
# Software Distribution with CernVM-FS

Jakob Blomer  CERN, EP-SFT

jblomer@cern.ch

https://github.com/cvmfs/cvmfs

2019-05-17

# 1. Accelerate. . .

# 2. Collide. . .

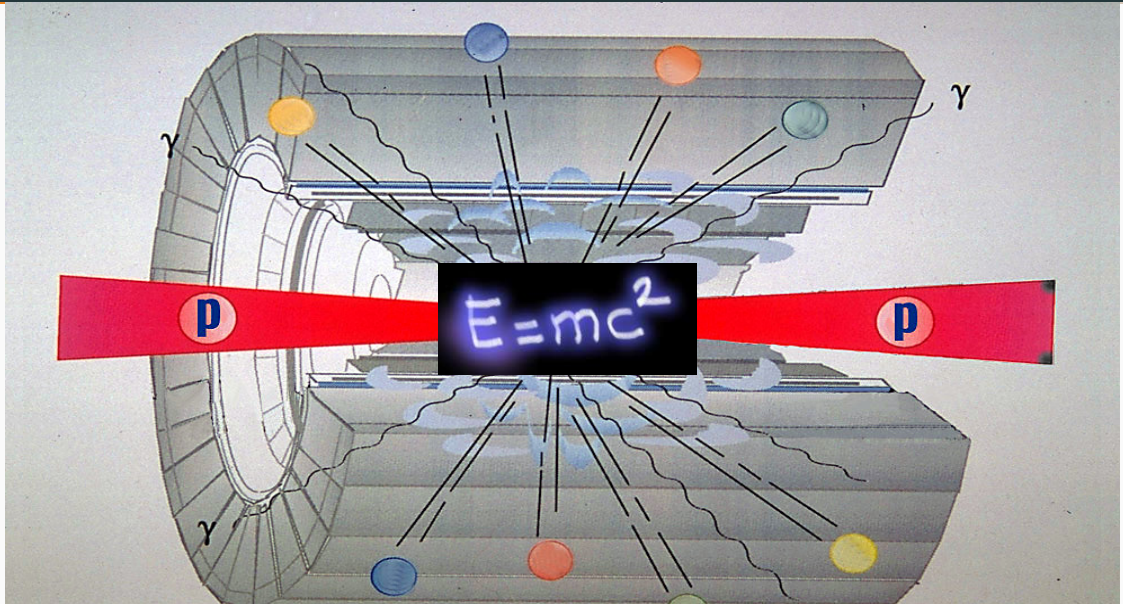

# 3. Measure. . .

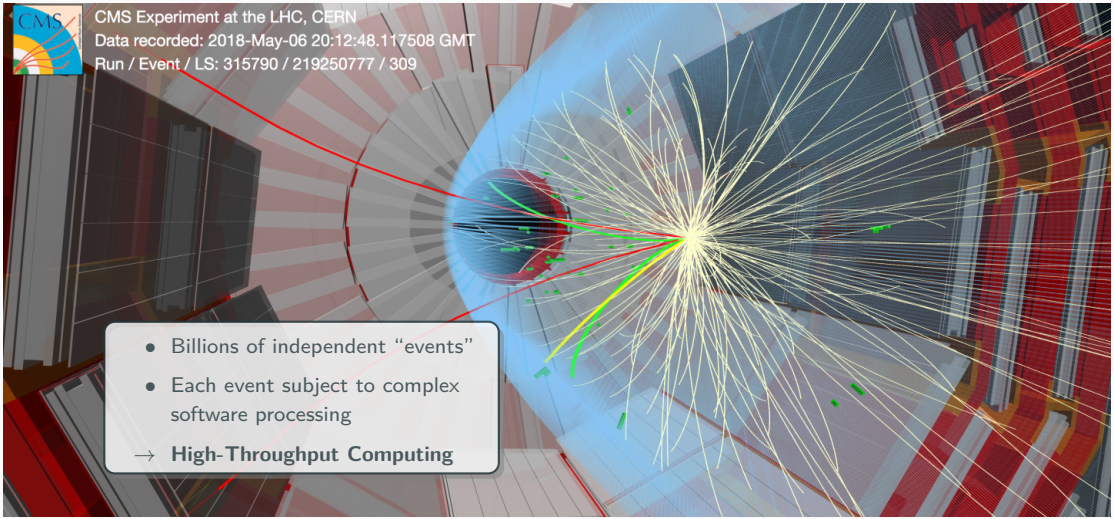


- Billions of independent “events”
- Each event subject to complex software processing

