



CernVM and the CernVM File System

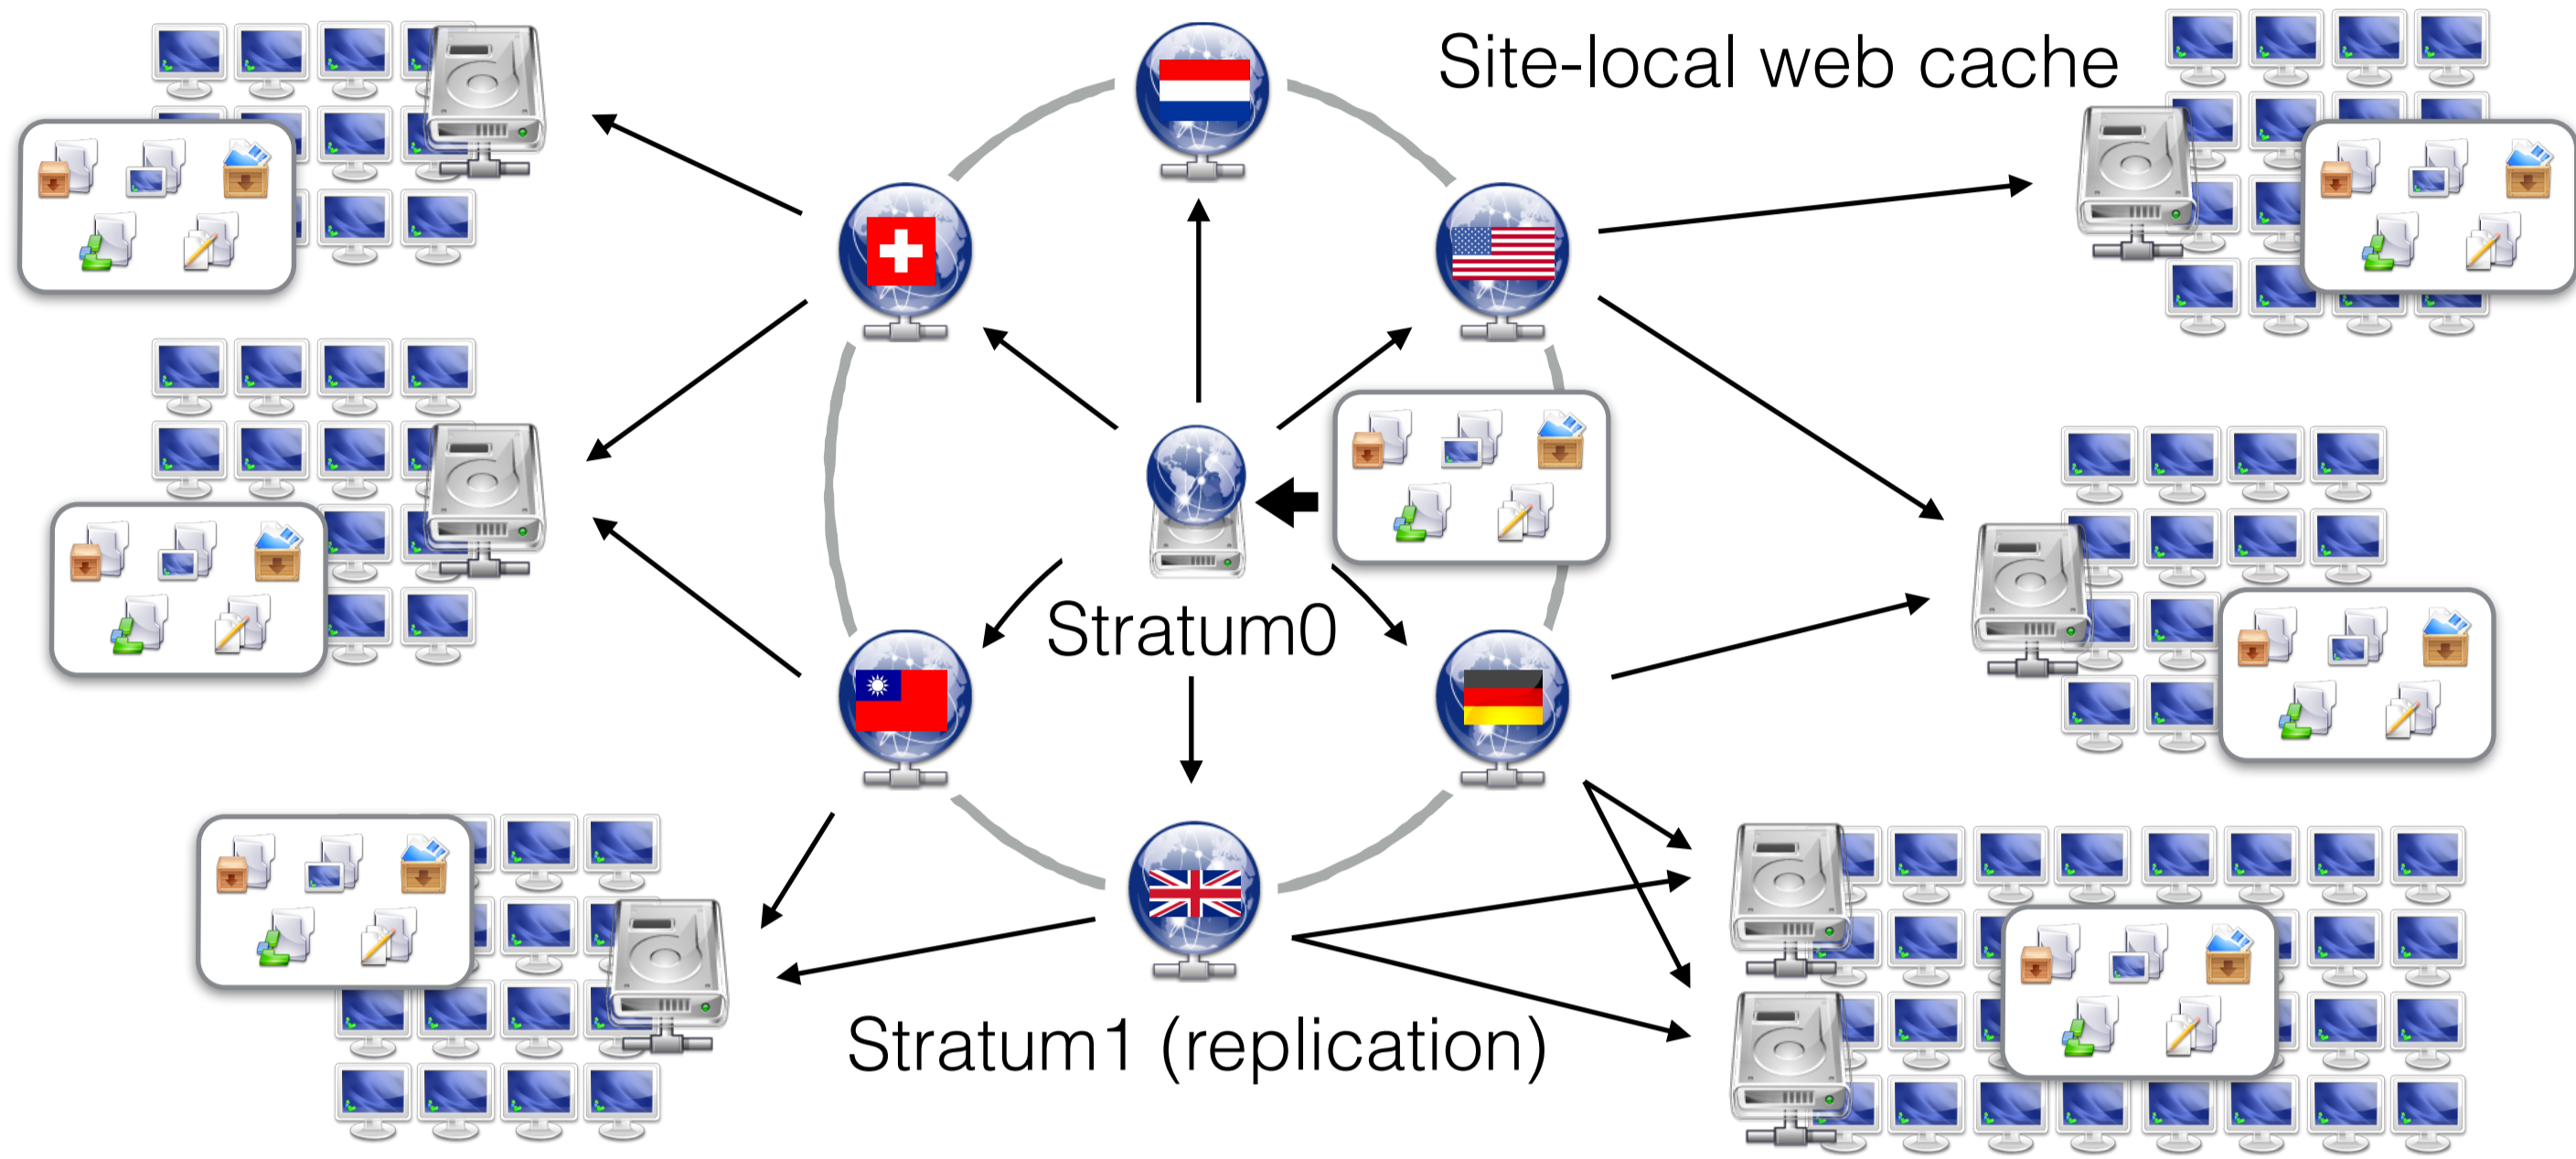
Current Status and Future Roadmap

CernVM-File System

- Scalable and reliable software distribution
- Using HTTP for transport
- Globally distributed full replication servers
- Aggressively cached on multiple layers

CernVM-FS Usage Statistics

Repository	Files	Volume
cms.cern.ch	40 M	1.2 TB
alice.cern.ch	9 M	0.8 TB
lhcb.cern.ch	14 M	0.6 TB
atlas.cern.ch	40 M	2.5 TB
lhcbdev.cern.ch	126 M	6.5 TB

