Container technology for HEPCloud

Panagiotis Spentzouris, Burt Holzman, Anthony Tiradani

MAGIC monthly meeting

Wednesday March 7, 2018

HEP user community

Many experiments and projects, both nationally and internationally distributed collaborations, 1000's of users

- Fermilab currently supports ~15 experiments and projects, largest experiment (CMS) ~3000 users, largest project (DUNE) ~1000 users, both international
  - Although HEP does have a "standard" software stack and Fermilab provides many software layers as common services, experiments lock on different versions for production and have experiment (or even user) specific libraries and environment for production and analysis
    - And, in most cases, limited resources or expertise to build on them on a multitude of different platforms.

HEPCloud

- HEP experiments need massive computing resources (e.g. CMS ~150k cores internationally)
  - Demands expected to increase x10 (or more) in the next decade
- Resource utilization pattern exhibits peaks and valleys
  - Driven by the science program, detector operations, …
- Elastic provisioning of resources a must, both for cost effectiveness and maximizing scientific output in timely fashion
- HEPCloud a concept for elastically expand the Fermilab facility
  - With the overall goal of offering a dependable set of services for automated, dynamic, on the-fly acquisition of heterogeneous computing resources which treats grid, cloud, and HPC resources uniformly.
  - Could be utilized by other SC programs with similar needs

