

Using Pilot Jobs and CernVM File System for Simplified Use of



Containers and Software Distribution

Namratha Urs¹, Marco Mambelli², Dave Dykstra²

¹Computer Science and Engineering, University of North Texas; ²Scientific Computing Division, Fermi National Accelerator Laboratory




MOTIVATION

- High Energy Physics (HEP) experiments require massive computing resources (sites) to process data and are heavily reliant on data/software releases.
- Growing experimental needs and trends of heterogeneity in resources makes resource handling difficult.
- Scalability, reliability, and maintenance becomes critical w.r.t distributing data and software.
- Sites like HPC resources currently lack the capability to support data and software distribution, thereby requiring manual effort by system administrators to install/maintain.

BACKGROUND

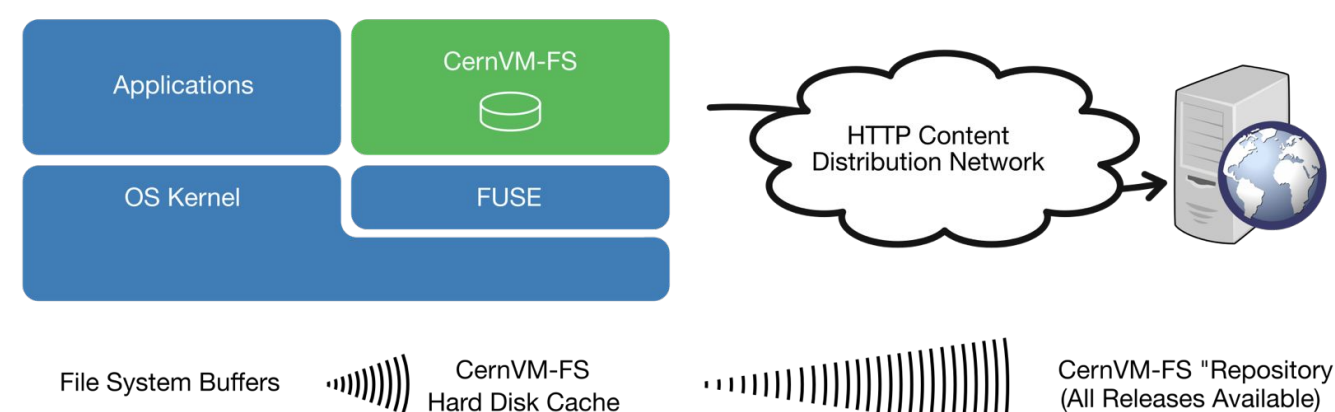
GlideinWMS framework for resource management



- Uses *pilot jobs*, aka **Glideins**, to provision reliable elastic virtual clusters
- Submits glideins to unreliable heterogeneous resources
- Glideins validate and customize the worker node for end-user jobs

CernVM File System for software distribution

- HTTP-based, read-only globally distributed file system.
- Optimized to distribute and deploy scientific software/data.
- Distribute experiment software and data such as calibrations.
- Facilitate containerization by hosting container images.
- Provide containerization software such as Singularity.