→ **High-Throughput Computing**

# 4. Analyze!

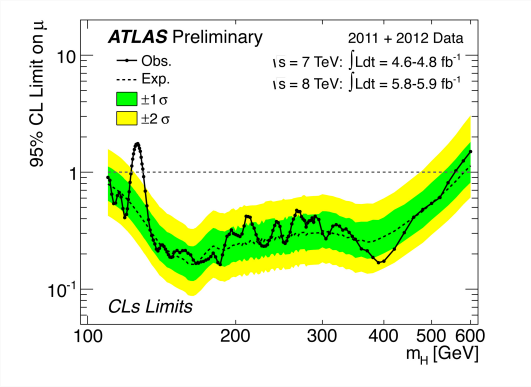
ATLAS Preliminary
2011 + 2012 Data
Obs.
Exp.
±1σ
±2σ
√s = 7 TeV: ∫Ldt = 4.6-4.8 fb-1
√s = 8 TeV: ∫Ldt = 5.8-5.9 fb-1
95% CL Limit on μ
CLs Limits
mH [GeV]

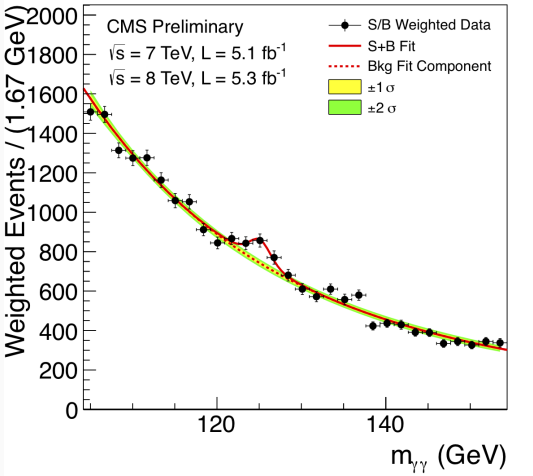
CMS Preliminary
√s = 7 TeV, L = 5.1 fb-1
√s = 8 TeV, L = 5.3 fb-1
S/B Weighted Data
S+B Fit
Bkg Fit Component
±1σ
±2σ
Weighted Events / (1.67 GeV)
mγγ (GeV)

# Computing, Federated

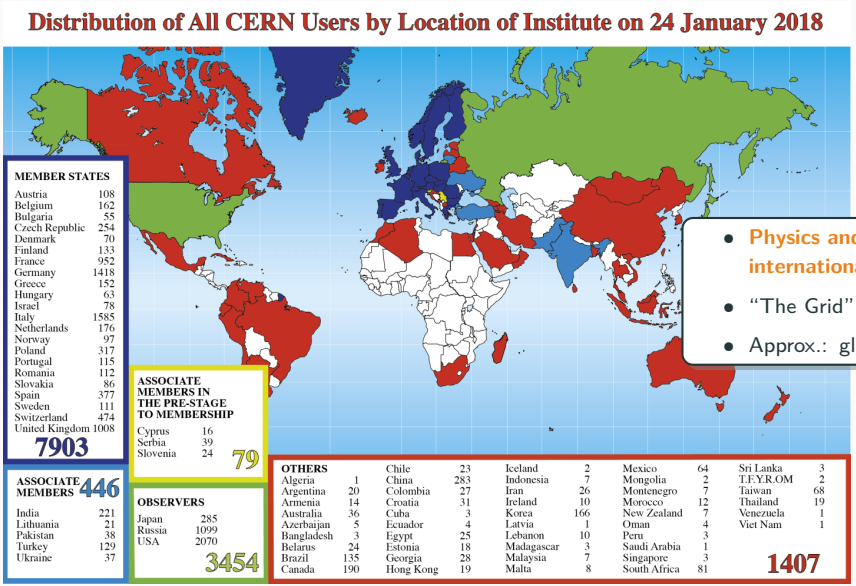


- **Physics and computing: international collaborations**
- "The Grid": ∼160 data centers
- Approx.: global batch system