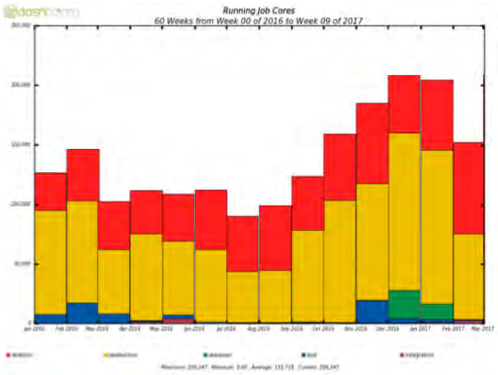

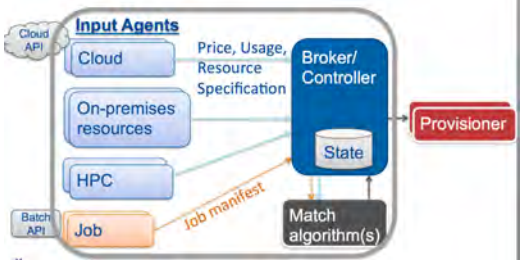

Fermilab

HEPCloud Architecture

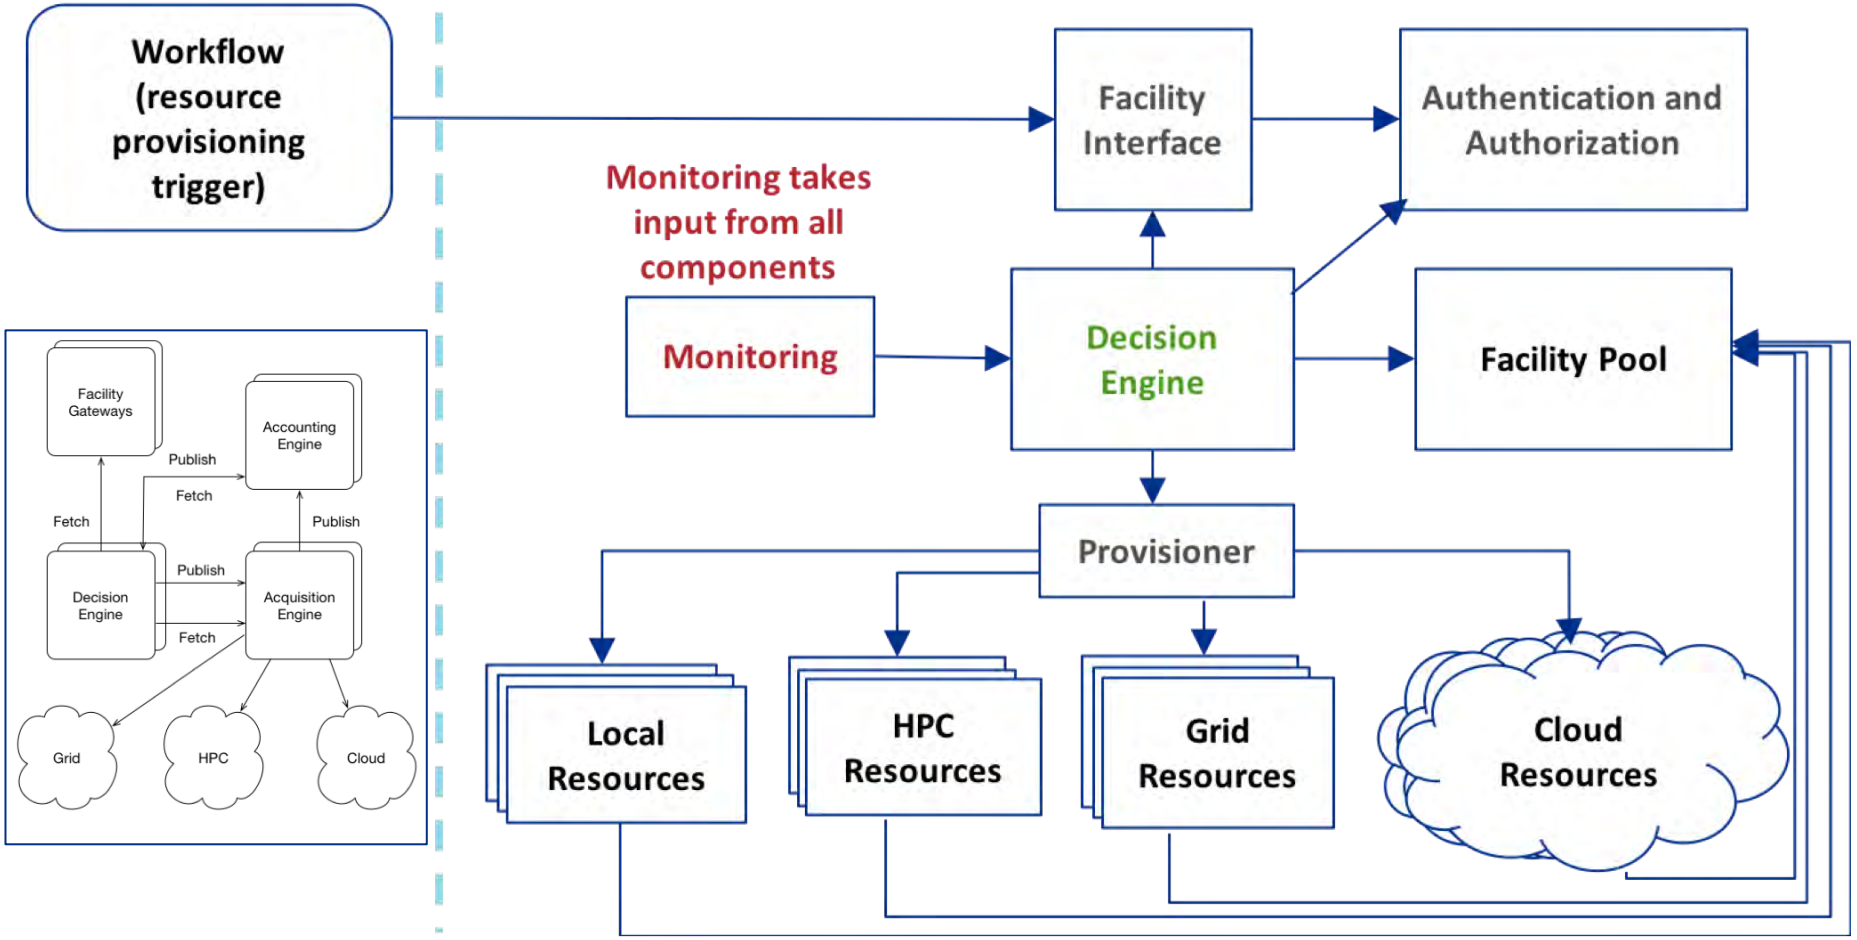

Provides access to heterogeneous computing resources, enables elastic provisioning of resources. Insert shows a high level view including the Acquisition and Accounting Engines.

