



CernVM Program of Work 2019

J Blomer for the CernVM Team

SFT Meeting

21 January 2019

14,362 commits

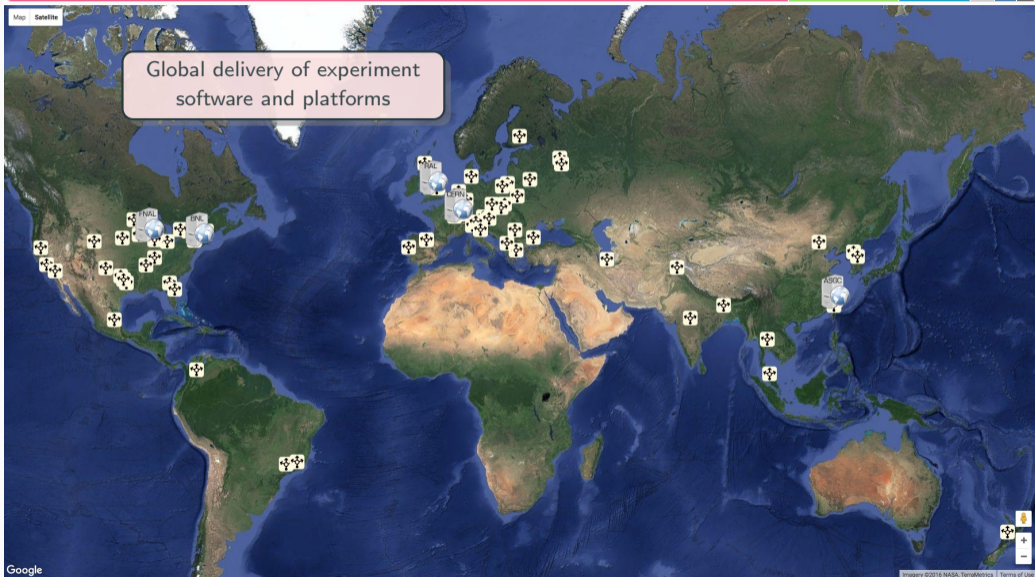
42 branches

39 releases

35 contributors

BSD-3-Clause

Global delivery of experiment software and platforms



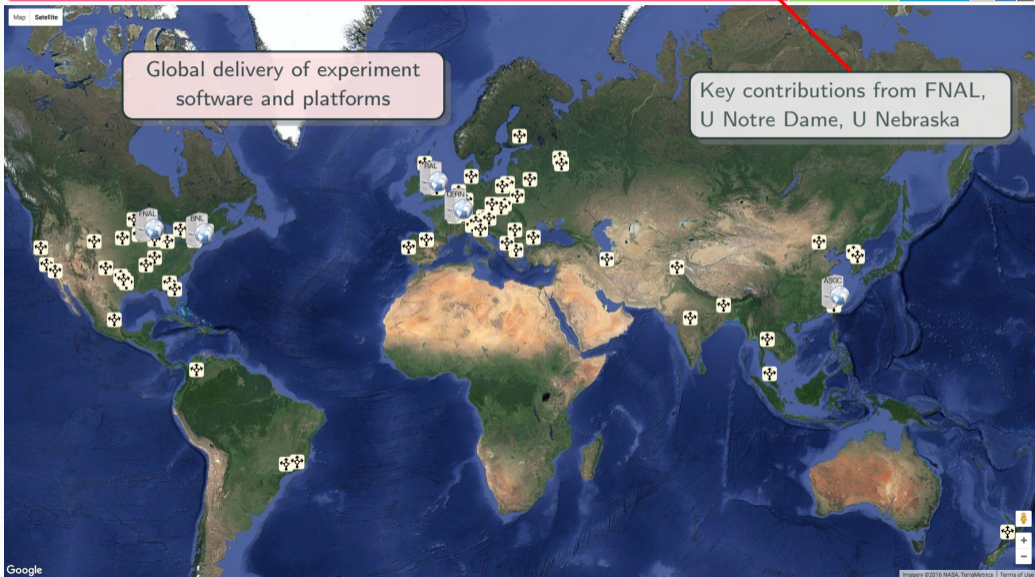
14,362 commits

42 branches

39 releases

35 contributors

BSD-3-Clause



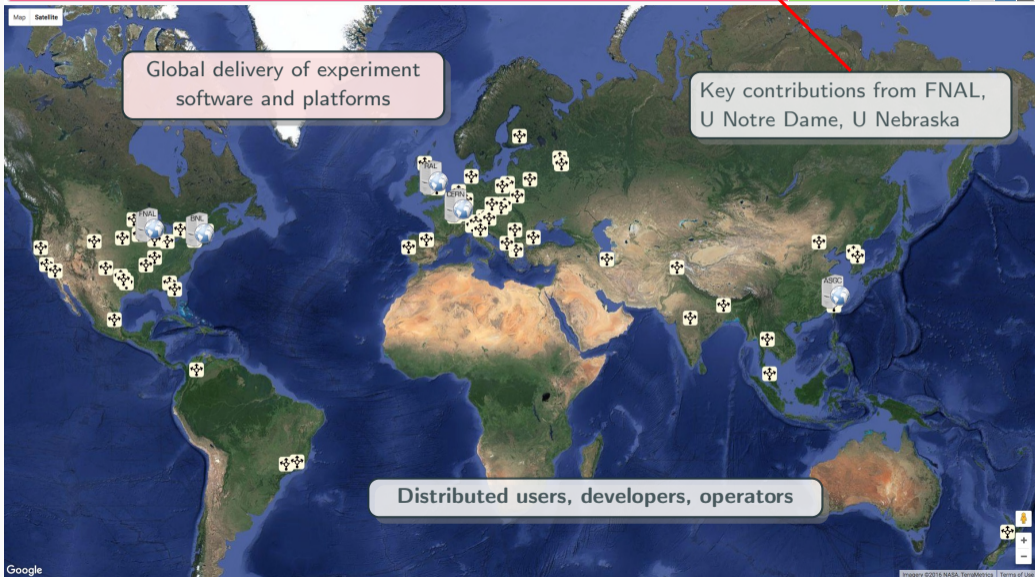
14,362 commits

42 branches

39 releases

35 contributors

BSD-3-Clause



Global delivery of experiment software and platforms

Key contributions from FNAL, U Notre Dame, U Nebraska

Distributed users, developers, operators



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers: new cvmfs_shrinkwrap utility
- Introspection for repository content responsables: publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona), IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers:
new `cvmfs_shrinkwrap` utility
- Introspection for repository content responsables:
publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona),
IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers:
new `cvmfs_shrinkwrap` utility
- Introspection for repository content responsables:
publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona), IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers:
new `cvmfs_shrinkwrap` utility
- Introspection for repository content responsables:
publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona),
IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers:
new `cvmfs_shrinkwrap` utility
- Introspection for repository content responsables:
publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona), IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers: new cvmfs_shrinkwrap utility