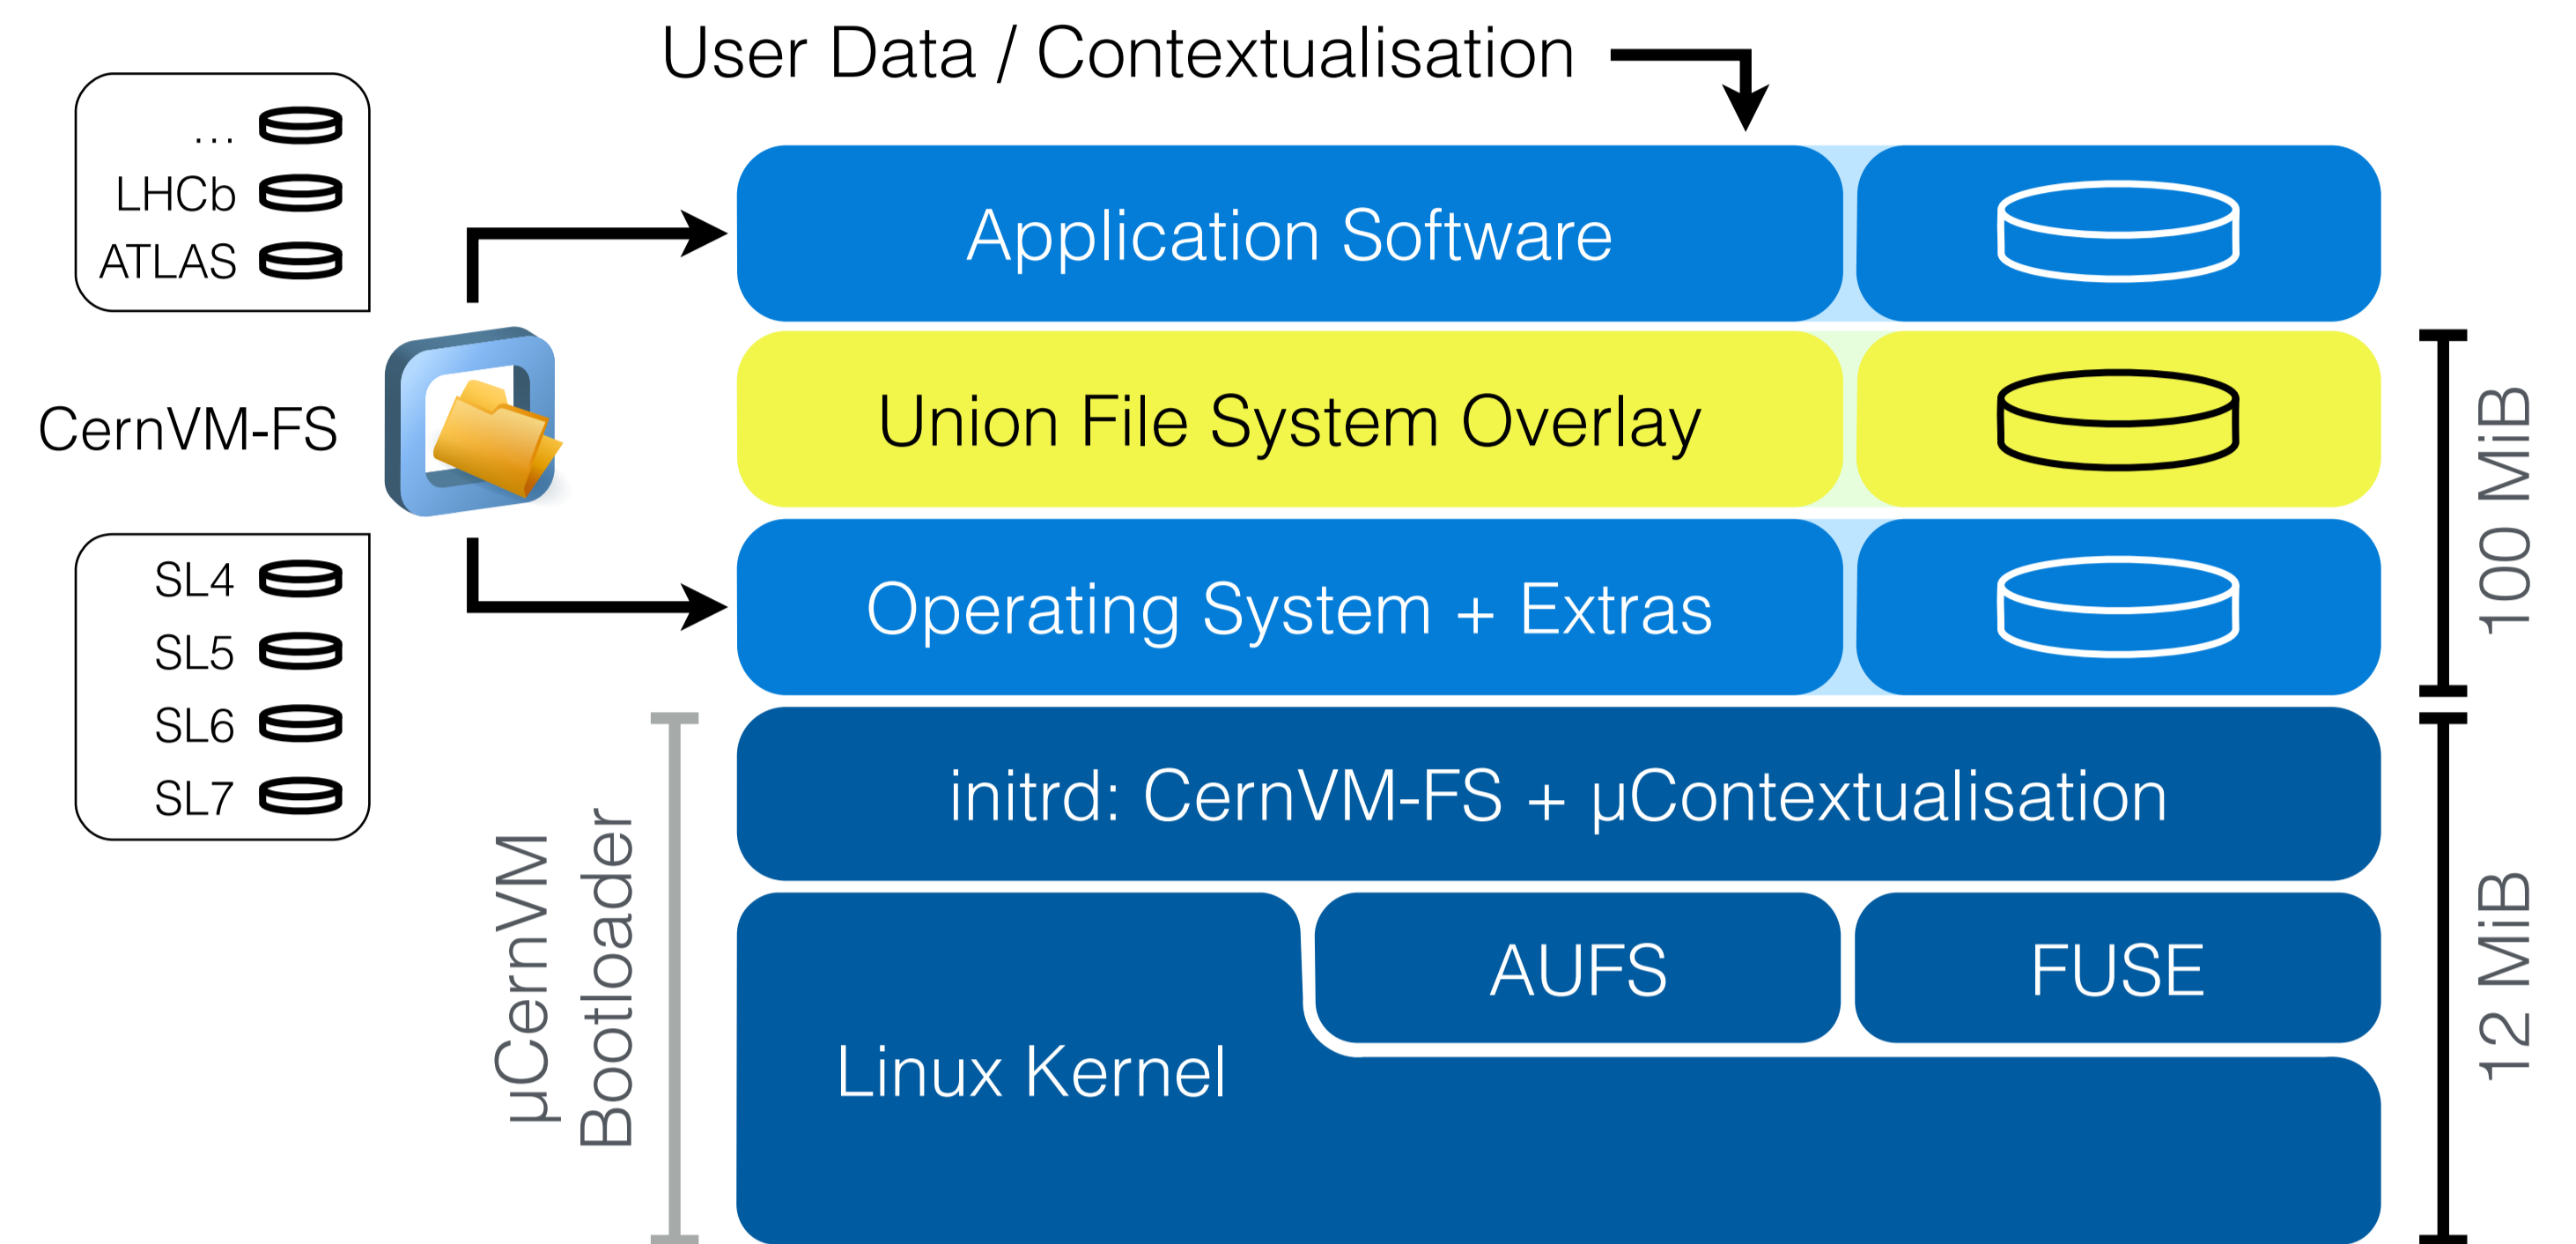


CernVM Virtual Appliance

- Minimal virtual machine image
- Linux Kernel + CernVM-FS
- Operating system "streamed" on-demand from CernVM-FS
- Can impersonate SL 4, 5, 6, 7
- Application software on CVMFS


Use Cases

- IaaS Clouds
- Development Environment
- Volunteer Computing
- Long-Term Analysis Preservation
- Outreach & Education



History-less Repositories (Garbage Collection)

- Insert-only storage to conserve full repository history by default
- Frequent and large repository changes clutter storage quickly
- *Example:* Nightly build distribution does not need history
- Garbage collection can provide history-less repositories




Configuration Improvements

- Increasing adaption of CernVM-FS needs facilitation of configuration:
 - 1) Custom Configuration Packages (by default: *cvmfs-config-default*)
 - 2) Central Configuration Repository bootstrapping of supported repos with configuration and public keys
 - 3) Automatic Stratum1 Ordering
- CernVM-FS clients use GeoIP web-service to find closest Stratum1

Pluggable Storage

- Store repository data on a alternative storage system
- Currently supported:
 - POSIX file systems
 - S3 compatible storage
- Both on Stratum0 and 1