- HEPCloud Fermilab deployment
  - Project phase, operations planned for end of CY2018
    - Already completed running impactful use cases from different experiments on different resources
      - AWS, Google, NERSC

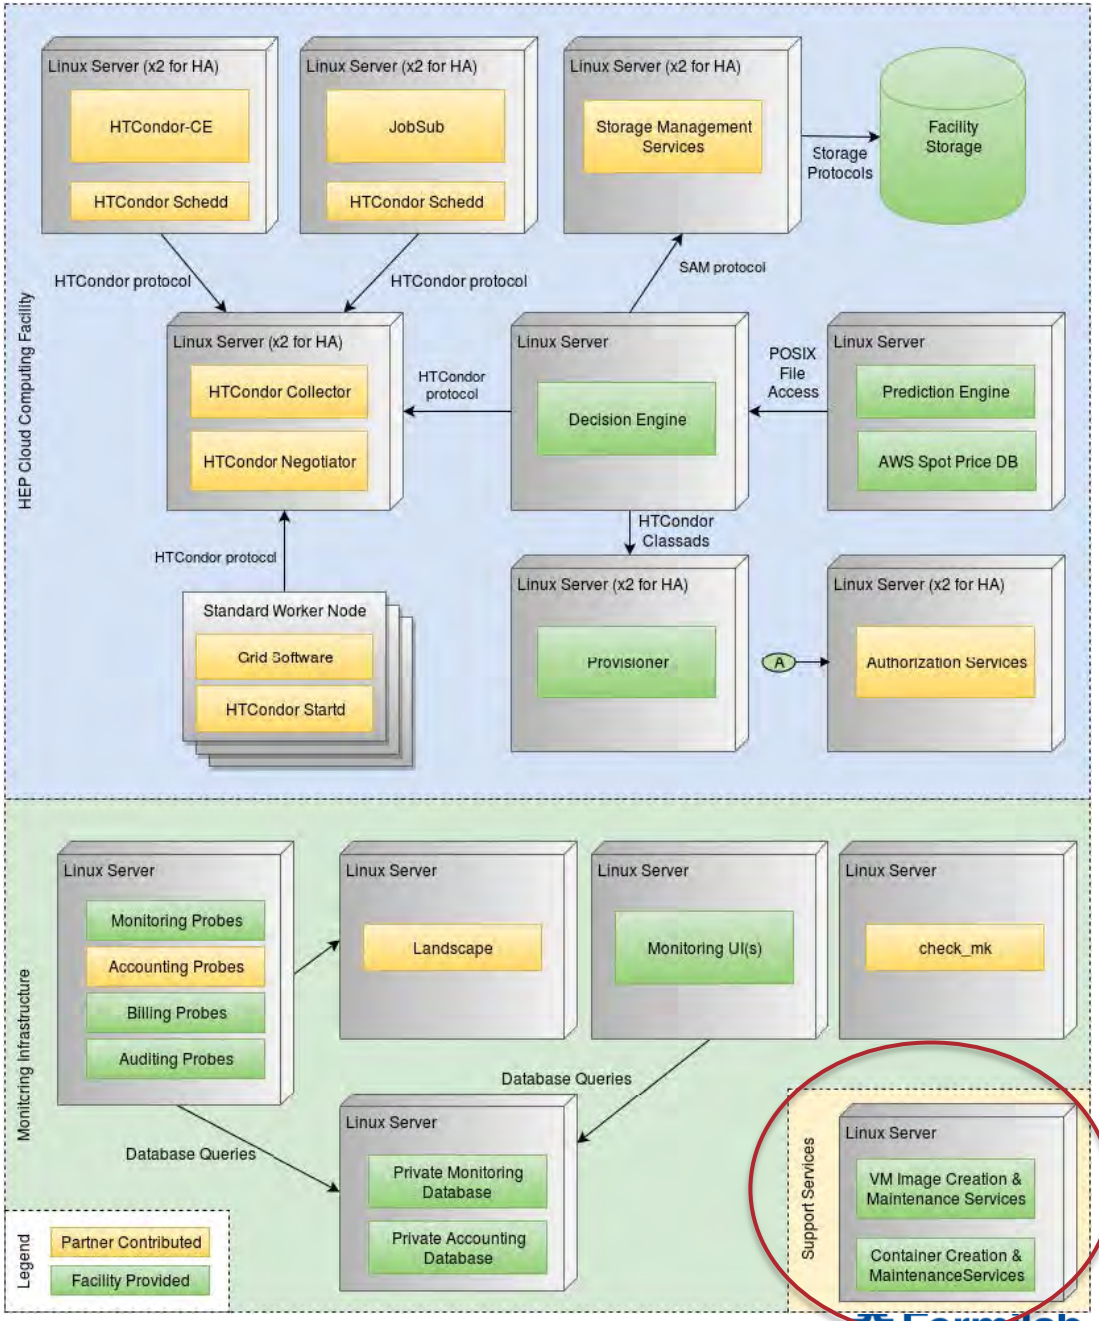

HEPCloud Decision Project Use cases

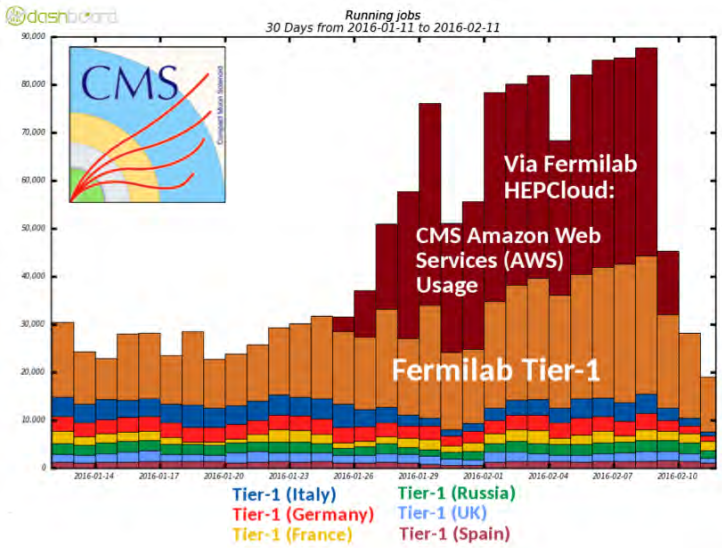

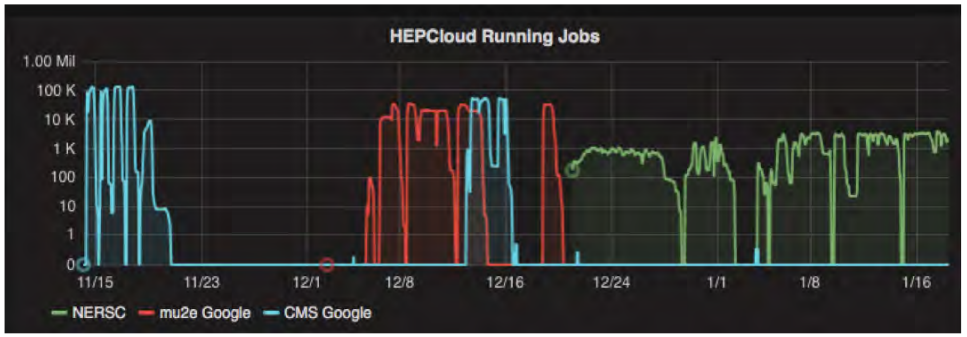

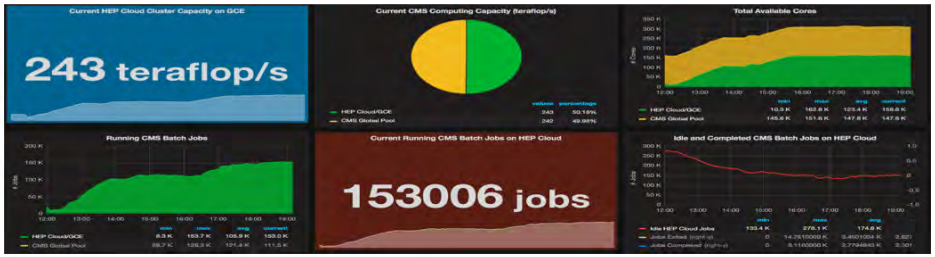

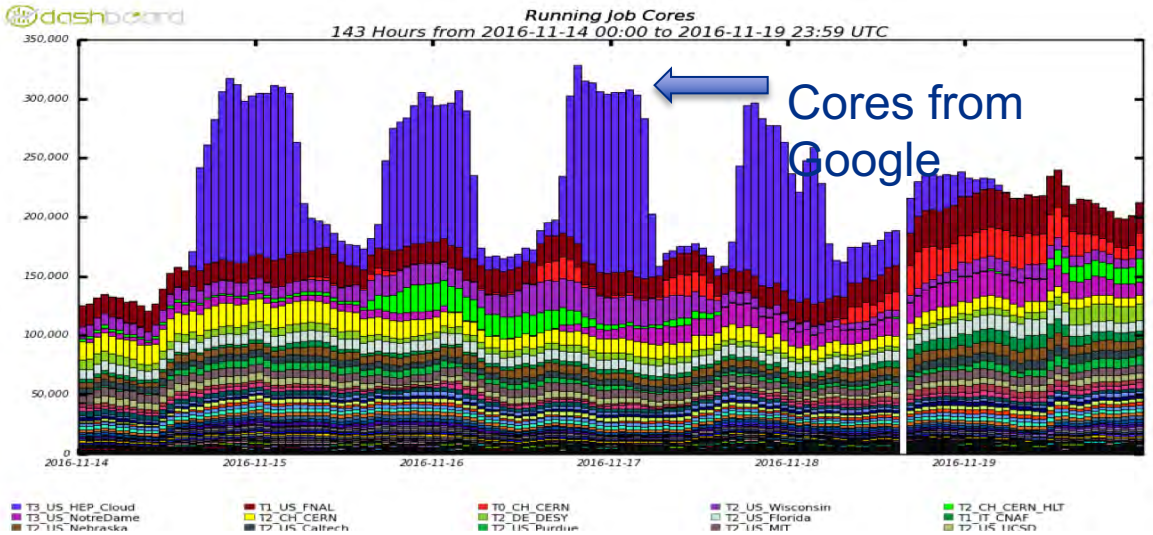

Fermilab

Containers in HEPCloud

- Containers necessary to decouple user environment from the facility resource environment (local –owned– or acquired)
- Current deployment is driven by security requirements of the different resources, user needs, availability of services (e.g. CVMFS) on the different resources, and HEPCloud support and distribution service maturity and capabilities.
  - Current and near term future compute resource targets
    - Fermilab local Grid (FermiGrid)
    - OSG
    - Commercial Cloud (AWS, Google, …)
    - NERSC
    - XSEDE
    - LCFs (?)

Fermilab ”local” resources

We provide a short menu of options for docker containers, consulting for singularity

CVMFS (read-only file system in user space, hosted on standard web services) used as software distribution service.

