


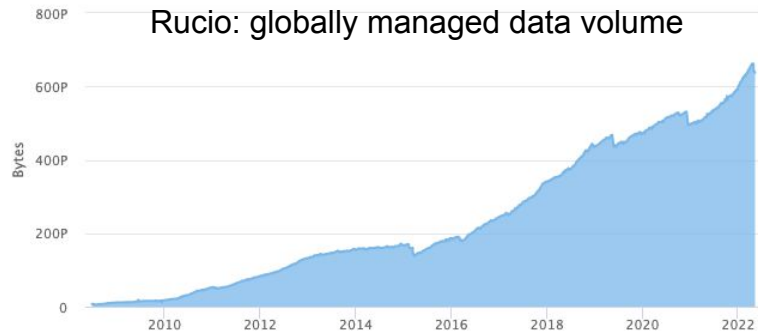
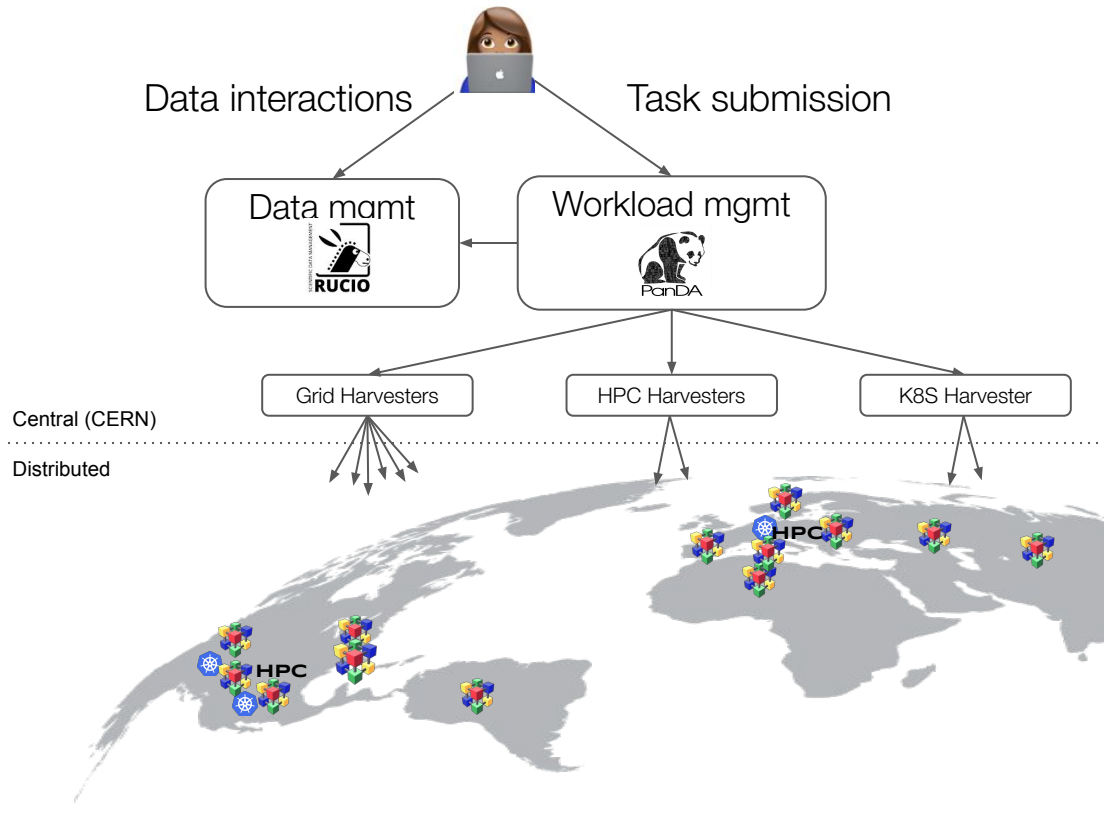
ATLAS Cloud R&D, Kubernetes and CVMFS