Building Blocks of CernVM

Boot Loader (VM Image)

- 3.10 Linux Kernel + "Guest Additions"
- File system is CernVM-FS with a writable AUFS overlay
- µContextualisation to select specific OS flavour to boot
- Relatively stable < 5 minor releases

OS on CernVM-FS

- Based on Scientific Linux + Configuration and tuning + Contextualisation + Extra packages
- Fully versioned installations
- Available OS flavours:
 - SL4** Prototype
 - SL5** Stable
 - SL6** Production
 - SL7** In Development

CernVM-FS over Parrot

- Without FUSE; CernVM-FS can be accessed with Parrot *in user space*
- *However:* Parrot-connector is not always stable and lacks features
- High demand for opportunistic computing resources
- We plan to improve on the CVMFS-Parrot connector

Push Replication

- *Now:* Stratum1s regularly poll for repository updates
- *Plan:* Stratum0 will actively announce updates to Stratum1s through an API
- Will significantly lower the update dissemination latency to CVMFS clients

Scientific Linux 7 Support

- Working prototype exists, but:
- Necessary CernVM-FS changes:
 - File capability support (xattr)
 - systemd integration
- System tuning and integration
 - Mainly systemd related issues
- To be released as CernVM 4

CernVM Online

- Manages VM contexts on the web
 - "Pair" CernVMs with Context IDs to instantiate
 - Easy to configure entire cluster setups online
- cernvm-online.cern.ch

CernVM-FS Repositories as a Service

- *3rd-party developments* to provide CernVM-FS as a Service to small VOs
- *Idea:* a shared umbrella repository to host software of many user groups
- We plan to facilitate automatic repository management in collaboration with those projects by means of an API

Lightweight Virtualisation with CernVM

- Docker and 1xc are available (Containers in CernVM)
 - *cernvm-fork* for simple spawning of isolated runtime environments
- ```
$ cernvm-fork fork03 --new --fast --cvmfs=sft.cern.ch --run=/cvmfs/sft.cern.ch/exe
```
- 
- local file cache of CernVM-FS is shared between containers
  - CernVM as a Docker container needs CernVM-FS inside:
    - Mount inside the container (needs privileged container)
    - Map mounts from outside (needs host collaboration)
    - Use CernVM-FS over Parrot (stability concerns - cf. left side)
  - Union file system might not be available inside
    - Requires mapping/copying of the root file system (/) layout
  - This is an often requested feature - more development needed

