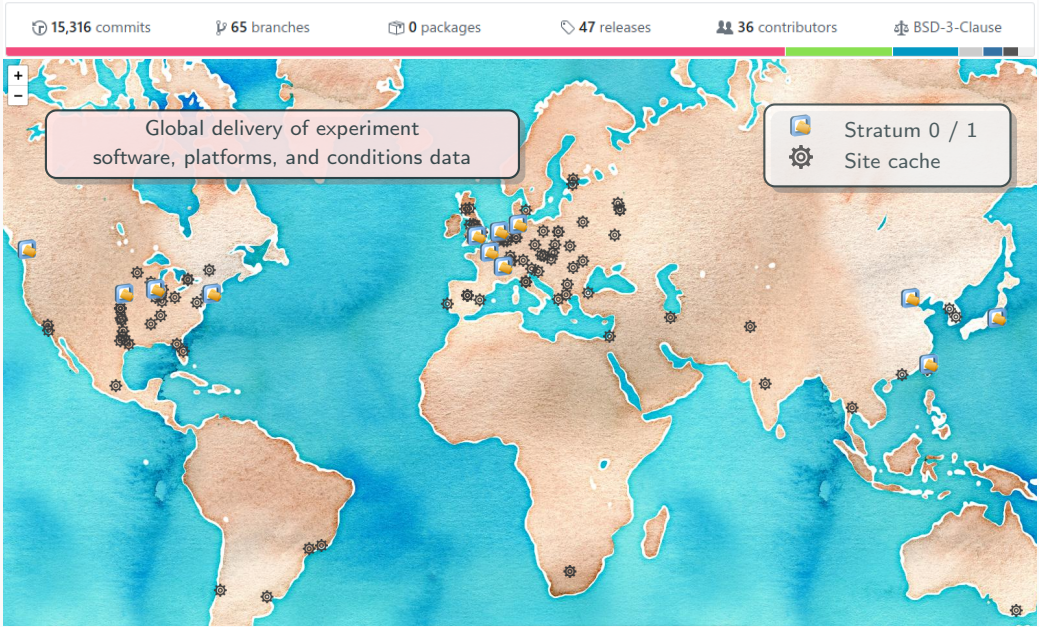
15,316 commits
65 branches
0 packages
47 releases
36 contributors
BSD-3-Clause
Global delivery of experiment
software, platforms, and conditions data
Stratum 0 / 1
Site cache

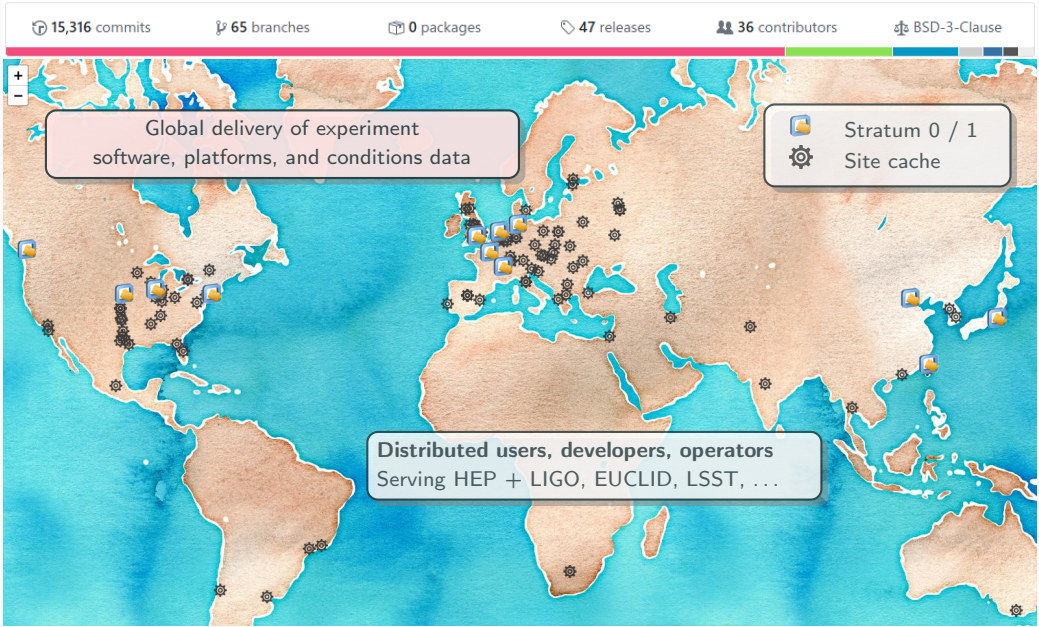
15,316 commits
65 branches
0 packages
47 releases
36 contributors
BSD-3-Clause
Global delivery of experiment
software, platforms, and conditions data
Stratum 0 / 1
Site cache
Distributed users, developers, operators
Serving HEP + LIGO, EUCLID, LSST, ...

