



## Using CVMFS to distribute LCG Releases

---

Johannes Heinz

CERN, EP-SFT-SPI

CernVM Workshop 2019

03/06/2019

## Introduction to LCG Releases

---

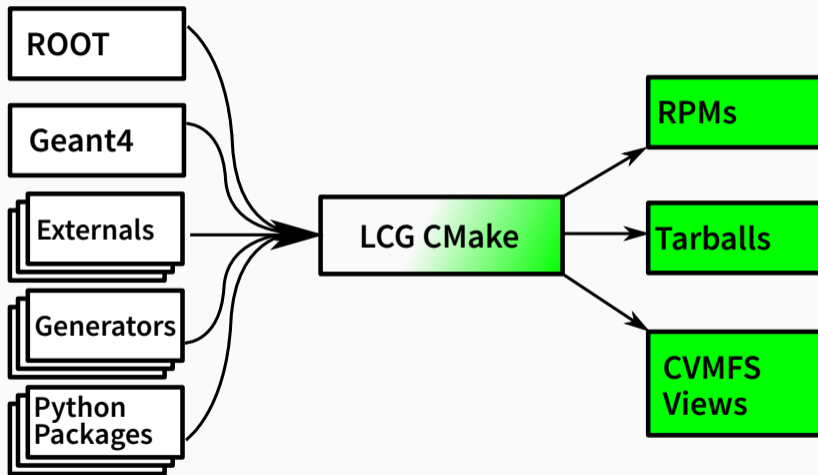


## SPI group mandate:

- Service group for experiments
- Solve common tasks only once
- Provide easy access to commonly used programs
- Simplify setup of work environments

## Clients:

- ATLAS
- LHCb
- SWAN
- IT-DB-SAS
- BE department





Use the view of the latest release on `lxplus.cern.ch`

```
1 source /cvmfs/sft.cern.ch/lcg/views/LCG_95a/  
2 x86_64-centos7-gcc8-opt/setup.sh
```

Look up releases, packages and platforms: [LCG Info](#)



## Configure Environment

Software stack [more...](#)

95a

Platform [more...](#)

CentOS 7 (gcc7)

Environment script [more...](#)

e.g. \$CERNBOX\_HOME/MySWAN/myscript.sh

Number of cores [more...](#)

2

Memory [more...](#)

8 GB

Spark cluster [more...](#)

None



- **Architecture:** x86\_64
- **Operating system:** Scientific Linux 6, CentOS 7, Ubuntu 16.04 LTS, Ubuntu 18.04 LTS
- **Compiler:**
  - **GCC:** native, 6.2.0, 7.3.0, 8.2.0
  - **Clang:** 7.0.0, 8.0.0
- **Flags:** sse3, avx2, fma, CUDA support
- **Build type:** Debug, Optimized
- **Python version:** 2.7, 3.6
- **ROOT version:** master, 6.18
- **Toolchains:** BE (Beams), Experiments




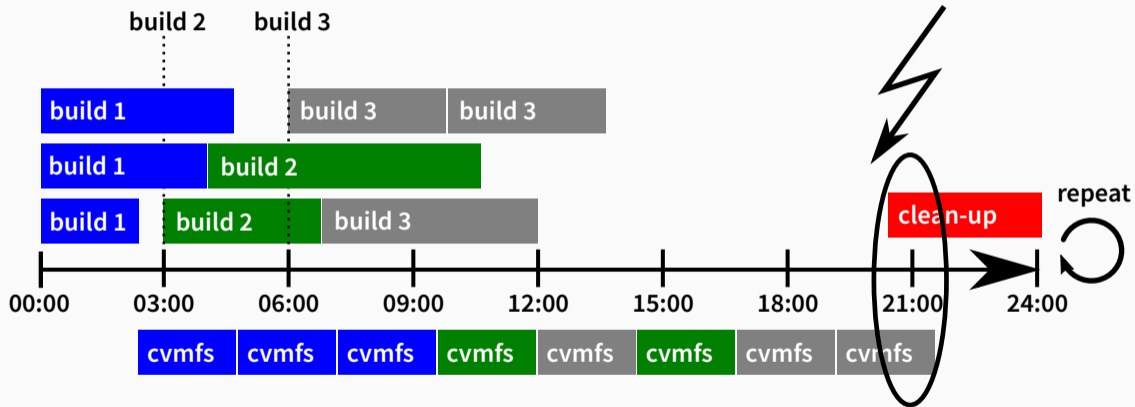
- **Architecture:** x86\_64, aarch64 (ARM), PowerPC
- **Operating system:** Scientific Linux 6, CentOS 7, Ubuntu 16.04 LTS, Ubuntu 18.04 LTS, **Fedora ?**
- **Compiler:**
  - **GCC:** native, 6.2.0, 7.3.0, 8.2.0, **9.1.0**
  - **Clang:** 7.0.0, 8.0.0
- **Flags:** sse3, avx2, fma, CUDA support
- **Build type:** Debug, Optimized
- **Python version:** 2.7, 3.6
- **ROOT version:** master, 6.18
- **Toolchains:** BE (Beams), **ATLAS**, **LHCb**, ...