Fernando Barreiro Megino
CVMFS Workshop
13 September 2022

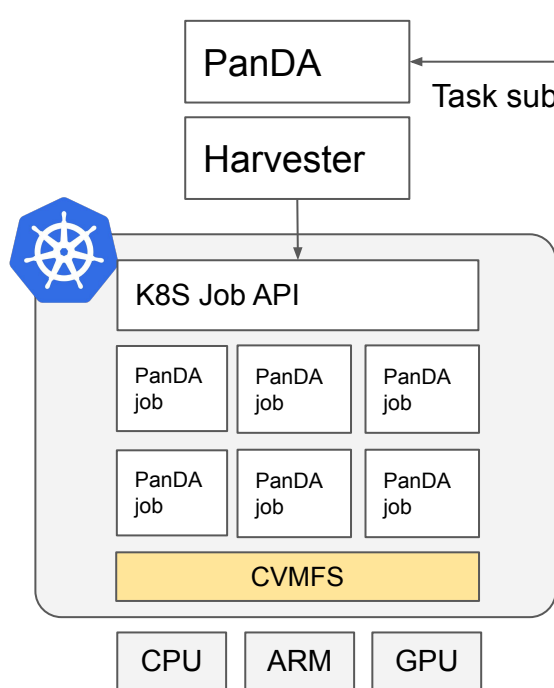


UNIVERSITY OF
TEXAS
ARLINGTON

ATLAS Distributed Computing



Harvester-Kubernetes integration for batch




- Integrate K8S to run jobs like any Grid site
- Keep setup as simple as possible
- Work on any K8S cluster
- K8S Job controller
- Harvester plug-ins for job submission & monitoring
- Usage of K8S options for scheduling customization
- Only dependencies:
 - CVMFS
 - Optionally Frontier Squid



Batch processing

K8S queues





Mini K8S Grid scaled from few hundred vCPUs to many thousands vCPUs




Integration of heterogeneous architectures

- Cloud queues backed by resources not commonly available on prem
- Straightforward to offer different architectures, e.g. ARM, GPU