# Software Delivery by a File System

❶ Provide uniform, consistent, and versioned POSIX file system access to `/cvmfs`

```
$ ls /cvmfs/cms.cern.ch
slc7_amd64_gcc700  slc7_ppc64le_gcc530  slc7_aarch64_gcc700  slc6_mic_gcc481
...
```

on **grids**, **clouds**, **supercomputers** and **end user laptops**

*read*

---

*publish*

❷ Populate and propagate new and updated content

- A few "software librarians" can publish into /cvmfs
- Transactional writes as in `git commit/push`

# Scale of Deployment

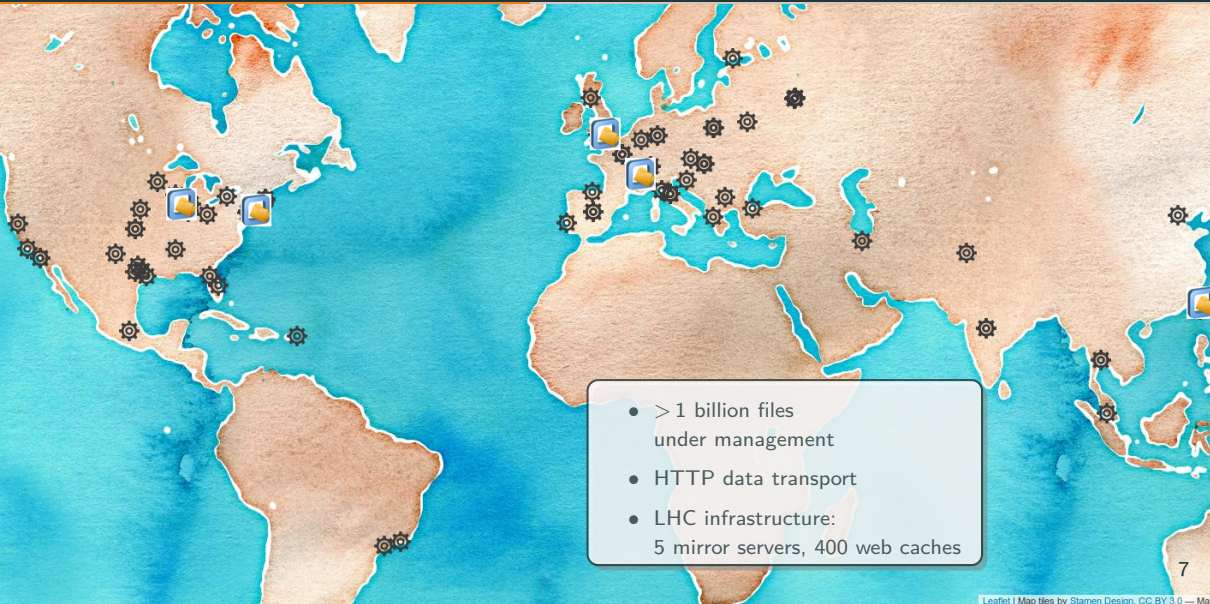


- $>1$ billion files under management
- HTTP data transport
- LHC infrastructure: 5 mirror servers, 400 web caches

# Key Design Choices

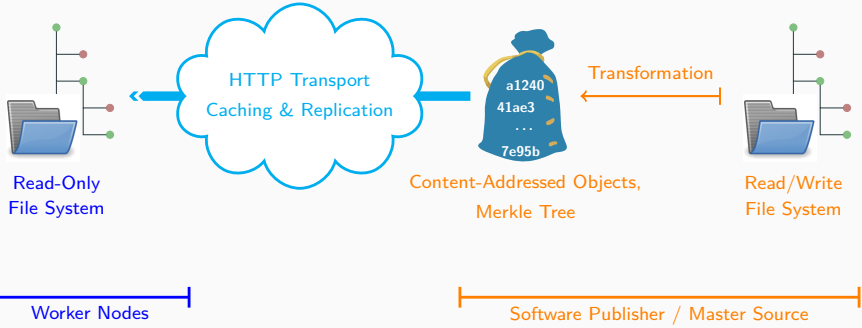


- Reading and writing treated asymmetrically
- Immutable objects, stateless services
- HTTP transport + caching
- Consistency over availability