- Possible combinations per night: 2 304
- Actual platforms per night: 52 (37 published to CVMFS)
- Packages per Platform: Typically more than 390
- Repositories:
  - `sft.cern.ch` – CVMFS 2.6.0, 67 086 catalogs
  - `sft-nightlies.cern.ch` – CVMFS 2.6.0, 19 299 catalogs
- Available releases: From **LCG\_71** to **LCG\_95a** (plus nightly builds)





## New approaches

---



## Current approach:

- Puppetized VM builds for releases, Docker images for nightly builds
- Bind mount CVMFS from host system (CERN CentOS 7)
- Irregular problems while publishing to EOS as staging area
- Active waiting for CVMFS transaction (blocking slots in Jenkins)
- Assumptions about stratum 1 state

## New approach:

- Move to a unified approach for all builds  $\Rightarrow$  **Kubernetes orchestration**
- Use CSI driver for read access
- Test CERN S3 as staging area
- Asynchronous deployment using conveyor
- Use notification about nightly builds to trigger LCG Info update and tests



- Reference use case for new CVMFS features
- Future **HSF** (HEP Software Foundation) repositories:
  - `sw.hsf.org` – Future releases repository
  - `sw-nightlies.hsf.org` – Future nightly repository
  - `test-sw-nightlies.hsf.org` – Experimental repository for testing



```
1 /cvmfs/sft-nightlies.cern.ch/lcg/nightlies/ # base (1)
2   dev4/ # toolchain (5)
3   Mon/ # day (1)
4   Python/ # package (390)
5   2.7.15/ # version (1+)
6   x86_64-slc6-gcc8-opt/ # platform (19)
```

⇒ Take lease at platform level



- Publish each package separately
- Collect publication IDs

```
1     id=$(conveyor submit
2         --repo test-sw-nightlies.hsf.org
3         --lease-path ${PACKAGE_VERSION_PLATFORM}
4         --payload "script |
5             https://${S3_BUCKET}/${LCG_TRANSACTION_SCRIPT}
6                 ?checksum=sha1:${SHA1_HASH}|
7             https://${S3_BUCKET}/${$package}.tgz"
8         | tail -1
9         | jq -r '.job_id')
10    ids="$ids $id"
```





- Create view with a dependency to each package

```
1 conveyor submit
2   --repo test-sw-nightlies.hsf.org
3   --deps "$ids"
4   --wait
```

# Read from CVMFS within Kubernetes

---



Gain read access to CVMFS within Kubernetes using the CSI driver:

- CernVM-FS CSI driver
- CERN IT: Cloud Docs
- Kubernetes documentation: Run a job to completion

At CERN we can take advantage of the pre-configured Kubernetes plugin and the support by IT



```
1  openstack coe cluster create my-kubernetes-cluster \  
2      --cluster-template kubernetes-1.13.3-1 \  
3      --master-flavor m2.medium \  
4      --node-count 3 --flavor m2.xlarge \  
5      ...  
6      --labels cvmfs_csi_enabled=True \  
7      --labels cvmfs_csi_version=v0.3.0 \  
8      --labels cvmfs_tag=qa \  
9      ...
```



```
1 apiVersion: storage.k8s.io/v1
2 kind: StorageClass
3 metadata:
4   name: csi-cvmfs-sft
5 provisioner: csi-cvmfsplugin
6 parameters:
7   repository: sft.cern.ch
```



```
1 apiVersion: v1
2 kind: PersistentVolumeClaim
3 metadata:
4   name: csi-cvmfs-sft-pvc
5 spec:
6   accessModes:
7   - ReadOnlyMany
8   resources:
9     requests:
10      storage: 1Gi
11   storageClassName: csi-cvmfs-sft
```



```
1 apiVersion: batch/v1
2 kind: Job
3 metadata:
4   name: cvmfstest
5 spec:
6   template:
7     # (on next slide)
8   backoffLimit: 4
```



```
1     containers:
2     - name: cvmfs-container
3       image: cern/cc7-base
4       command: ["ls", "-halt", "/cvmfs/sft.cern.ch/lcg"]
5       volumeMounts:
6       - mountPath: /cvmfs/sft.cern.ch
7         name: sft-pvc
8 restartPolicy: Never
9 volumes:
10  - name: sft-pvc
11    persistentVolumeClaim:
12      claimName: csi-cvmfs-sft-pvc
13      readOnly: true
```



## Execute the job and get results:



### Send job configuration to Kubernetes:

```
kubectl apply -f ./job.yaml
```

### Get pod that runs the job:

```
kubectl describe jobs/cvmfstest
```

### Get the output from the job:

```
kubectl logs cvmfstest-mhvng
```



## Challenges:

- Growing number of supported platforms (new compilers, new OSs, ARM support)
- Diversification of experiments' requirements
- Sequential publication to CVMFS too slow
- Build environment consistency hard to control (puppetized VMs vs. Docker images)

## Approaches:

- New toolchains for each experiment (ATLAS, LHCb)
- Use multiple release machines with clear defined lease paths
- Use exclusively Docker to define the build environment
- Use Kubernetes to make better use of our build resources