

cvmfs at GSI

cvmfs Workshop @ CERN 2024

Sören Fleischer

2024-09-17

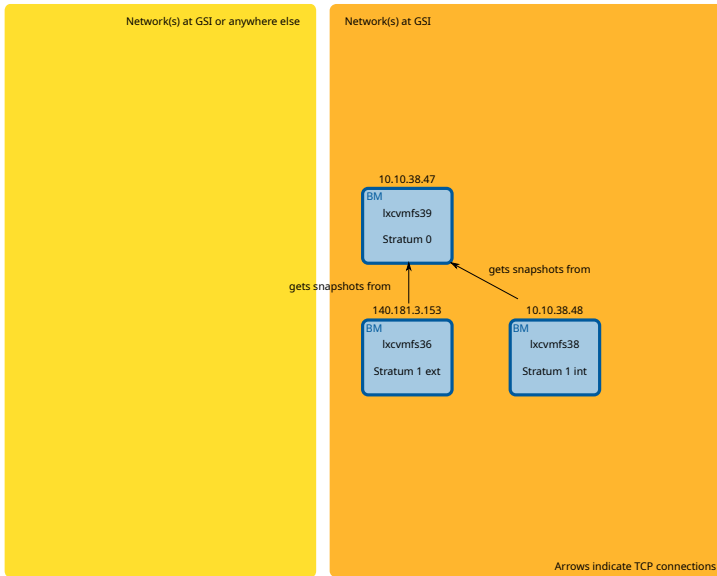
Operations

- OS: Debian 10 11
- OS-level configuration: Chef/CINC
- Updates: unattended-upgrades
- cvmfs workflows: Ansible