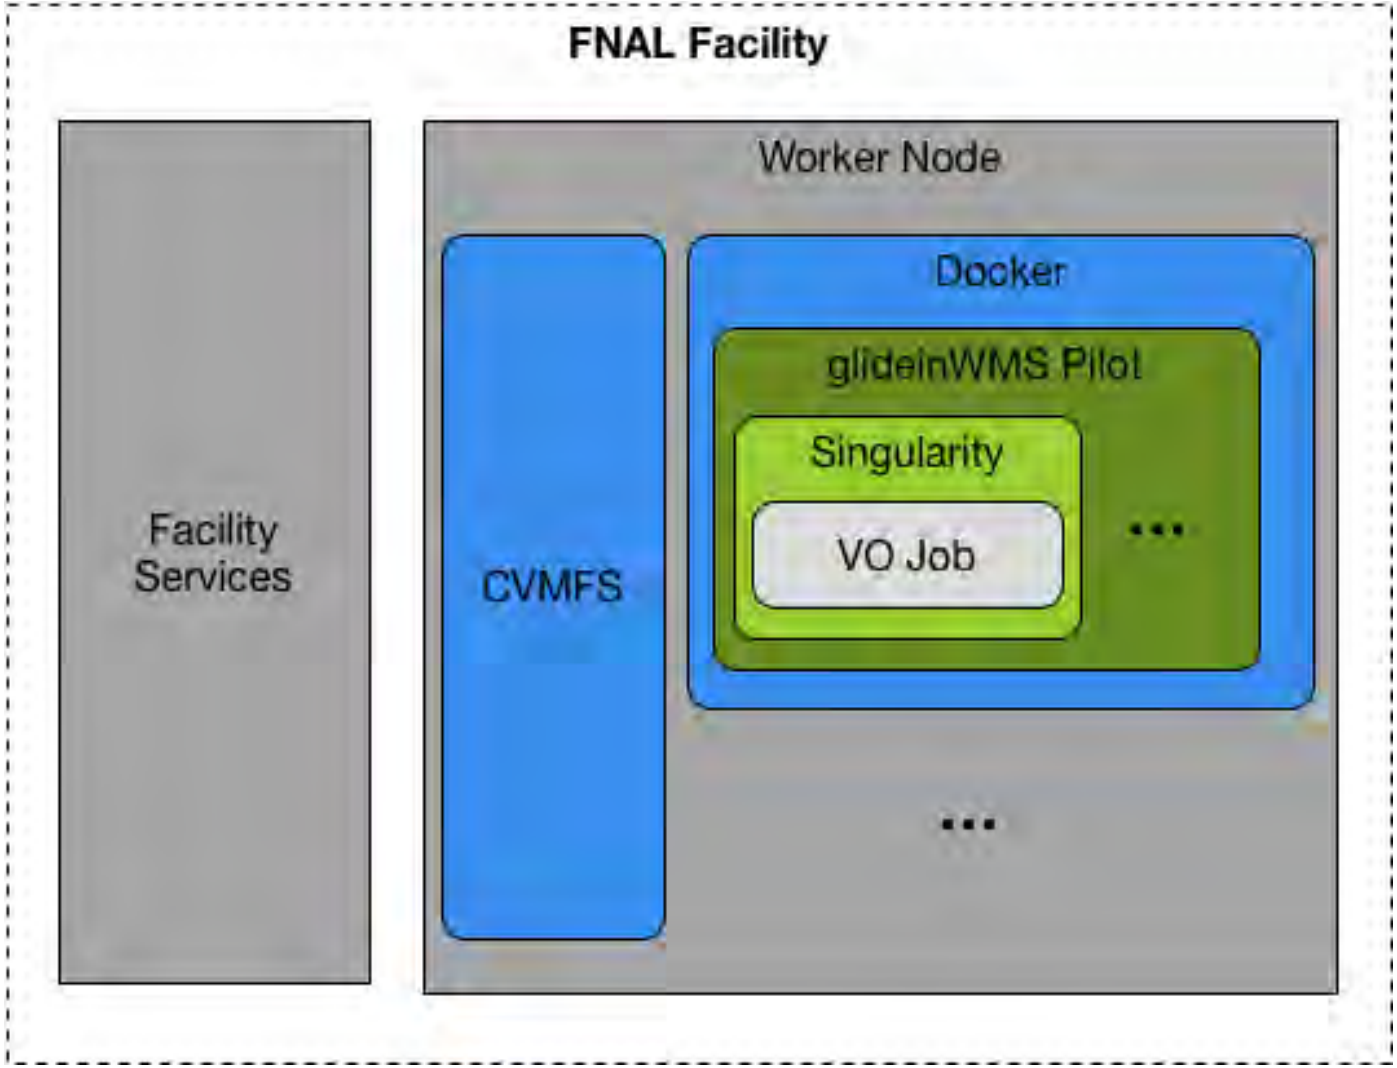

Commercial clouds

Since HEPCloud started work with AWS, AWS added support for Docker and Windows containers via the Elastic Container Service.

AWS also support Kubernetes as a service.

HEPCloud is investigating these services.

