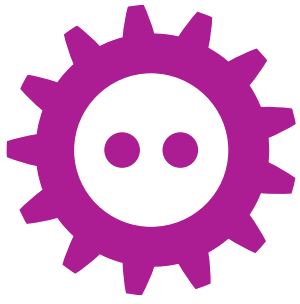

FOSDEM

Multi-Petabyte Data Distribution in Industry & Science with CernVM File System

Georgios Christodoulis, CERN
Andriy Utkin

2026

Slides:

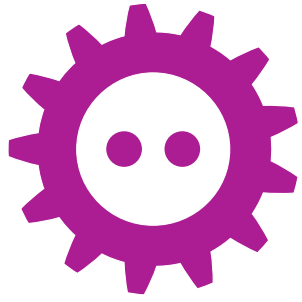

FOSDEM

CVMFS

Georgios Christodoulis, CERN
Andriy Utkin

2026

Slides: 

CVMFS INSTALL PARTY + ADOPTION Q&A (BOF)

FOSDEM 2026 / Schedule / Events / BOF/Unconference / CVMFS install party + adoption Q&A

CVMFS install party + adoption Q&A

Track: BOF/Unconference
Room: UA4.222
Day: Sunday
Start (UTC+1): 10:00
End (UTC+1): 10:55

Get help in understanding and setting up CVMFS, a highly scalable network filesystem.

The session is initiated by a contributor to CVMFS. Core maintainers of CVMFS may be joining us.

Developers and adopters of other Software Defined Storage technologies are also welcome to join and lead their discussions.

Event:

OPEN SOURCE AT CERN

stand @ K level 1 (group A)

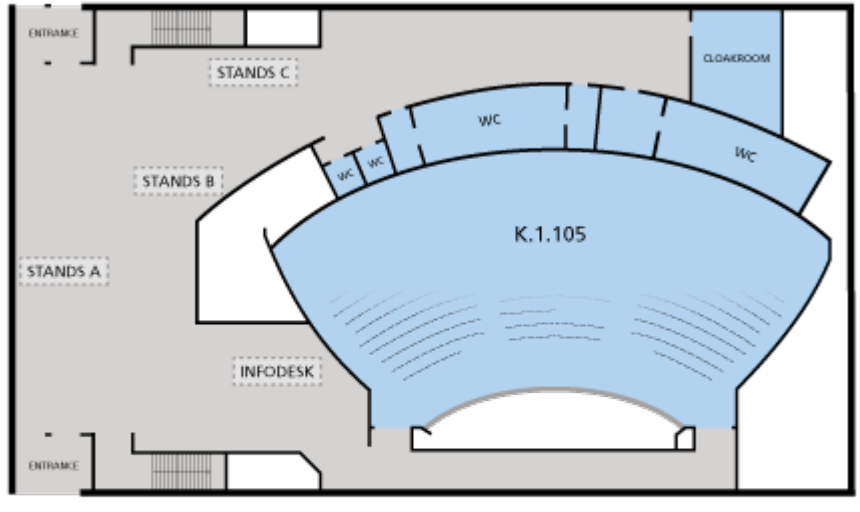

BUILDING K
LEVEL 1

SOFTWARE DISTRIBUTION IN HEP

Speaker notes

Something that people may not be very aware of is the amount of data that experiments in HEP are producing! The CERN Data Center stores more than 30 petabytes of data per year from the LHC experiments. These data along with the software used for their processing and analysis needs to be further distributed to laboratories all over the world.