- Introspection for repository content responsables: publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona), IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers: new cvmfs_shrinkwrap utility
- Introspection for repository content responsables: publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona), IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers: new cvmfs_shrinkwrap utility
- Introspection for repository content responsables: publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona), IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM-FS gateway services
- Commissioning CernVM 4, cernvm-prod.cern.ch 1st user of the gateway services
- Established CernVM-FS as one of the first users of the CERN S3 storage service
- Projecting partial /cvmfs trees in HPC containers: new cvmfs_shrinkwrap utility
- Introspection for repository content responsables: publish and garbage collection metrics
- Integration test suite ported to macOS
- R&D: CernVM-FS – XRootD integration, WebAssembly client, containerd integration
- Users workshop with speakers from Red Hat, Microsoft, Docker, Singularity
- Dissemination: CHEP (plenary), HPC Knowledge Meeting (Barcelona), IEEE eScience (Amsterdam), ALICE T1-T2 Workshop (Derby)



Highlights

- Commissioning of the CernVM gateway services
- Commissioning CernVM 4, `cernvm-prod.cern.ch` 1st user of the gateway services
- Projecting partial `/cvmfs` trees in HPC containers:
new `cvmfs_shrinkwrap` utility
- Introspection for repository content responsables:
publish and garbage collection metrics
- Integration test suite ported to macOS
- ...