# Key Design Choices

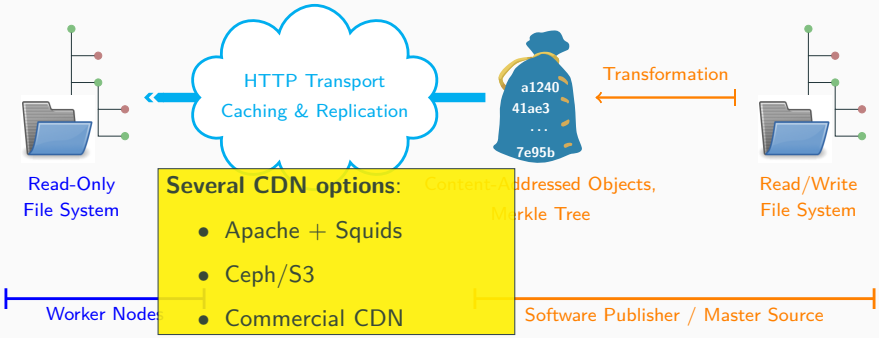


- Reading and writing treated asymmetrically
- Immutable objects, stateless services
- HTTP transport + caching
- Consistency over availability

# Reading

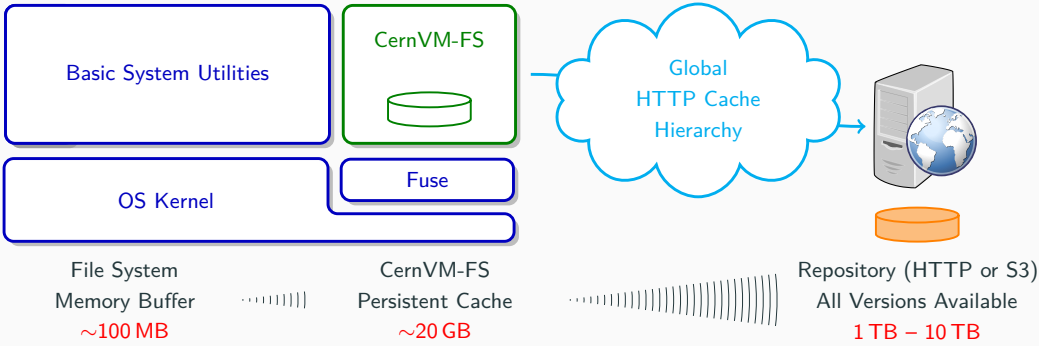


- Fuse based, independent mount points, e. g. /cvmfs/atlas.cern.ch
- High cache effiency because entire cluster likely to use same software

# Writing

- Kernel-level union file system:
  AUFS, OverlayFS

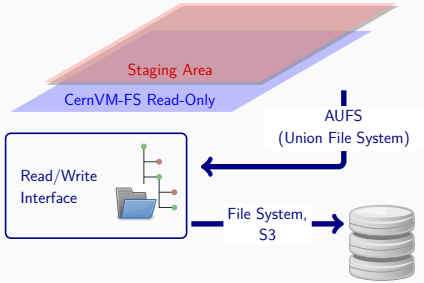


**Publishing new content**

```
[ ~ ]# cvmfs_server transaction containers.cern.ch
[ ~ ]# cd /cvmfs/containers.cern.ch && tar xvf ubuntu1610.tar.gz
[ ~ ]# cvmfs_server publish containers.cern.ch
```

# Use of Content-Addressable Storage

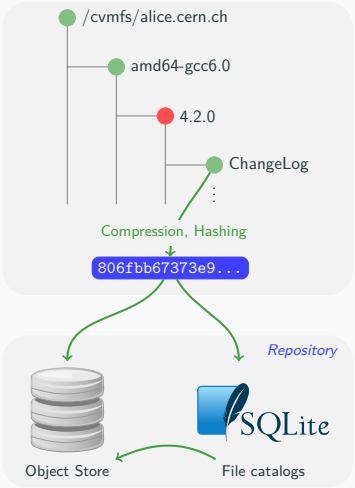


**Object Store**

- Compressed files and chunks
- De-duplicated

**File Catalog**

- Directory structure, symlinks
- Content hashes of regular files
- Digitally signed
  $\Rightarrow$ integrity, authenticity
- Time to live
- Partitioned / Merkle hashes
  (possibility of sub catalogs)