LAZY-PULLING CONTAINER IMAGES

Lazy pulling in action

https://asciinema.org/a/E5MV9aVZBBF1Qp6W

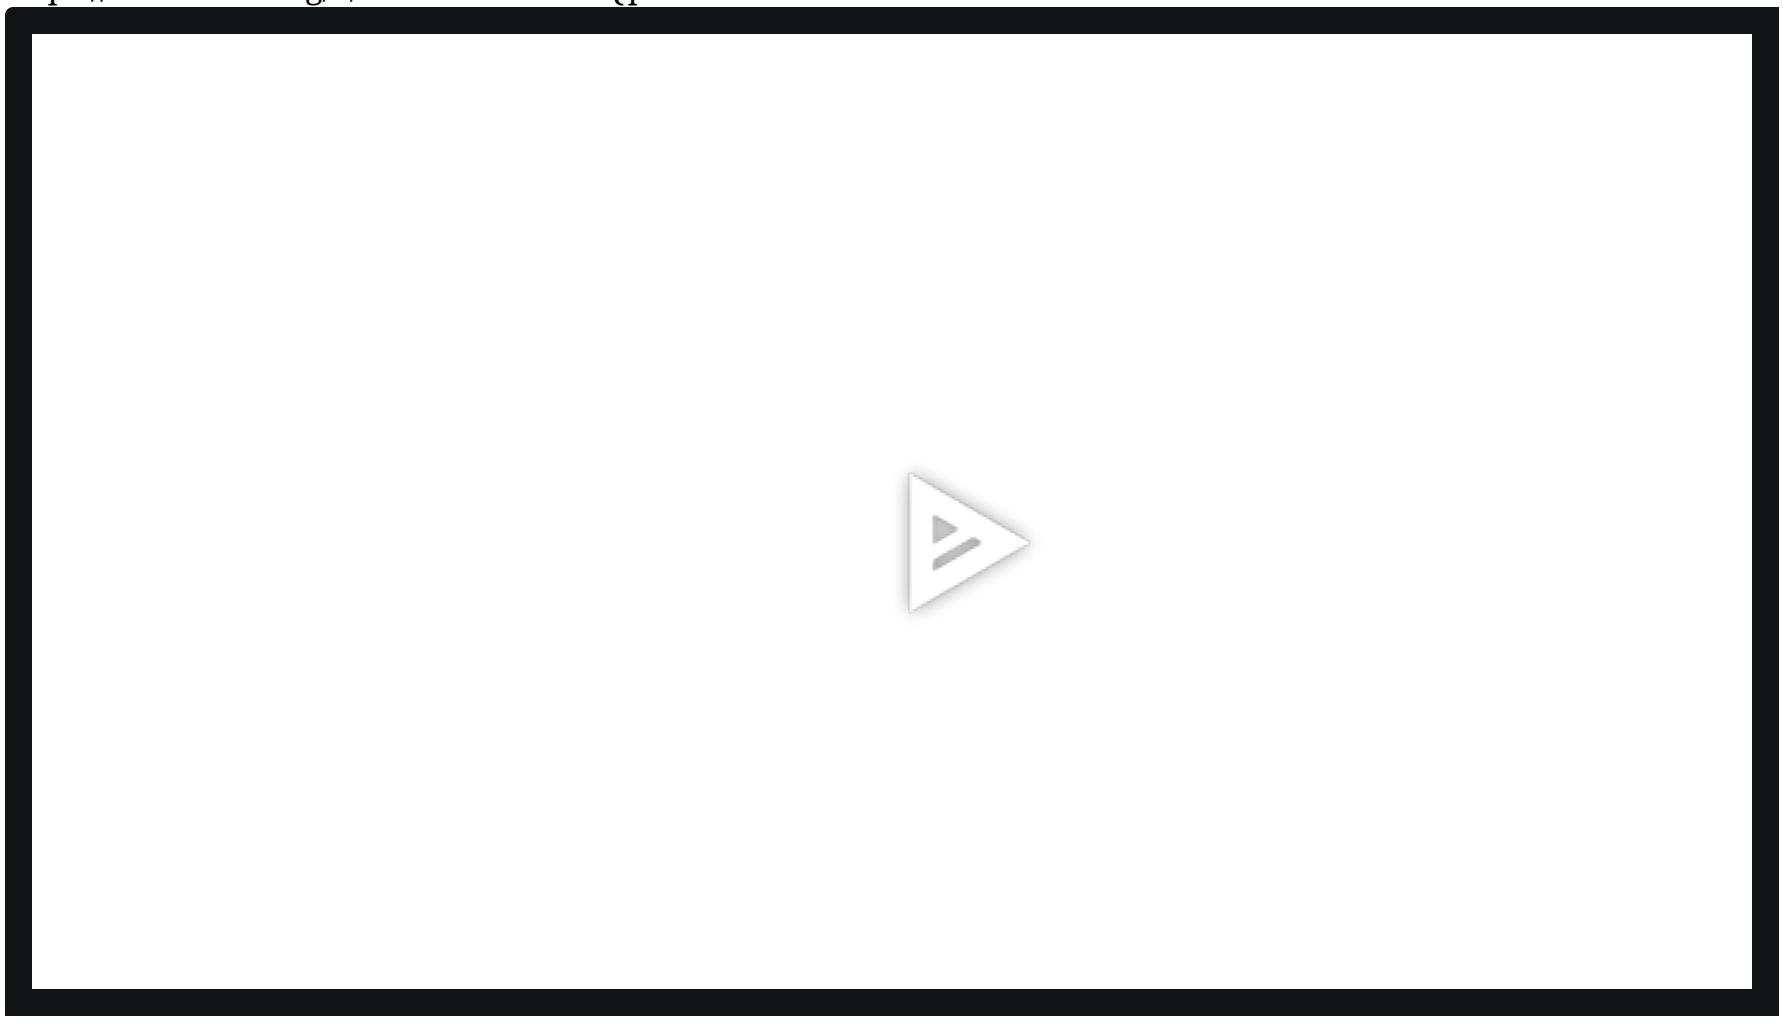

Speaker notes

Here we use the lazy pulling. Here we use a popular image slightly bigger than 1 Gigabyte. This image is developed and distributed by and for scientists, and is somewhat popular. It takes 25 seconds to fetch whatever is strictly necessary and execute the python process which loads the target module. Afterwards, the CVMFS cache is 600 Megabytes. So we skipped downloading quite a lot.