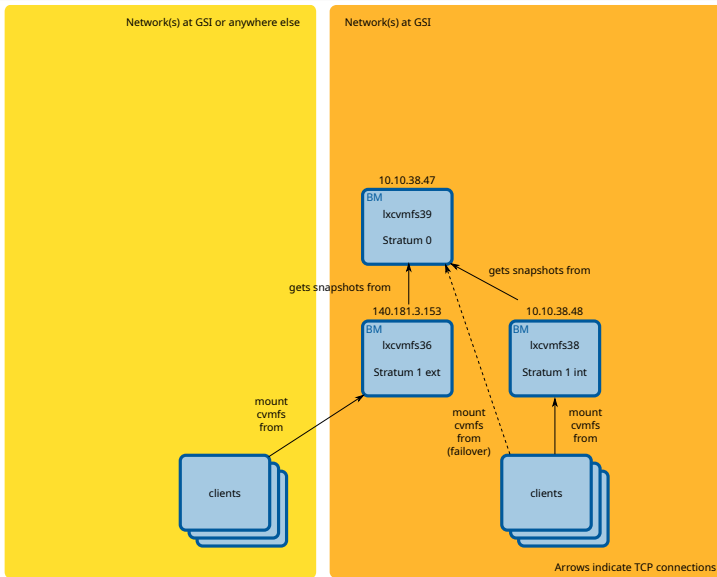
Current cvmfs Design at GSI



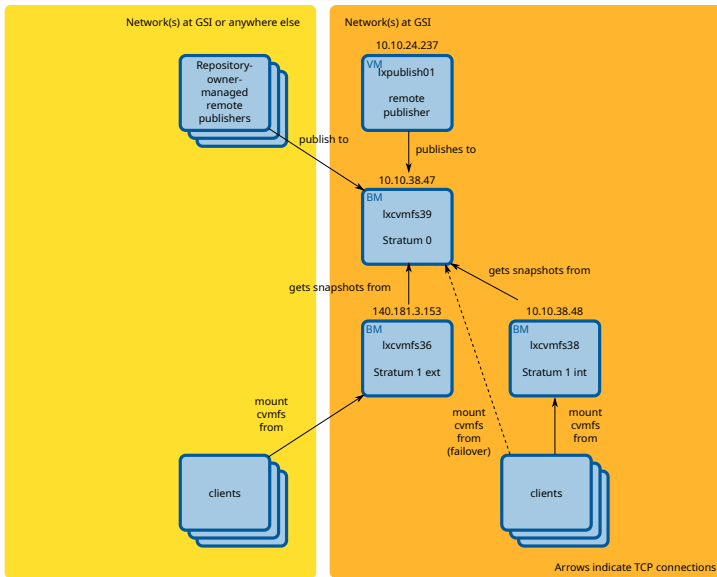
Current cvmfs Design at GSI



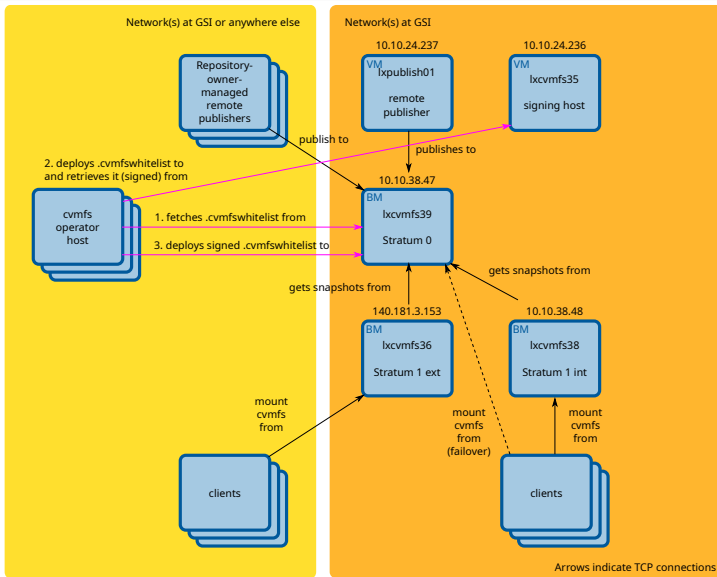
Current cvmfs Design at GSI



Current cvmfs Design at GSI



Current cvmfs Design at GSI



Items to be backed up

- `/etc/cvmfs/`
- `/srv/cvmfs/`
- `/var/spool/cvmfs/`
- `/etc/apache2/`
- `/etc/fstab`

Backup Workflow Stratum 0

- Every day a script
 - stops **cvmfs-gateway** and **apache2**
 - umounts all cvmfs repositories
 - creates a **btrfs** snapshot
 - starts **apache2** and **cvmfs-gateway**
 - mounts all cvmfs repositories
- Effective downtime each day: $\approx 5..10$ seconds
- The 2 most recent snapshots are kept on the Stratum 0
- Those 2 snapshots are picked up by **ISP** (Formerly TSM) every day

Stats

- ~~16~~ 20 repositories
- ~~7~~ 9 repositories on external Stratum 1
- ~~4~~ 5 repositories use their own publisher host
- \approx ~~270~~ 600 GB total

Main Use Case

- Making executables/libraries accessible to batch farm / cluster nodes
 - Traditional "copying pre-compiled software packages to cvmfs"
 - Spack [package management tool]
 - "VAEs" (Virtual Execution Environments): Apptainer Images

VAEs / Containers

Main(?) purpose of containers:

- Making a specific set of libraries visible to an executable
 - Without installing those libraries to the host system
 - Eliminating or greatly reducing Dependency Hell
 - Saving resources compared to proper virtualisation

VAEs / Containers

- The host system can be different than what the users see when they log in
- The host system can be upgraded/replaced without breaking every user's workflows
- Users who need different library environments can be accommodated, given they supply an Apptainer container image
- However, the Linux Kernel seen by the application is always that of the host system

VAEs / Containers

Submit Nodes	VAE	Support EOF
virgo.hpc.gsi.de	No	-
vae23.hpc.gsi.de	Yes	2025
vae24.hpc.gsi.de	Yes	2026

Figure: <https://hpc.gsi.de/virgo/user-guide/access/submit-nodes.html>

- Interactive: ssh into the VAE/container
- Batch: For each batch job submission, the user can configure which VAE they want to use