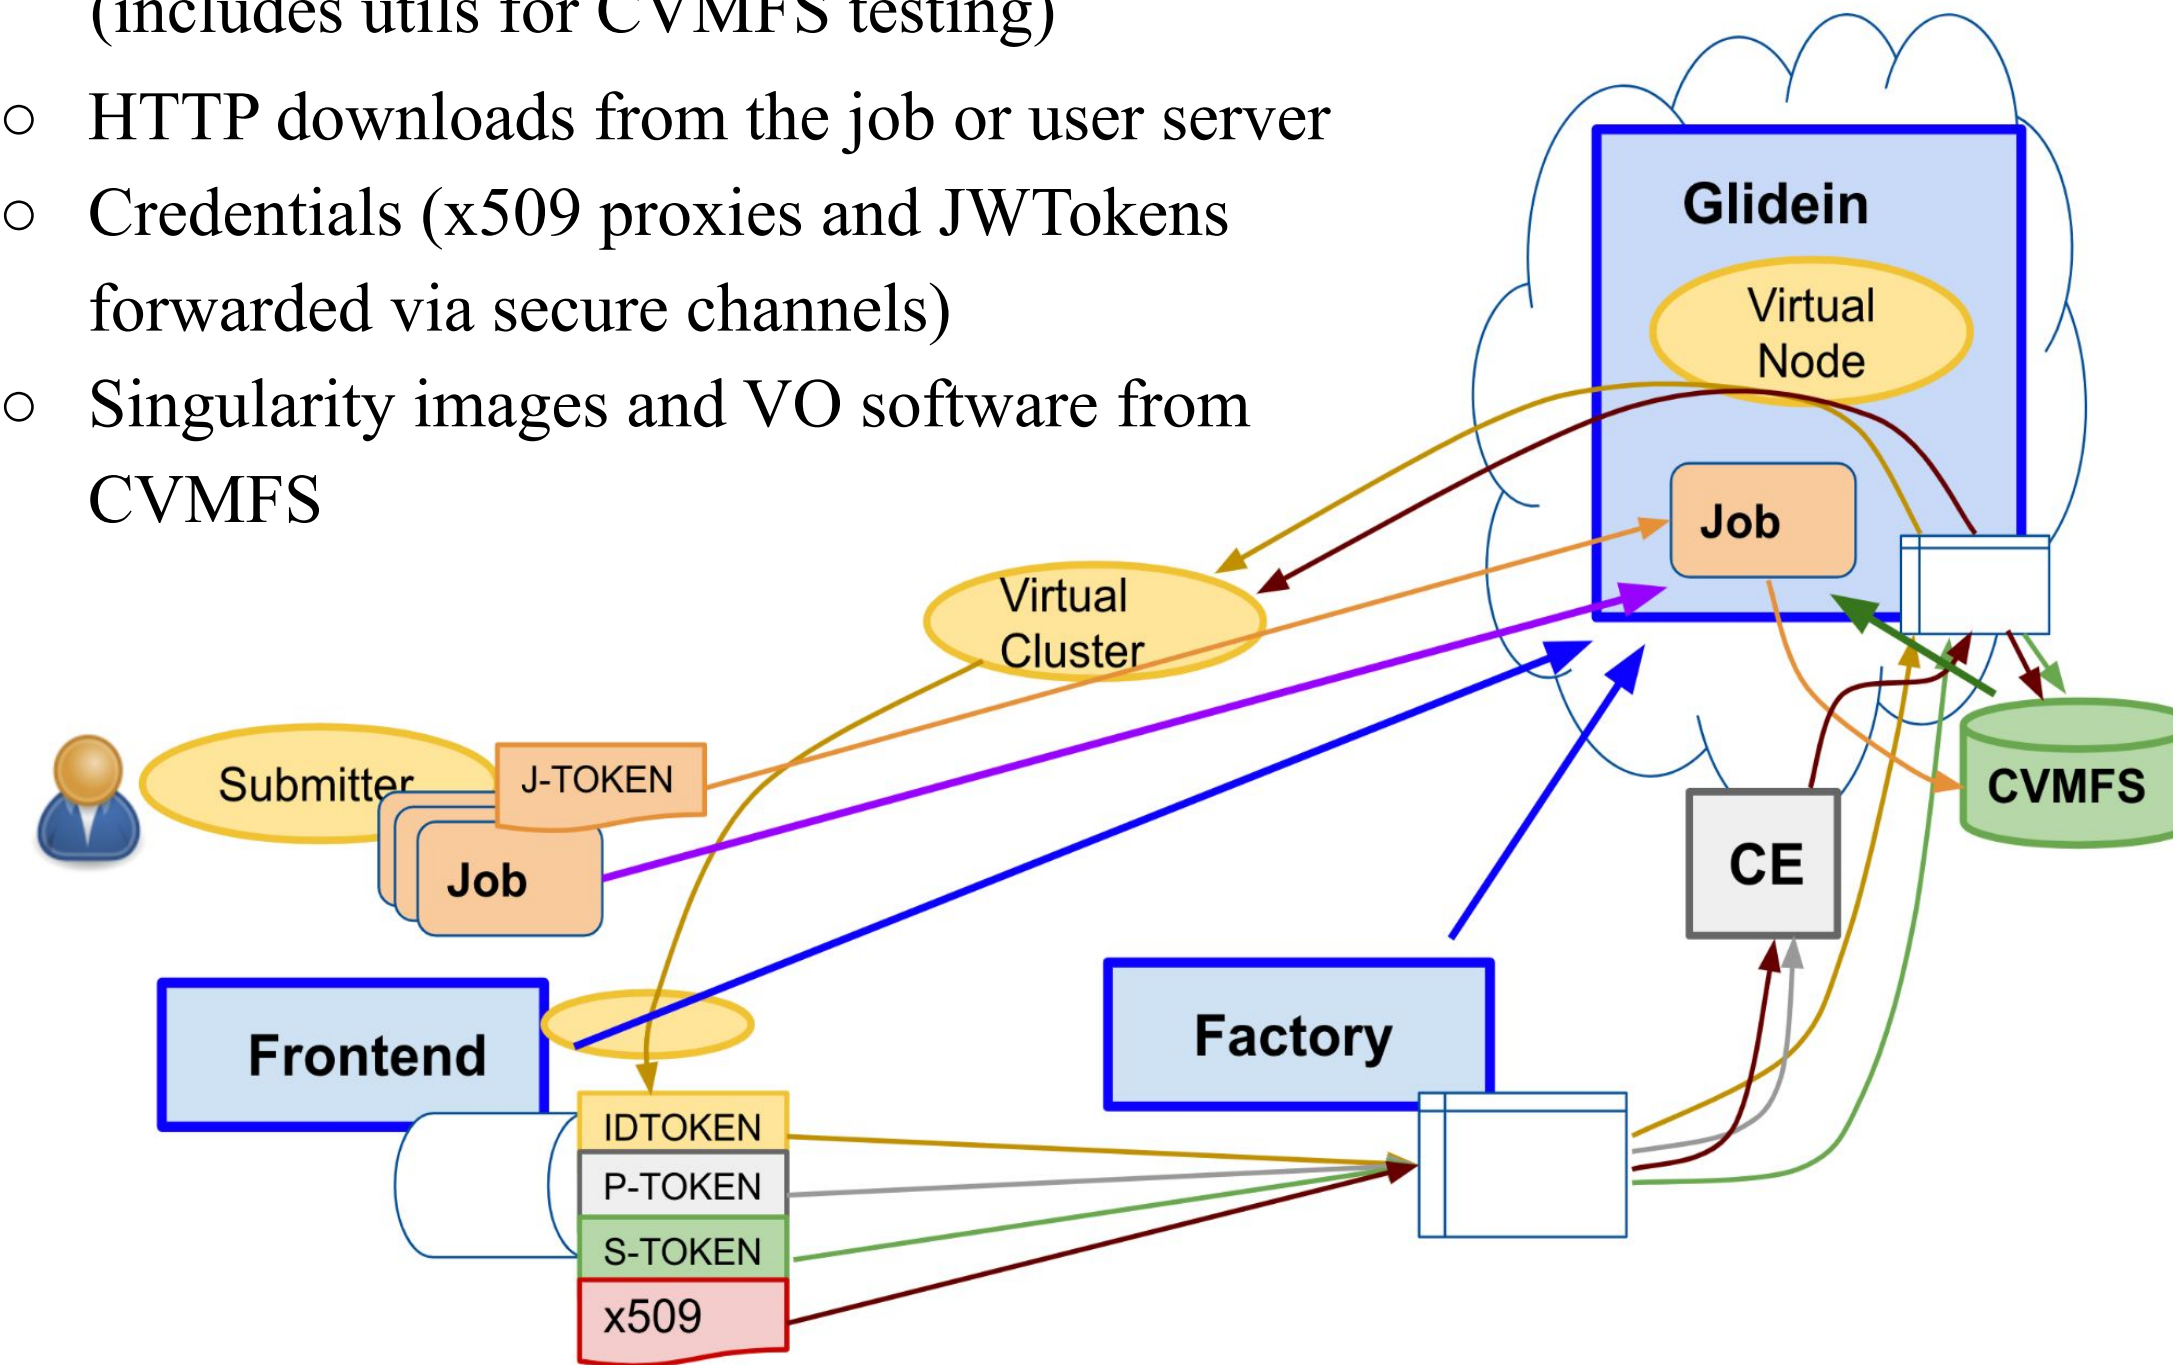


Graphic Credits: <https://cvmfs.readthedocs.io/en/stable/>

PROTOTYPE AND FEASIBILITY STUDY

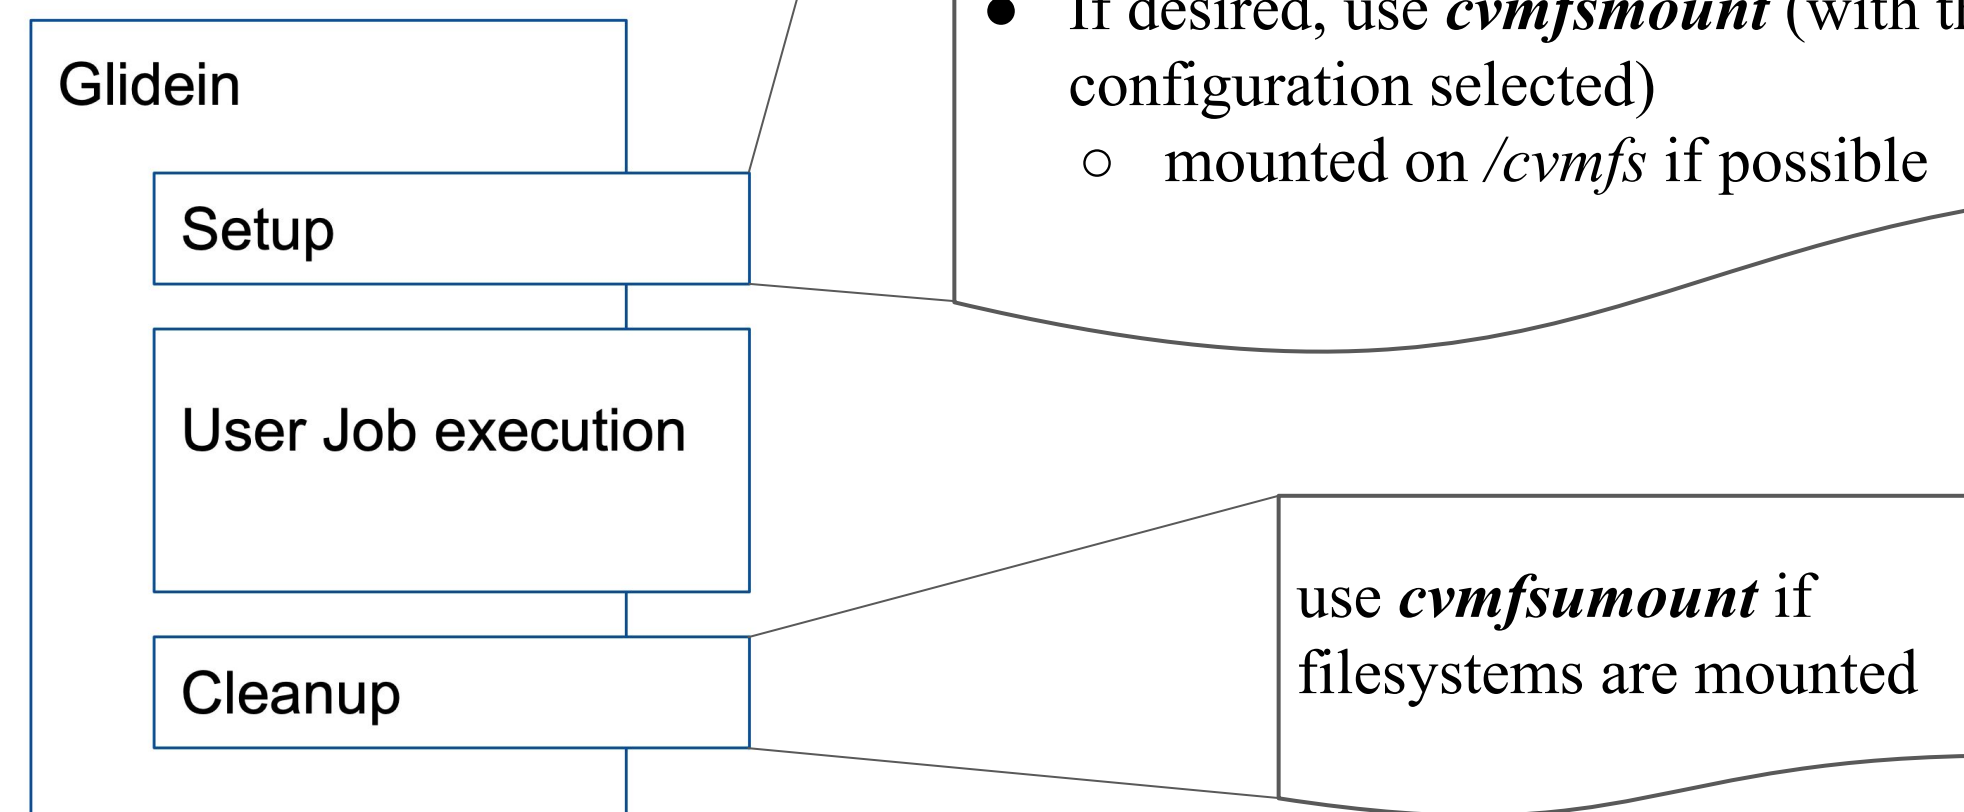
Data Movement within the Glidein

- Glidein gets submitted by Factories upon requests from Frontends that monitor users' needs
- Glidein downloads:
 - Signed HTTP downloads for validating and customizing the worker node (includes utils for CVMFS testing)
 - HTTP downloads from the job or user server
 - Credentials (x509 proxies and JWTokens forwarded via secure channels)
 - Singularity images and VO software from CVMFS



†*cvmfsexec* - supports unprivileged CVMFS use: CVMFS can be mounted by regular users without being installed by a system admin

- Checks the following:
 - OS (platform and kernel info)
 - Support for **unprivileged user namespaces**
 - **FUSE** (packages installed, user in fuse group)
- Download *cvmfsexec*[†] package (both OSG and EGI distributions)
- If desired, use *cvmfsmount* (with the configuration selected)
 - mounted on */cvmfs* if possible



FULLY AUTOMATED SOLUTION

The complete automated solution will include:

1. A tool to generate and update different downloadable distributions of the *cvmfsexec* package per resource type
2. A Glidein component to test the computing resource and trigger the proper CVMFS installation, including the download on-demand of the *cvmfsexec* package
3. A Glidein component to complete the proper cleanup at the end of the Glidein's lifespan

CONCLUSION / FUTURE DIRECTION

- Glideins now provide CVMFS to all jobs running via GlideinWMS, eliminating the need for a local installation
 - CVMFS is available on sites where it was not
 - Reduced administration effort for sites supporting CVMFS
- Necessary step towards running user jobs in containers on sites where neither CVMFS nor Singularity are available
- Plan to integrate the multi-modality aspect of CVMFS installation offered by the *cvmfsexec* package

ACKNOWLEDGEMENTS

This work was carried out under the GlideinWMS project as part of the 2020 Fermilab Computational Science Internship program.

This work has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

REFERENCES

1. Jakob Blomer, Predrag Buncic, René Meusel, Gerardo Ganis, Igor Sfiligoi, and Douglas Thain. 2015. *The evolution of global scale filesystems for scientific software distribution*. Computing in Science & Engineering 17, 6 (2015), 61–71.
2. Jakob Blomer, Dave Dykstra, Gerardo Ganis, Simone Mosciatti, and Jan Priessnitz. 2020. *A fully unprivileged CernVM-FS*. In EPJ Web of Conferences, Vol. 245. EDP Sciences, 07012.
3. Gregory M Kurtzer, Vanessa Sochat, and Michael W Bauer. 2017. *Singularity: Scientific containers for mobility of compute*. PloS one 12, 5 (2017), e0177459
4. Igor Sfiligoi, Daniel C Bradley, Burt Holzman, Parag Mhashilkar, Sanjay Padhi, and Frank Wurthwein. 2009. *The pilot way to grid resources using glideinWMS*. In 2009 WRI World congress on computer science and information engineering, Vol. 2. IEEE, 428–432.