$\Rightarrow$ Immutable files, trivial to check for corruption, versioning

# Why a file system?

# HEP Software Challenge

```
$ cmsRun DiPhoton_Analysis.py
```

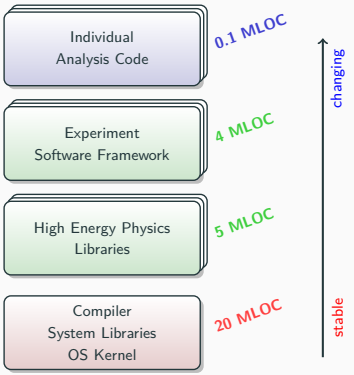


**Key Figures**

- Hundreds of (novice) developers
- Hundred million binaries
- 1 TB / day of nightly builds
- ∼100 000 machines world-wide
- Daily production releases, remain available “eternally”

→ **too much for a packaging approach**

# Actual new content

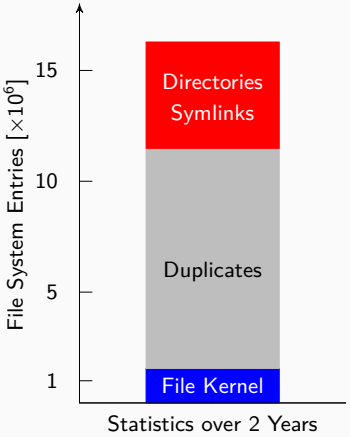


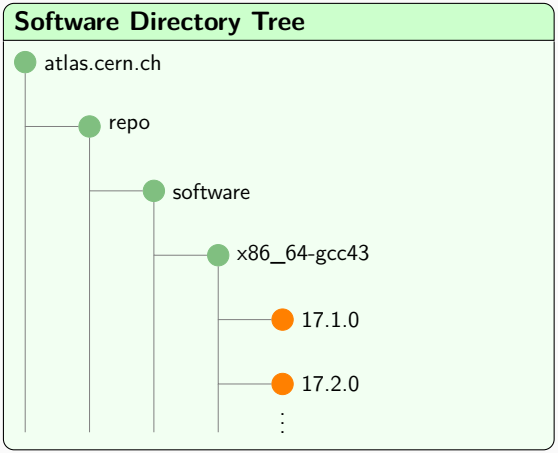


Between consecutive software versions: only ∼15 % new files

At runtime only tiny fraction of files actually accessed

# Container Amplification

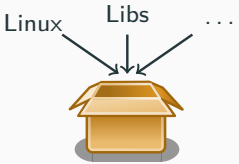


- Containers are easier to create than to role-out at scale
- Ideally: containers for isolation and a software file system for distribution

# Custom Graph Driver / Snapshotter Approach

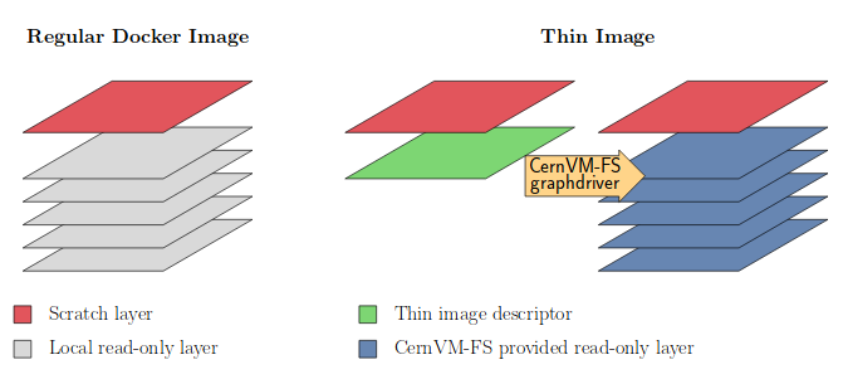


Ideally: native support for (some) unpacked layers on a read-only file sytem