Unfinished Tasks

- Refurbishment of the CernVM-FS server tools
- Consolidation of web services:
`cernvm-online.cern.ch`,
`cernvm-monitor.cern.ch`
- In progress:
 - publishing through portals
 - fast content propagation assisted by a notification service

CernVM / CernVM-FS Program of Work 2019



Ready to use platform (VM, container) for HEP experiment software
production and **development** use case

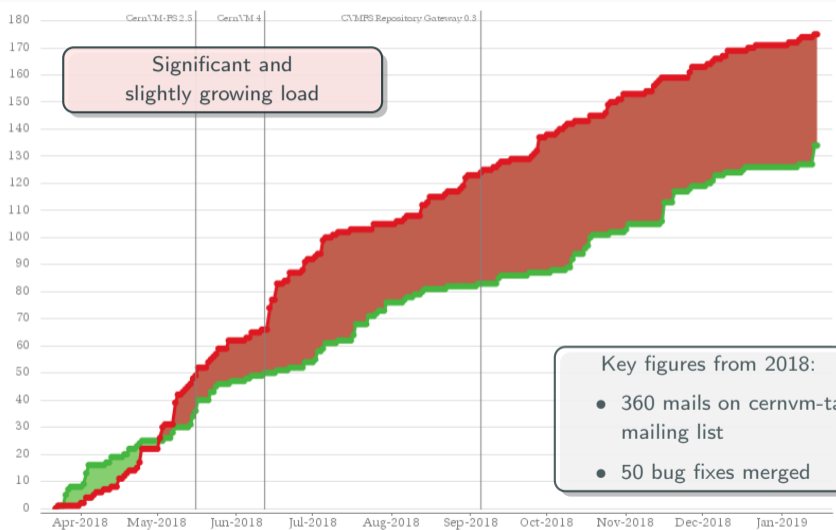
- 12 000+ VMs / day
- 300+ *interactive* users
- EL6 and EL7 compatible versions supported

2019 Plan of Work

- CernVM 4.2 following EL 7.6
- Define a roadmap for deprecation of CernVM 3
- Consolidation and continuous delivery: minimize maintenance costs
(planned as Portuguese trainee project)
- R&D: EL 8 based CernVM 5 proof-of-concept following the RHEL 8 production release



1. Maintenance and support
2. Reference deployments showcasing the latest developments
3. Code works
4. Releases
5. Community interaction





Idea: showcase latest developments,
learn for the adoption of new features into production environment

- sw.hsf.org, sw-nightlies.hsf.org

Send publish jobs from Jenkins to CernVM-FS (CernVM-FS Queue Service), use multiple concurrent release manager machines (Gateway Service) to publish into S3 storage (S3 backend) and monitor the change sets over time (Publisher Metrics)

- unpacked.cern.ch (in collaboration with WLCG container working group)


Publish unpacked container images (Tarball Ingestion) automatically based on a declaration of interest of images (Docker Registry Connector) and use them with Singularity, Docker, or containerd (Docker graph driver)

- Conditions data repository (depending on experiment interest):

Start validation jobs immediately after publishing conditions data (Notification Service)



- Revisit and refurbish publisher tools
 - Drop requirement of exclusively owning a release manager machine, instead: use ephemeral publish container
 - Main blocker in further scaling up write loads
 - Future-proof client:
investigate implications of root-less FUSE and latest Linux kernel namespaces features
-
- Incremental (faster) garbage collection (*approved technical student project*)
 - cernvm-online.cern.ch: transition to client-side monitoring based on JavaScript client (*planned as a summer student project*)
 - Deep performance analysis using modern Linux tools (eBPF) (*planned as a openlab student project*)
 - Proof-of-concept Windows client (*planned as a GSoC project*)



Features for 2.6

- Shrinkwrap utility for HPC
- Publish metrics
- Direct tarball ingestion
- Container publishing service (*experimental*)
- Notification service (*experimental*)

Features for 2.7

- Renovated `cvmfs_server` tool suite
- Publish in ephemeral container (*experimental*)



- Developers and operators meet in a monthly coordination call (no changes for 2019)
- Plan to prototype a more modern web presence under the `cvmfs.io` domain
- Plan to suggest a ~ 2 hours training for repository responsables

- **CernVM Workshop 2019**

Dates: June 3 – 6

Themes: Serverless (publishing), HPC integration, container integration

Ideas for talks: CloudFlare Workers, Ceph-FS, federated Ceph S3



- Streamline CernVM operations
- Showcase latest developments in pre-production environments

Linux and cloud system fundamentals changed over the last years
(containers, namespaces, overlays, root-less FUSE, ...).

It's time to revisit some central parts of the CernVM-FS server and client code
to make it future-proof and to lower future maintenance load.