Review of 2019

Highlights

- Established `unpacked.cern.ch` as a CernVM-FS powered hub for container images
- Established a set of high-level services in Go
  - Gateway services (rewrite based on first production experience)
  - Container publishing service (DUCC)
  - Notification service for instant content propagation
- Commissioned CERN S3 integration for nightly build use case
- Support for Fuse 3, including root-less client on modern kernels
- R&D: Prototype CernVM-FS-Kubernetes integration
- Users workshop with speakers from Cloudflare, Singularity, Jülich SC, Canonical
- Dissemination: ACAT (poster), CHEP (poster),
  Massive Storage Systems and Technology (MSST, invited talk),
  Ceph Day, KT/HUG workshop on Information Technology for the Hospitals

Review of 2019

Highlights

- Established `unpacked.cern.ch` as a CernVM-FS powered hub for container images
- Established a set of high-level services in the Go ecosystem
  - Gateway services (rewrite based on first production experience)
  - Container publishing service (DUCC)
  - Notification service for instant content propagation
- Commissioned CERN S3 integration for nightly build use case
- ...

Unfinished Tasks

- Consolidation of CernVM services (student project not granted)
- CernVM 5 prototype
- CernVM website migration
- Showcase repository: `sw.hsf.org`
- In progress:
  - New, solid underpinning of the CernVM-FS server tools
  - On-demand container for publishing
  - New cernvm-monitor web site

CernVM / CernVM-FS Program of Work 2020

Developer Power

		2019	2020
Jakob Blomer	Staff	50 %	50 %
Radu Popescu	Fellow	66 %	—
Simone Mosciatti	Fellow (Tech)	66 %	100 %
Jan Priessnitz	Tech	60 %	40 %
TBS	Tech	—	40 %
		2.42+	2.3+

Significant contributors:
Shahnur Isgandarli (openlab summer student), Marcin Mokrzan (CERN summer student), Enrico Bocchi (IT-ST), Dan van der Ster (IT-ST), Dave Dykstra (FNAL)

CernVM Calendar

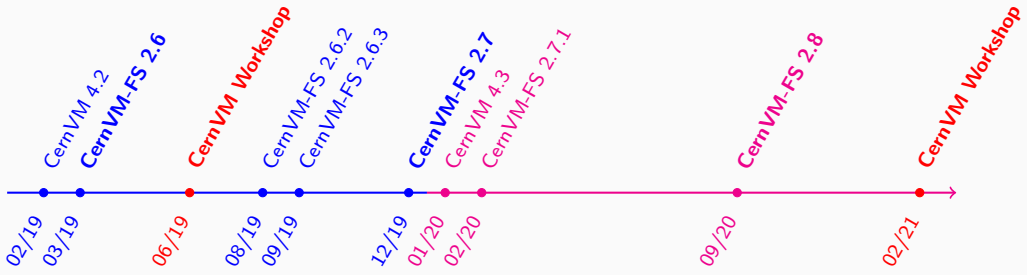

CernVM-FS integration with container engines (singularity, containerd, podman, . . .)
takes place in parallel to releases

CernVM Plan of Work for 2020

1. Ready to use platform for LHC experiment **production** and **development**
2. Reference platform for **long-term data preservation**

- 20 000+ VMs / day
- μCernVM bootloader + reference containers covering EL 4–7
- Interactive support: `cernvm-launch` and `cernvm-online.cern.ch`

2020 Plan of Work

- CernVM 4.3 following EL 7.7
- Freeze CernVM 3 together with RHEL 6 end-of-life
- In collaboration with SPI and the experiments, investigate the future of VM technology and the strategy for a common, containerized platform

CernVM-FS Plan of Work for 2020

1. Maintenance and support
2. Container images in CernVM-FS
3. General improvements
4. Infrastructure modernization
5. Community interaction

Maintenance and Support

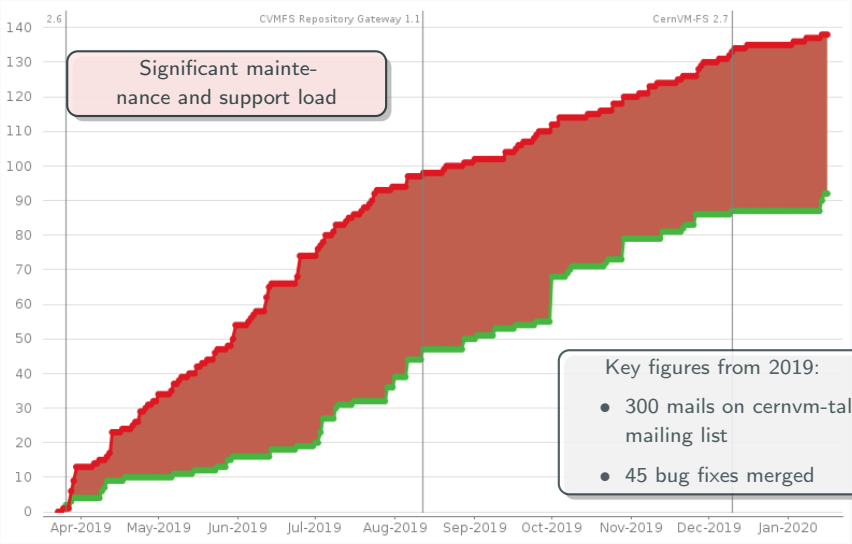

Container Images in CernVM-FS

❶ CernVM-FS in containers

- Bind mount:

```
docker run -v /cvmfs:/cvmfs:shared ...
singularity exec -B /cvmfs ...
```