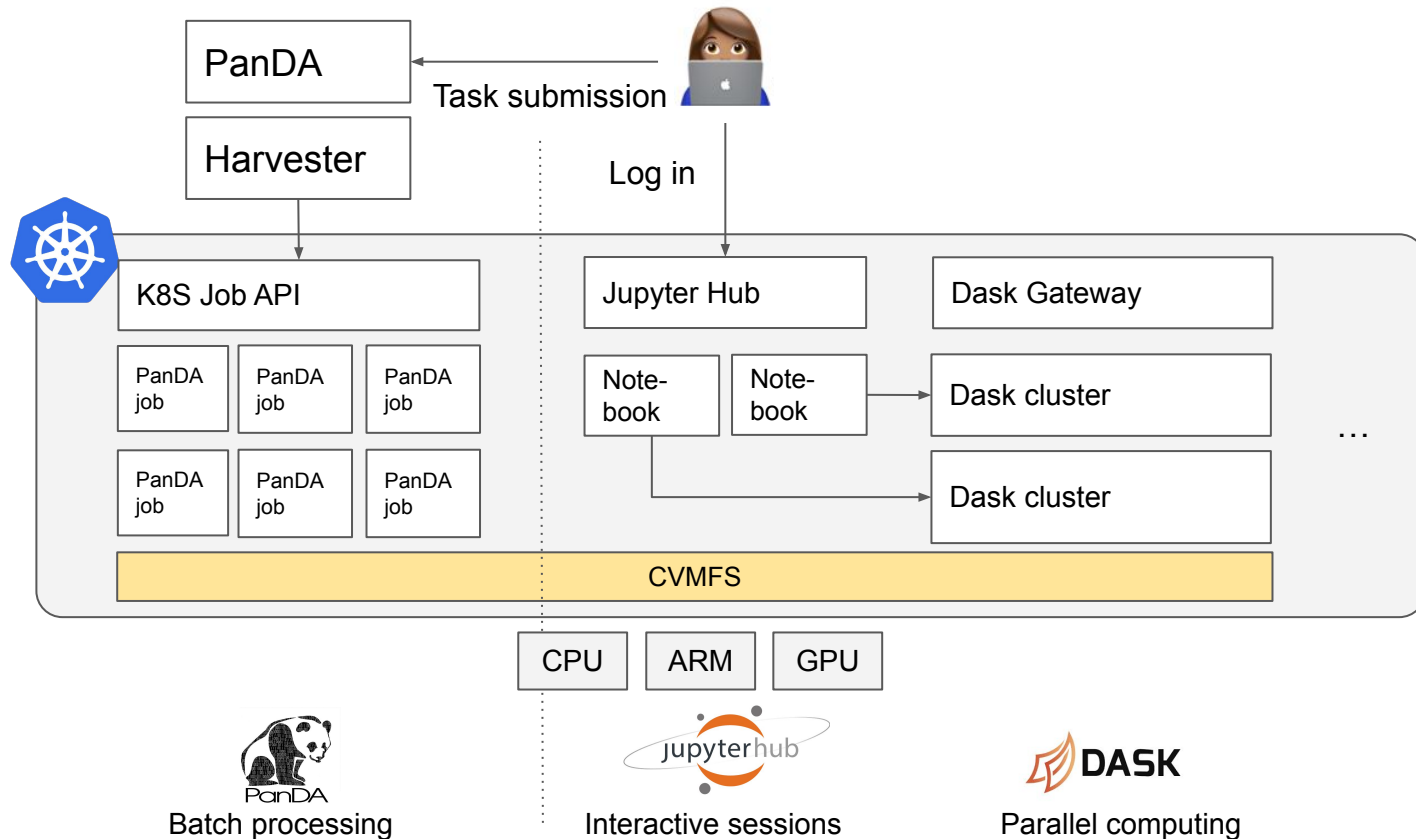


Plots by Johannes Elmsheuser



- First ATLAS simulation task on ARM processors with 10k events
- 1M events, official physics validation being signed off: [ATLPHYSVAL-872](#)
- Multi-arch Docker images for host OS and CVMFS K8S plugin
- Parts of ATLAS SW and grid MW compiled for ARM processors

