



LHCb containers status

Andrew McNab
University of Manchester,
GridPP, and LHCb

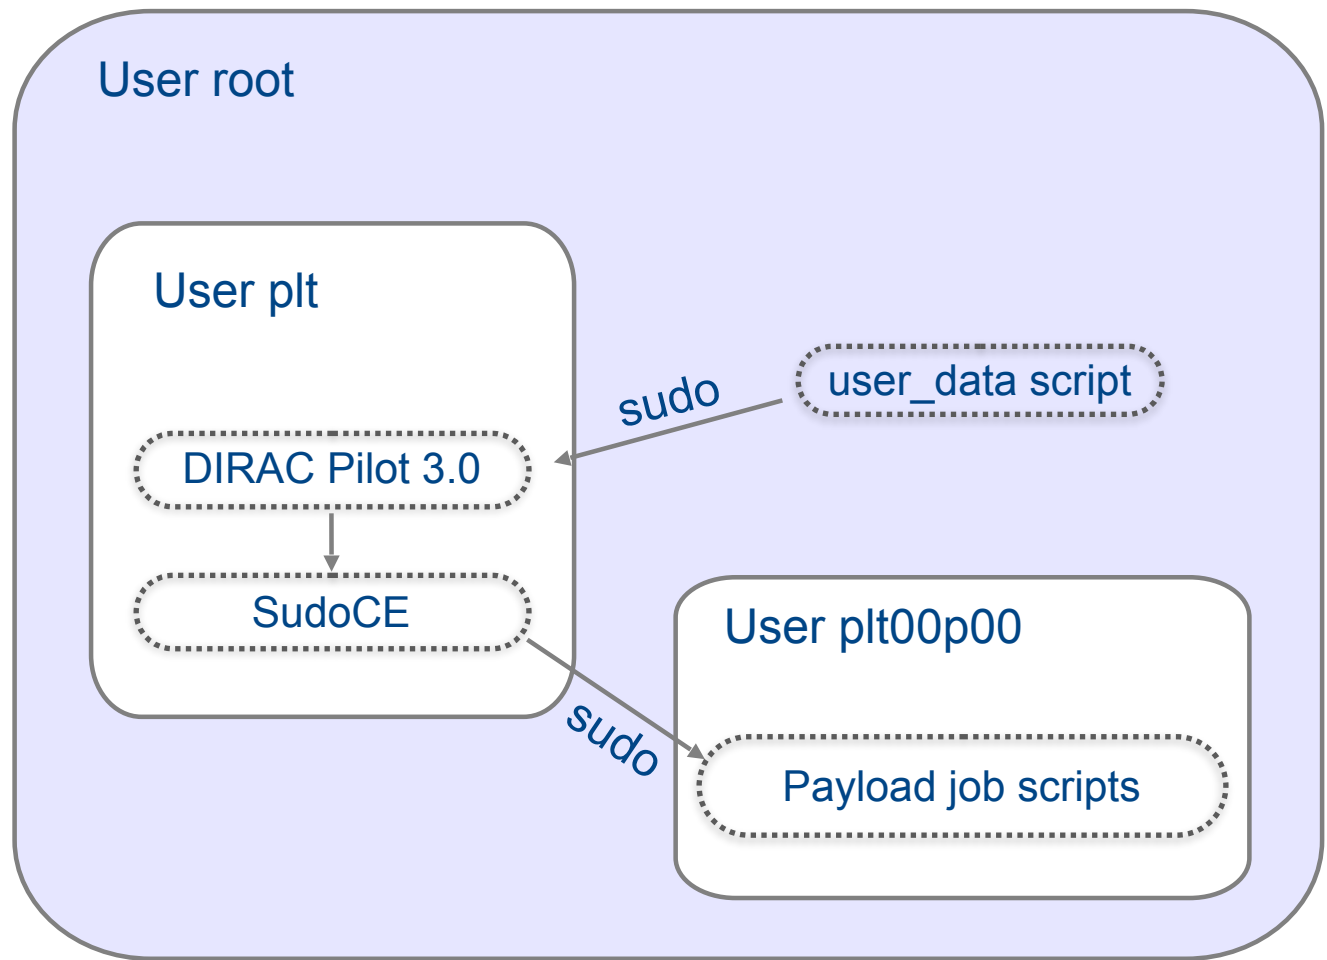


Docker and Singularity as logical machines

- This is not isolating payloads but is relevant
- LHCb has run at two sites which use LHCb containers
 - Andrew Lahiff's system at RAL, and Skygrid at Yandex
 - Both use (different) containers derived from LHCb DIRAC VMs
- Now also have generic Docker and Singularity container definitions
 - These have successfully run hundreds of ordinary LHCb production (MC) jobs
 - Use CernVM root image (ie via cvmfs)
 - LHCb cvmfs and /init script to run inside the container also provided via volumes
- So we know Singularity + CernVM root will broadly work for us