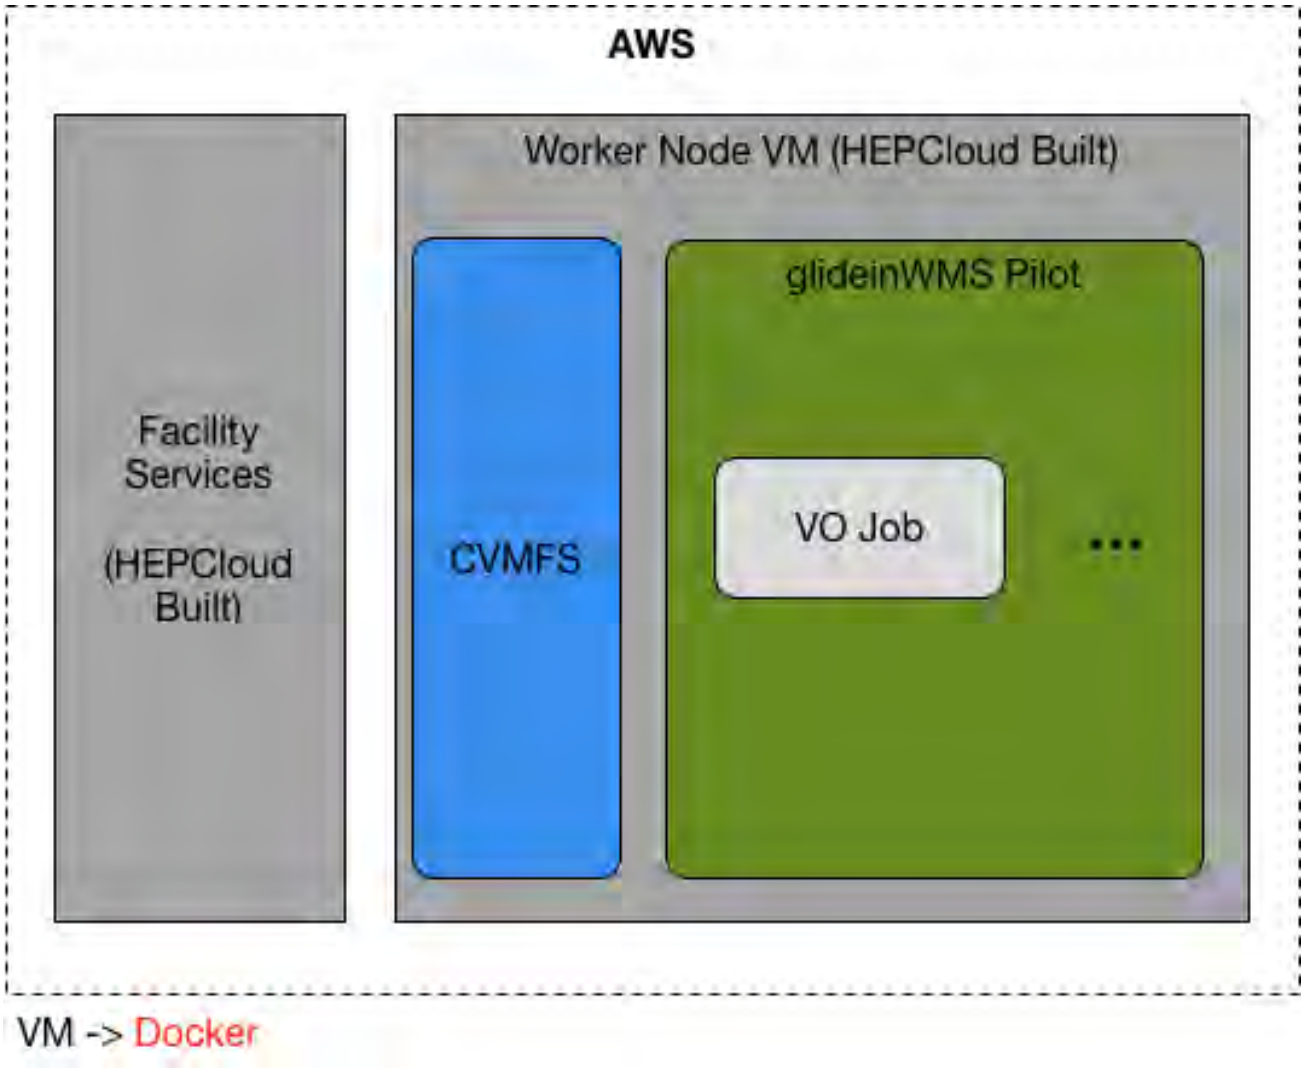

Fermilab

NERSC

- No CVMFS (policy)
- We send docker image to NERSC to be converted to shifter
- Currently edge services Squid services to deal with network load, expanding scope might assist with software distribution

