



CERNVM-FS

Web-based network file system.

Big data applications and large-scale scientific computing applications running on tens of thousands of individual computers, need access to specific data processing software. The distribution of such software is challenging, even when using virtual machines or virtual containers that can package an application and its dependencies, such as Docker.

AREA OF EXPERTISE

- ICT

IP STATUS

- Open Source Software

TECHNOLOGY READINESS LEVEL

- Mature

CONTACT PERSON

nick.ziogas@cern.ch

Find out more at:

kt.cern

CernVM-FS is a web-based, global, and versioning file system optimised for software distribution. The file system content is installed on a central web server from where it can be mirrored and cached by other web servers and web proxies. File system clients download data and meta-data on demand and cache them locally. Data integrity and authenticity is ensured by cryptographic hashes and digital signatures.

CernVM-FS is used, among others, by the LHC experiments for the distribution of 100 million files and directories of LHC experiment software onto tens of thousands of nodes distributed worldwide.

FEATURES

- Various Linux distributions (x86, AMD64, ARM) & Mac OS X (client) supported
- Global-scale open source file system optimised for software distribution.
- Data transport via standard HTTP protocol.
- Data integrity secured by cryptographic hashes and digital signatures.
- File system level versioning.
- Data de-duplication.
- Transparent data compression/decompression and file chunking.
- Capability to hot patch the file system client.
- Capability to work in offline mode providing that all required files are cached.

APPLICATIONS

- Large-scale computing, big data processing.