Existing LHCb VM structure

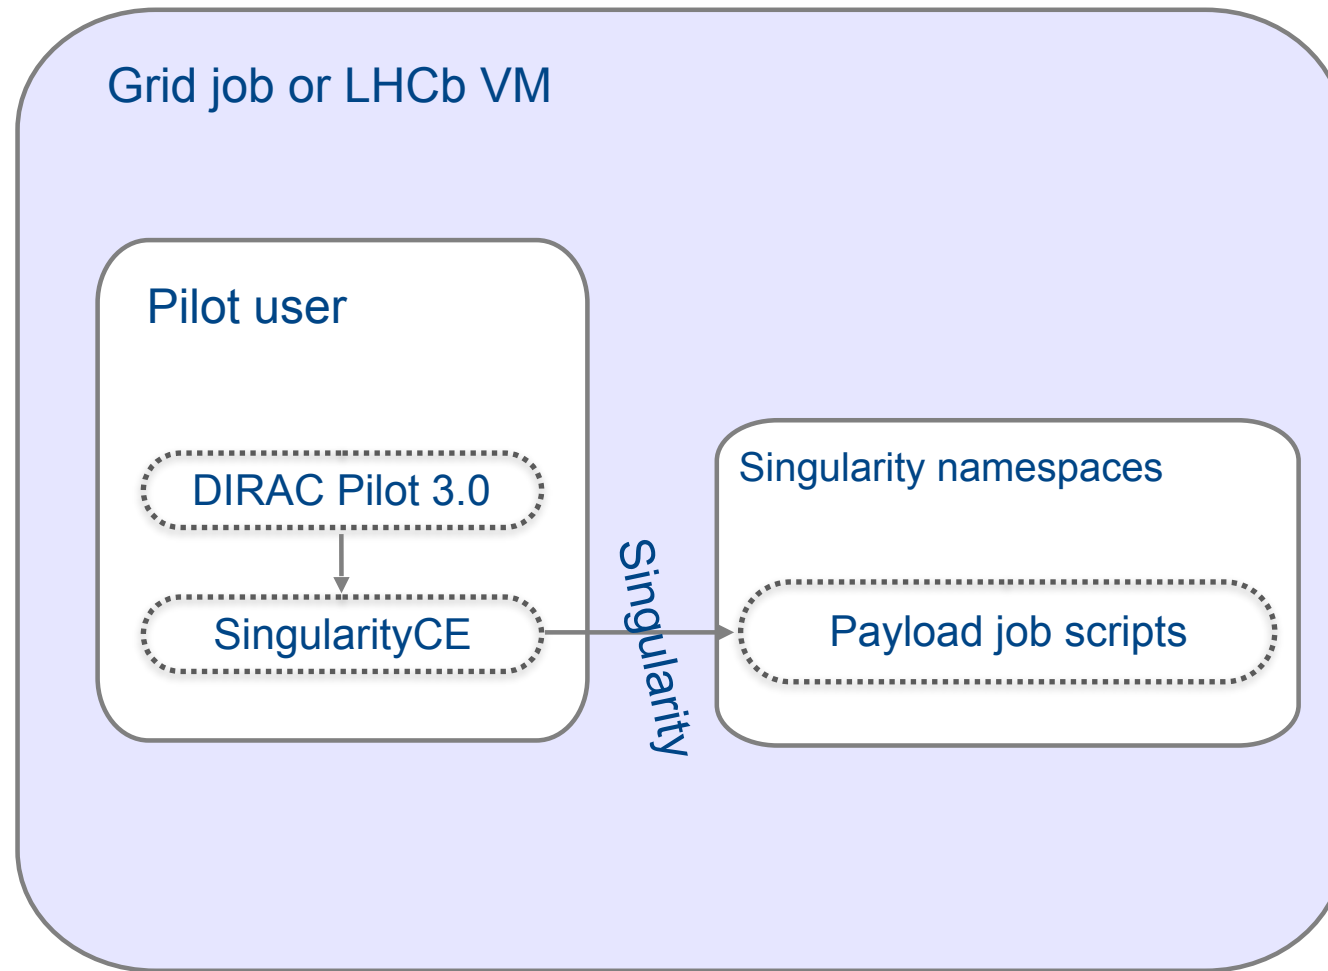
- Uses unix accounts and sudo to isolate root vs pilot vs payload
- DIRAC abstracts all forms of execution callout as “Computing Elements” (CREAM, ARC, ssh, glxexec, sudo, ...)
- We need a SingularityCE in DIRAC



LHCb payloads in Singularity containers

- SingularityCE for DIRAC exists in draft form
 - Developed by GridPP DIRAC rather than LHCb
 - Still needs testing and integration into DIRAC
 - Still needs testing with LHCb jobs
- We believe it will be possible to operate without OverlayFS or setuid Singularity
 - This probably requires coordination about the working directory paths so they can match inside and outside the container, as we discussed here previously
- We don't expect to need any special images beyond the CernVM root filesystems in /cvmfs
 - We don't want to have to support multiple OS images

Future job structure with Singularity





Summary

- We know we can run LHCb production jobs inside Singularity using CernVM root images from /cvmfs
- SingularityCE for DIRAC is being developed by GridPP
- We believe we will be able to run without OverlayFS and without setuid Singularity
- We do not need (or want) to run with custom images rather than the CernVM root images in /cvmfs