Normal image pulling

https://asciinema.org/a/wonwMFoArTR4i20P

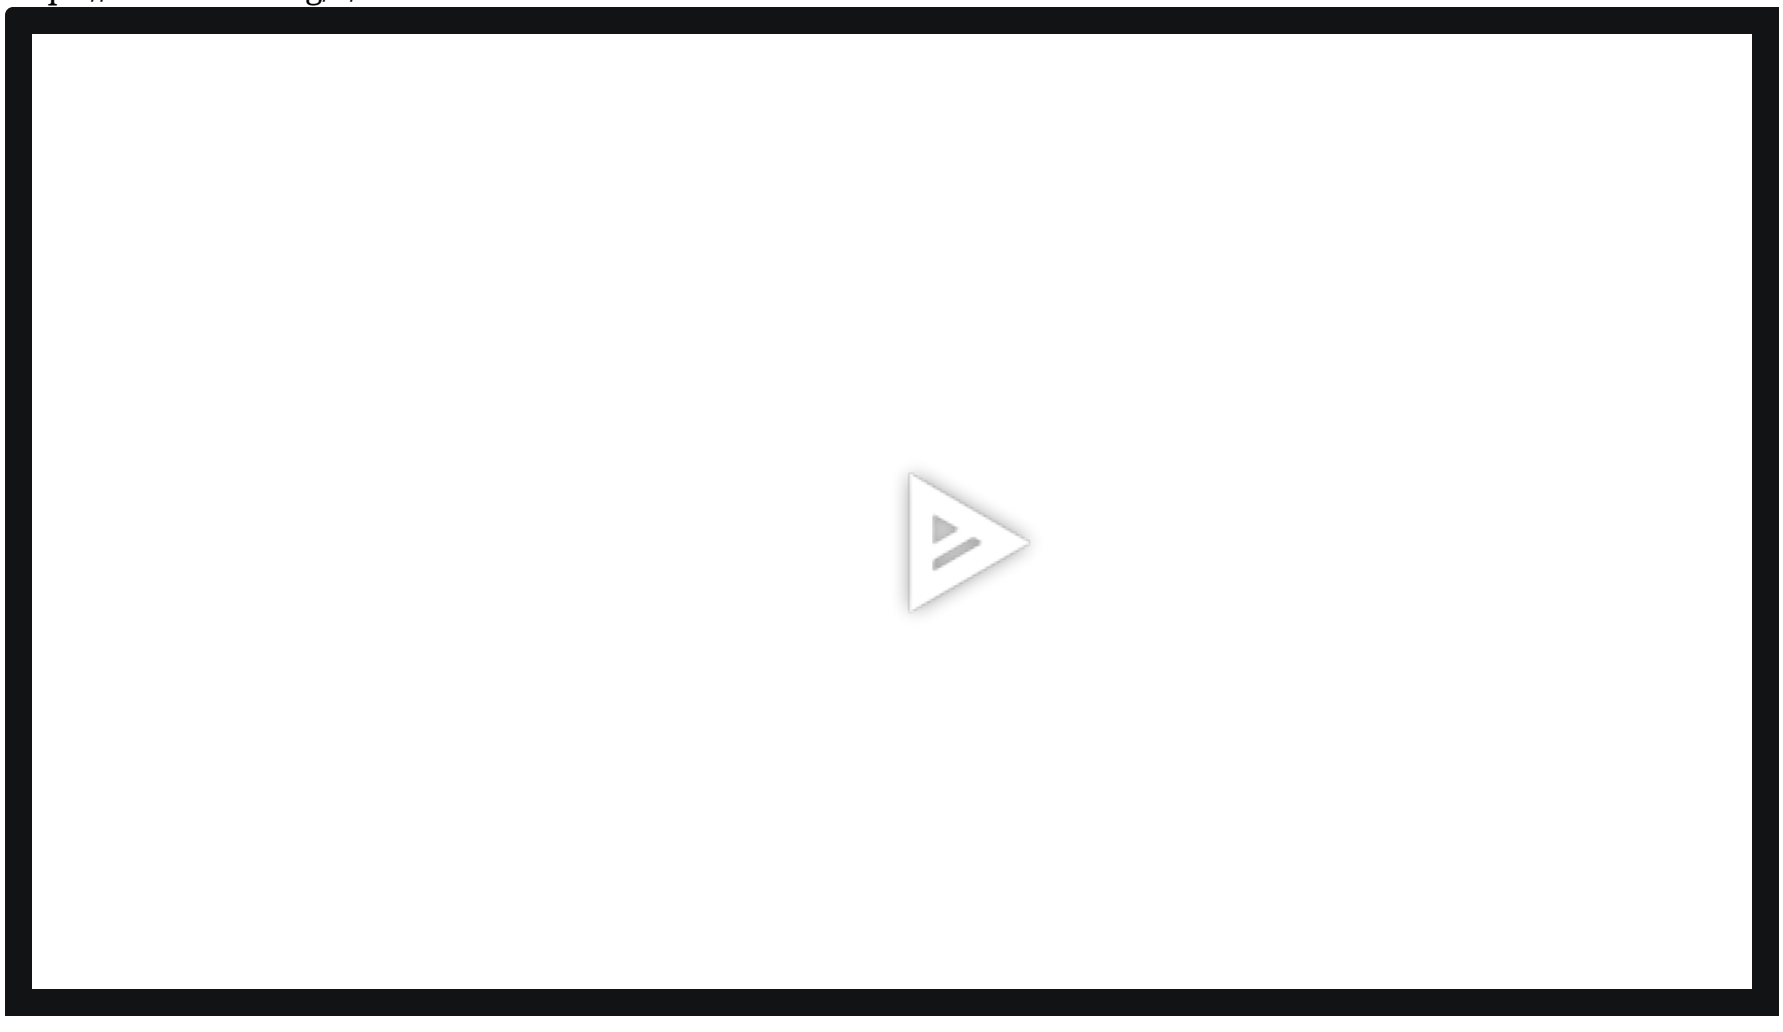

Speaker notes

You can, of course, use the same container image the usual way: issue a normal podman run command and let it download the image. This demo was recorded on a very fast connection, by the way, the uplink is upward of a gigabit. (Wait for the animation to end with the "no space left on device") Unless it doesn't fit on your disk, in which case you'd have to go on a cleanup diversion or resize your machine.

REFERENCES

doc, KubeCon25 slides, video

CVMFS-POSIX-TOOLS

cvmfs_**chgrp**
cvmfs_**chmod**
cvmfs_**chown**
cvmfs_**insert**
cvmfs_**ln**
cvmfs_**mkdir**
cvmfs_**rm**
cvmfs_**rmdir**
cvmfs_**rsync**
cvmfs_**setfacl**
cvmfs_**touch**

Speaker notes

These are new tools, modelled after popular file utilities, which offer the convenience of direct effect on CVMFS repo.
Can be seen as easier to use tools.
They provide speedups in some modes.
They can help avoid running out of disk space in scratch area.