- CSI driver ▸ Github repository
  "behind the scenes" bind mount, integrates with kubernetes (maintained by IT)
- Mount inside container,
  currently *no* sharing of the cache (see next slides), requires privileged container or modern kernel and tools

Some issues around autofs peculiarities

❷ Container images in CernVM-FS

Unpacked images on /cvmfs in order to benefit from de-duplication and on-demand caching

Requires:

1. DUCC: container unpacker and publisher
   → unpacked.cern.ch
2. Storage plug-in required for layer based engines
   - Docker driver available
   - Prototype driver for containerd/kubernetes available
   - In touch with podman developers (GSoC project)

Container Images in CernVM-FS – Plan of Work

Container Publishing

- Scale test: benchmark a repository with $\mathcal{O}(1000)$ images [est 1 FTM]
- Image grafting: fast merging of existing layers into new images [est 2.5 FTM]
- Dashboard: display current activity, list of hosted images, etc. [est 1 FTM]
- Wild-card image specification [est 2 FTW]
- Multi-arch images [est 2 FTW]
- Drop singularity dependency for image flattening [est 2 FTW]
- R&D: standardize unpacked file system layout
- R&D: investigate image sanity checks before publication

Container Engine Integration

Engine	Type	CernVM-FS Support
Singularity	flat	**native**
runc	flat	**native**
docker	layers	**graph driver** [est 2 FTW]
containerd	layers	**prototype** [est 1 FTM]
podman	layers	**GSoC** [est 1 FTM]

Documentation, examples, integration tests for different deployment options [est 1 FTM]

Next-Generation Server Code

Current state

A set of tools targeted for a dedicated release manager machine, and the interactive workflow *open transaction* + *copy* + *commit*

Where we are going

A common base library providing repository transformation primitives, on top of which higher-level publish abstractions can be built

Initial `libcvmfs_server` merged, first commands transparently ported to new architecture.
Est 2 FTM for complete migration of server command line interface.

General Improvements

- Parallel garbage collection; first tests suggest a 2× speed improvement
  mostly ready [est 2 FTW]
- Repository dashboard: web-accessible, auto-generated ROOT plots with publish statistics
  ongoing [est 1 FTM]
- Shared, external cache manager for multi-container host
  ongoing [est 1 FTM]
- Deployment of the new cernvm-monitor.cern.ch
  finishing off 2019 student project [est 2 FTW]
- R&D: investigate pre-caching to improve interactive use, collaboration with SWAN
  planned as summer student project [est 3 FTM]

Stretch Goal: CernVM-FS Conveyor

A high-level abstraction of writing based on interdependent publication jobs.

```
$ ssh cvmfs-sft.cern.ch
$ cvmfs_server transaction sft.cern.ch /lcg/ROOT
$ tar -xf ROOT-6.18.tar.gz
$ post-install.sh
$ cvmfs_server publish
```

Current approach

```
{
  "repository": "sft.cern.ch",
  "path": "/lcg/ROOT",
  "payload": "https://root.cern.ch/ROOT-6.18.tar.gz",
  "script": "https://spi.cern.ch/post-install.sh",
  "uuid": "e7b67a2...",
  "dependencies": ["f61d...", "a00e...", "..."]
}
```

- Send jobs to Conveyor API
- Conveyor distributes work to multiple publisher nodes

Goal: liberate CI pipeline from handling `cvmfs_server` intrinsics.
Prototype available, est 1–2 months to develop into a first usable version in collaboration with SPI

Infrastructure Modernization

- Migration of ∼15 build jobs to new Jenkins server [est 1 FTM]
- Migrate CernVM web site from Drupal to static page [est 1 FTM]
- Add pull request builder [est 1 FTW]
- Add SLES15 as a supported platform (HPC support) [est 1 FTW]
- Migrate yum and apt package repositories to S3 [est 1 FTW]
- Stretch goal: migrate CernVM central storage to S3 [est 2 FTW]

Community Interaction

- Developers and operators meet in a monthly coordination call (no changes for 2020)
- Weekly operations coffee with IT-SM (no changes for 2020)
- Conferences and workshops on the radar:
  experiment computing weeks, GDB, HEPIX, HSF-WLCG workshop, ACAT
- Modernize the web presence under the `cvmfs.io` domain

- **CernVM 2021 Workshop Preparation**

  Place: NIKHEF
  Dates: 1–3 Feb 2021

Outlook and Goals for 2020

- Improve workflows for containers in CernVM-FS
- Continue modernization of server code for more flexible publishing

The container ecosystem matures with the advent of root-less containers and root-less fuse file systems (EL8). For proper CernVM-FS integration, we need to work with external developers in order to find a viable path through the maze of container tools and standards.