K8S as a substrate for other ATLAS/HEP applications

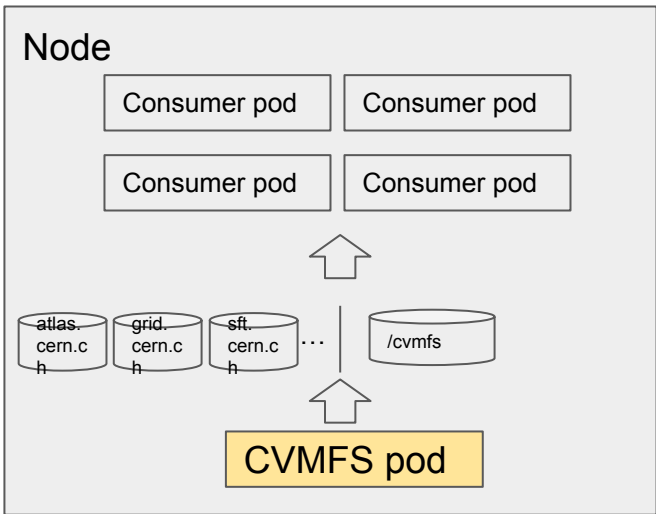


CVMFS installation options for K8S

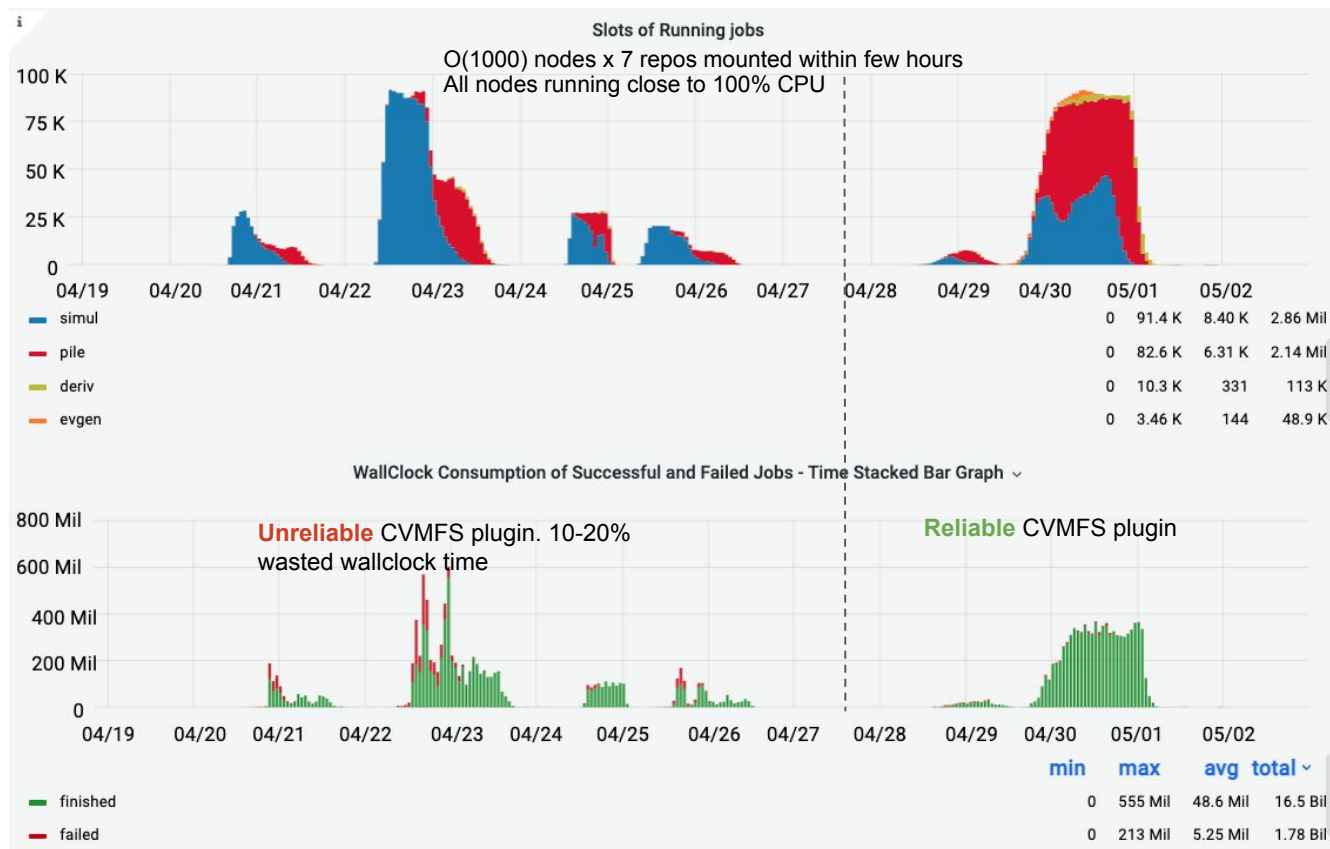
- Traditional installation: package manager installation on host
 - Share cvmfs folders through “hostPath”
 - Simplest method, but not always possible (in particular commercial clouds)
- K8S plugin installation: “daemonsets”
 - CERN CSI driver (v2 on the way): <https://github.com/cernops/cvmfs-csi>
 - CERN Sciencebox driver: <https://github.com/sciencebox/charts/tree/master/cvmfs>
 - PRP OSG driver: <https://github.com/sfiligoi/prp-osg-cvmfs>
 - ATLAS-specific fork: <https://github.com/PanDAWMS/prp-osg-cvmfs>
 - Maybe others

Plugin differences

- Different levels of complexity and generality
- Volumes: csi, local volumes, hostPath
- What to share: /cvmfs vs individual repositories
- Startup and monitoring: go, autofs, supervisord, bash script
- Autohealing: re-mount transparent for consumer vs restarting full pod
- Additional features, e.g. prefetching



CVMFS plugin stress test: fast scale out on GCP



Conclusions

- K8S usage quickly ramping up in the HEP community
- Case for official, tuned and supported CVMFS plugin
- Stepping stones and experience are there now
- Similar situation for Frontier Squid chart

Questions?