CHANGING CONTENTS, THE USUAL WAY

- `cvmfs_server transaction`
- make changes
- `cvmfs_server publish`

Speaker notes

To explain why they are needed, let's briefly see how changing contents **currently** works from the user perspective. The first method you learn from the documentation to change the contents of your repo is this:

```
* `cvmfs_server transaction`
* make changes
* `cvmfs_server publish`
```

Publish command in action

https://asciinema.org/a/0US0PO2VgJNYVuHJ

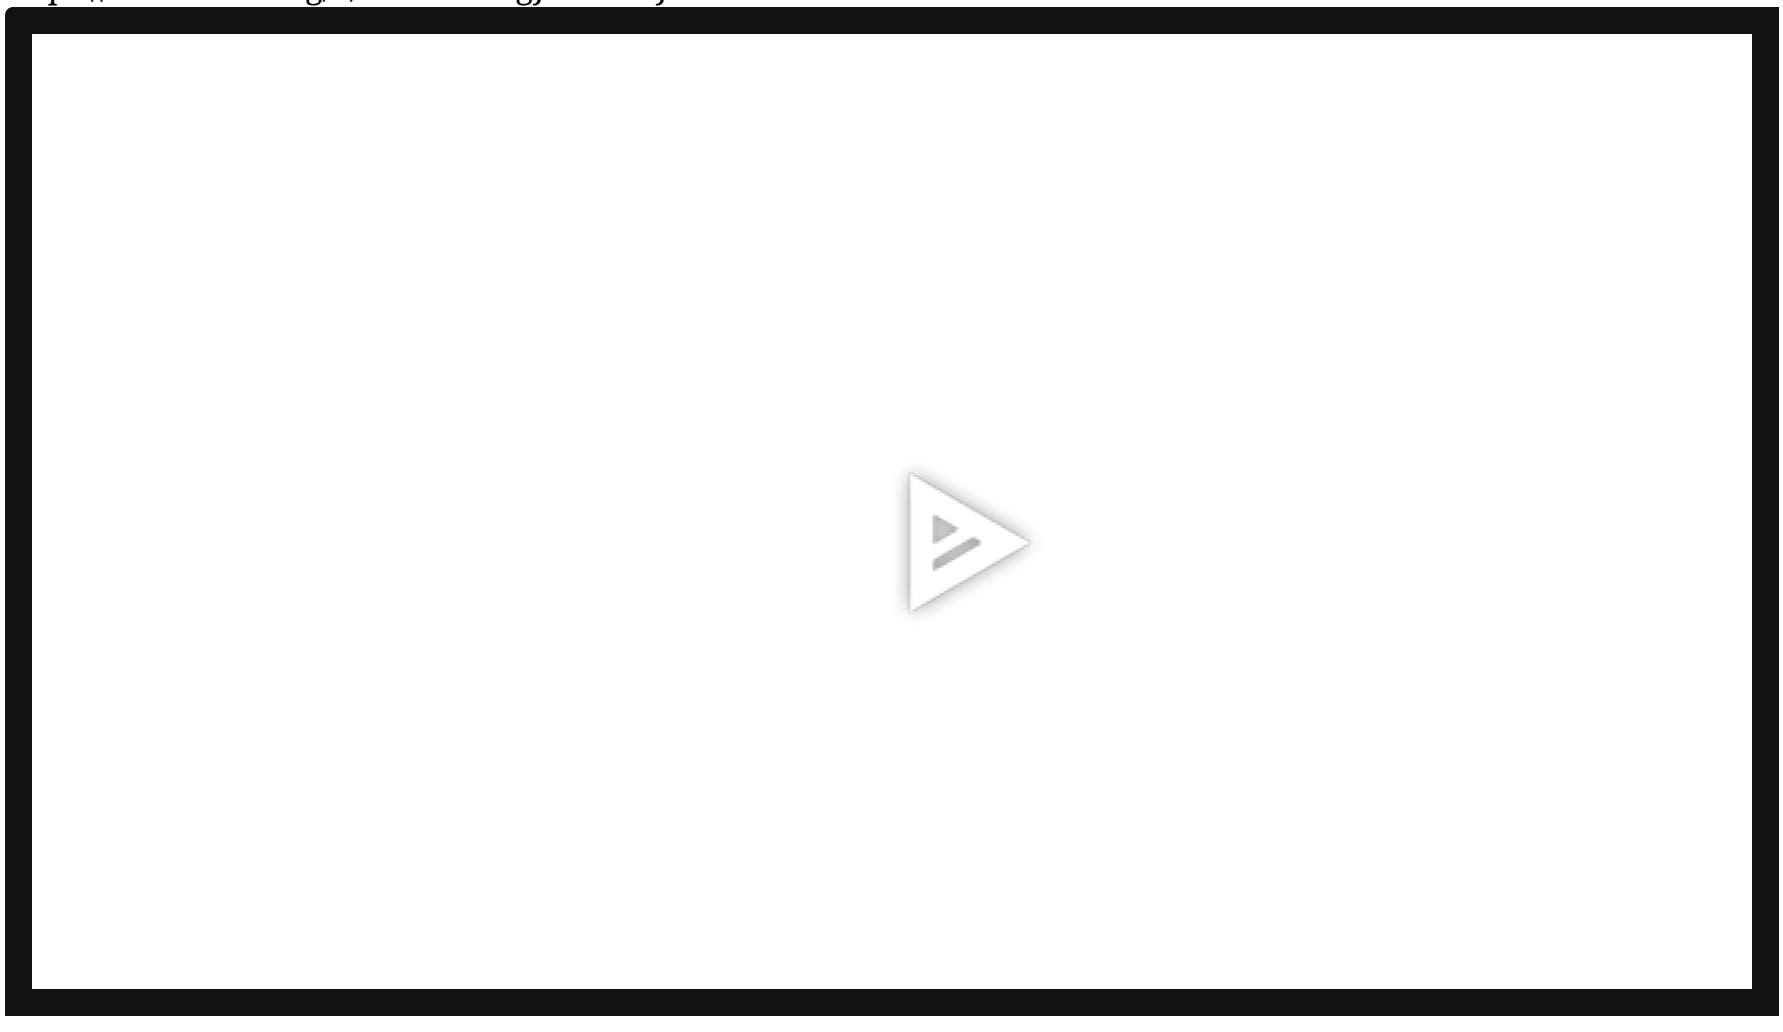

Speaker notes

This is publish method adding a 4 Gigabyte file into an S3-backed CVMFS repository.

`transaction` command changes the overlay mount from read-only to read-write.

The file first ends up in the scratch dir.

And after the `publish` command, the file is in the actual CVMFS mount.

`cvmfs_server transaction`	chang overla mode
`cp bigmodel.gguf /cvmfs/my.repo`	copy t scratc area
`cvmfs_server publish`	copy t primar storag

Speaker notes

Importantly, in this way, two expensive steps must occur to accomplish the task of publishing a new file.

First, the new file is copied into local scratch area.

SKIPPABLE BEGIN This is dealt by overlayfs, it provides a nice illusion of editing the filesystem directly, and it enables the next step, the publish command, to instantly see what are the changed files to upload. SKIPPABLE END

Second, the copy in the scratch area is transferred to the primary storage of CVMFS repo.

It can be a local directory if we are working directly on the main server, called Stratum 0.

It can be S3 location.

And it can be a remote CVMFS Gateway service.

cvmfs_rsync in action

https://asciinema.org/a/eyynoSWvecI5fwvh

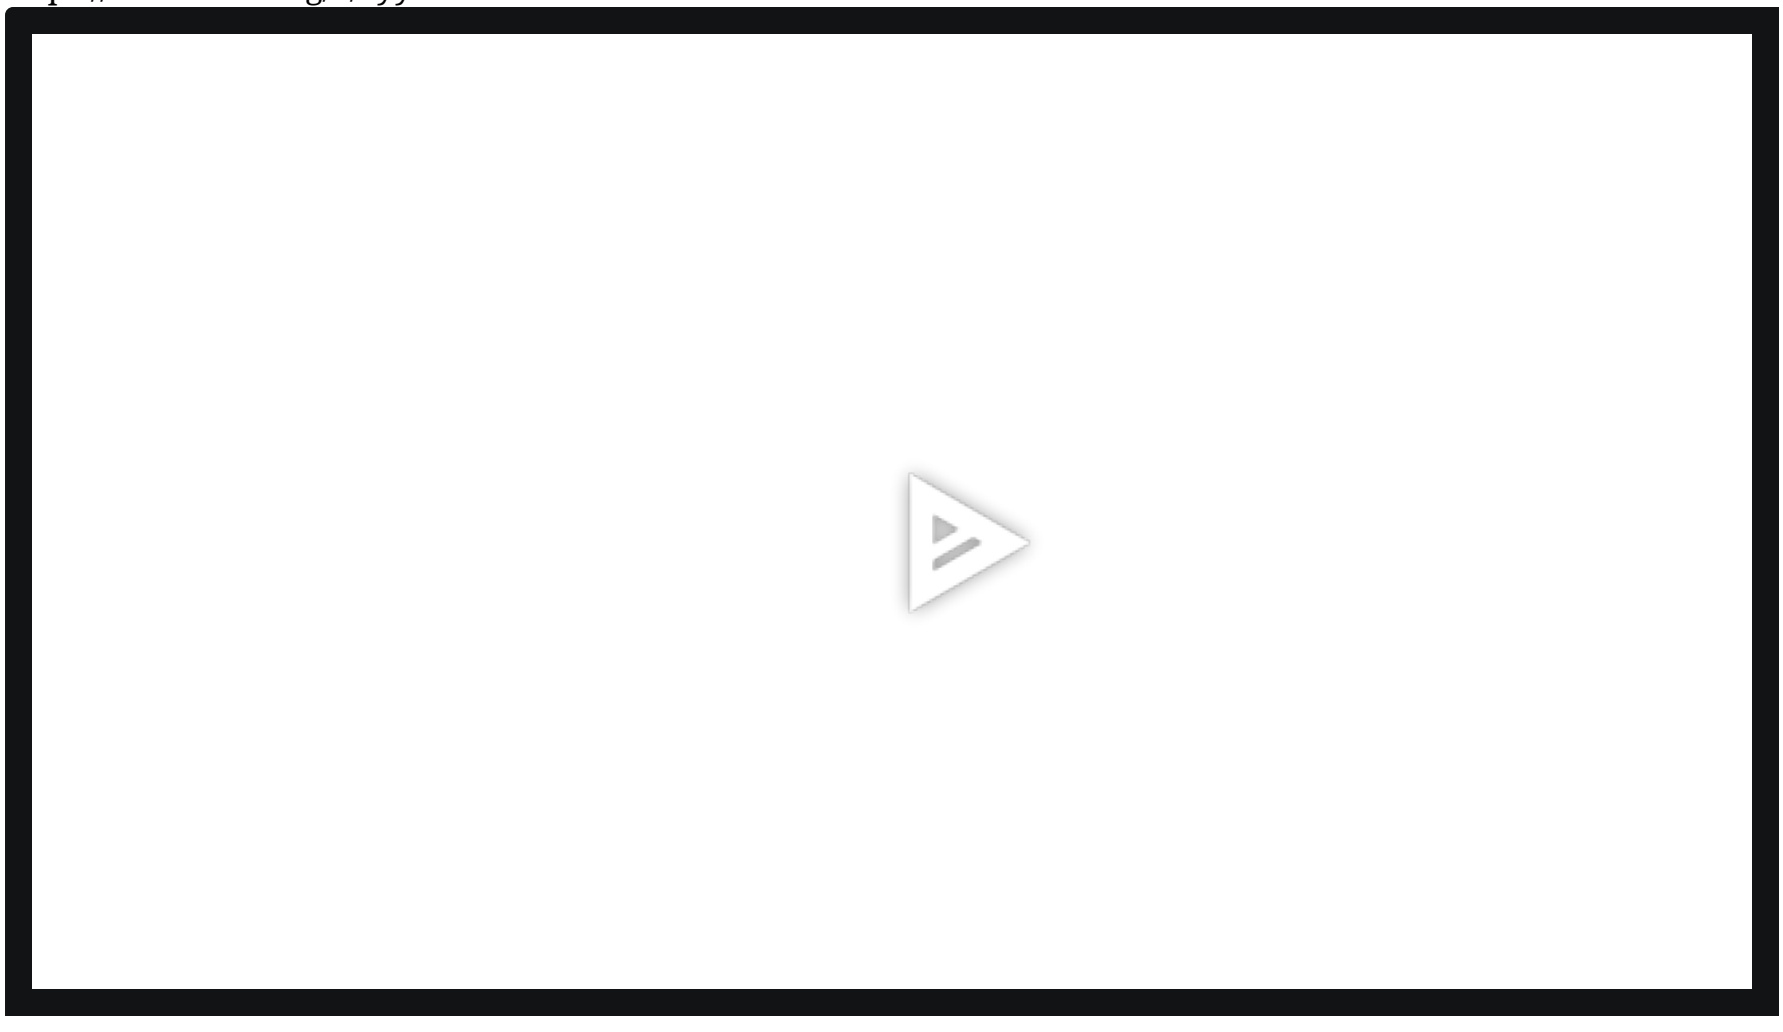

Speaker notes

You can also choose whether to store new files in Content Addressable Storage fashion, or represent the tree hierarchy and original filenames in S3 locations.

CVMFS-POSIX-TOOLS: UNDER THE HOOD

- uploads all the new files to S3 directly
- pushes metadata changes to CVMFS Gateway host
- waits for new CVMFS revision to become visible

Speaker notes

Note: this doesn't mean you need AWS, many protocol implementations work - Ceph, Garage, MinIO, Azure has some quirk but also works.

REFERENCES

Container Images and CVMFS: doc, KubeCon25 slides, video

cvmfs-posix-tools: pull request

Q & A

Georgios Christodoulis, CERN
georgios.christodoulis@cern.ch

Andriy Utkin
hello@autki

FOSDEM 2026 / Schedule / Events / BOF/Unconference / CVMFS install party + adoption Q&A

CVMFS install party + adoption Q&A

Track: BOF/Unconference
Room: UA4.222
Day: Sunday
Start (UTC+1): 10:00
End (UTC+1): 10:55

Get help in understanding and setting up CVMFS, a highly scalable network filesystem.

The session is initiated by a contributor to CVMFS. Core maintainers of CVMFS may be joining us.

Developers and adopters of other Software Defined Storage technologies are also welcome to join and lead their discussions.

Open Source at CERN

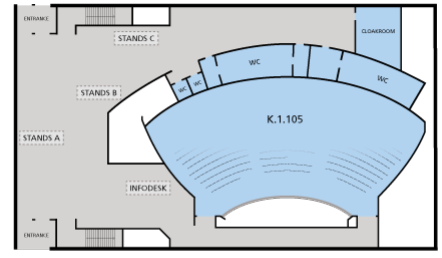

BUILDING K
LEVEL 1

CVMFS INSTALL PARTY + ADOPTION Q&A (BOF)

FOSDEM 2026 / Schedule / Events / BOF/Unconference / CVMFS install party + adoption Q&A

CVMFS install party + adoption Q&A

Track: BOF/Unconference
Room: UA4.222
Day: Sunday
Start (UTC+1): 10:00
End (UTC+1): 10:55

Get help in understanding and setting up CVMFS, a highly scalable network filesystem.

The session is initiated by a contributor to CVMFS. Core maintainers of CVMFS may be joining us.

Developers and adopters of other Software Defined Storage technologies are also welcome to join and lead their discussions.

Event:

OPEN SOURCE AT CERN

stand @ K level 1 (group A)

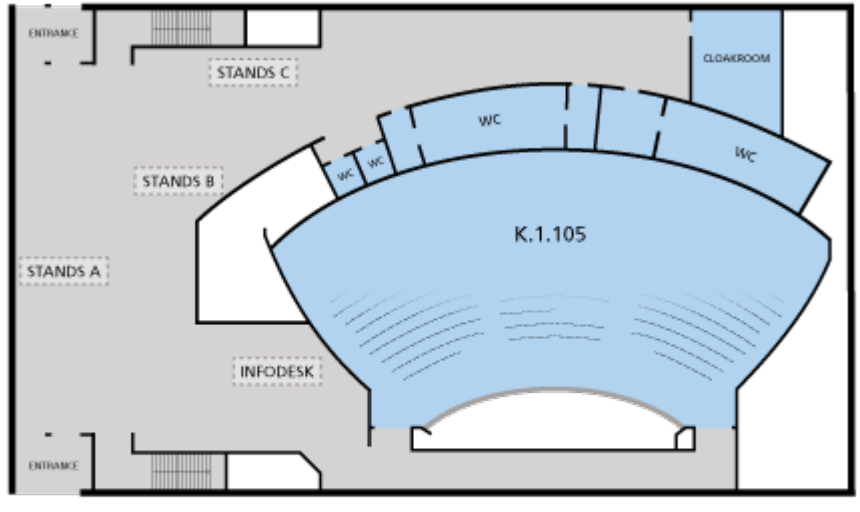

BUILDING K
LEVEL 1