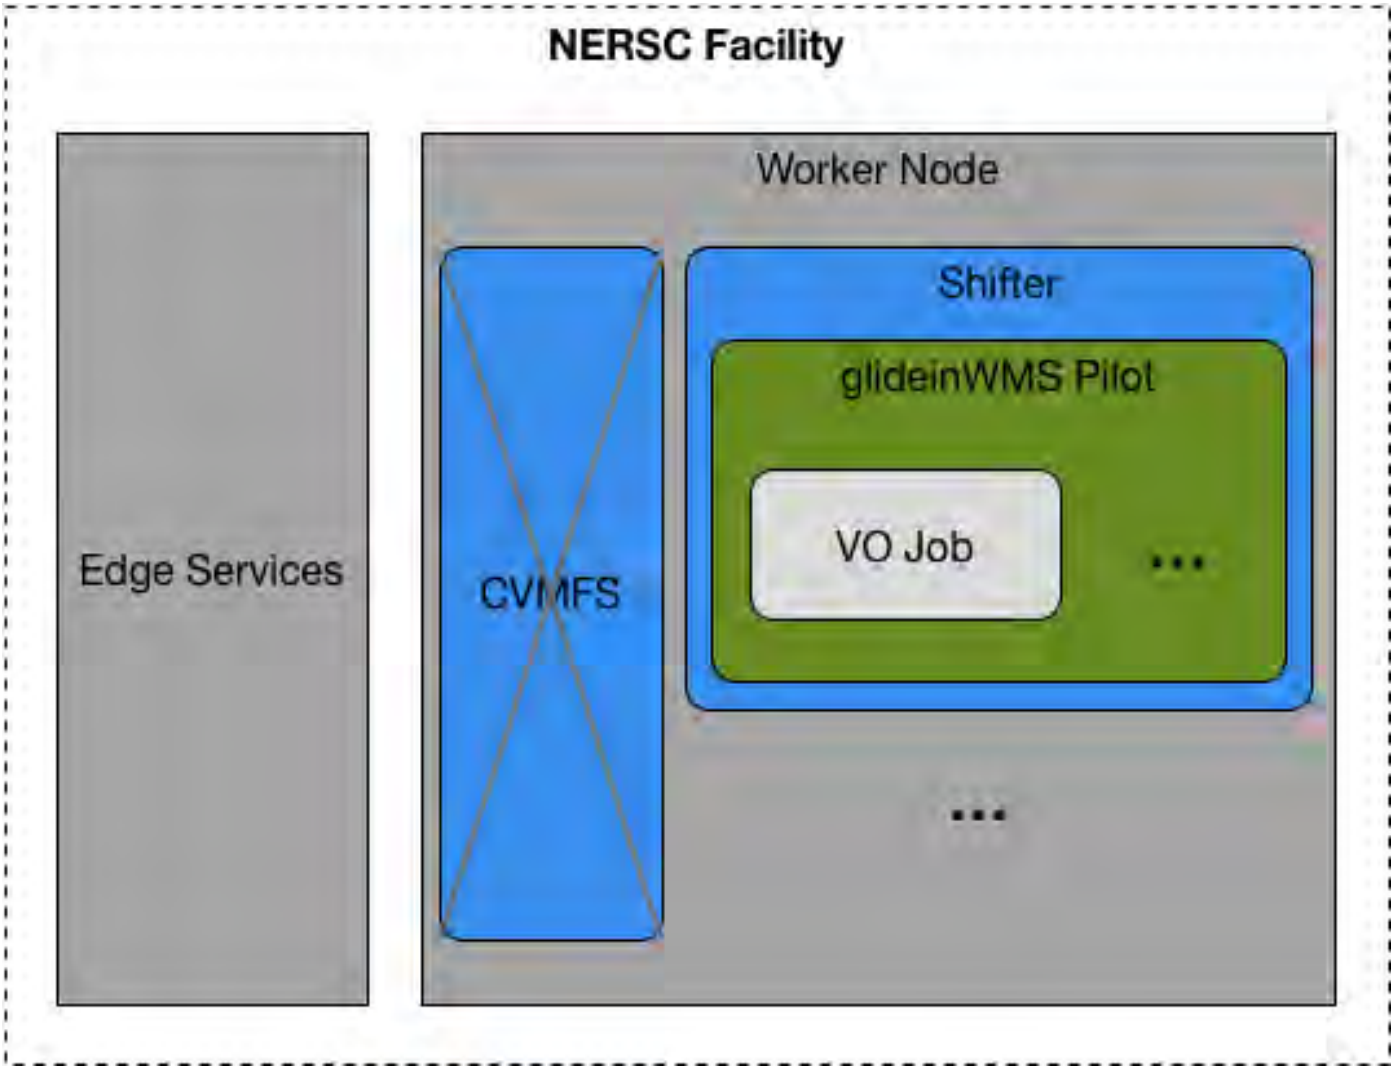

LCF (possible solution?)

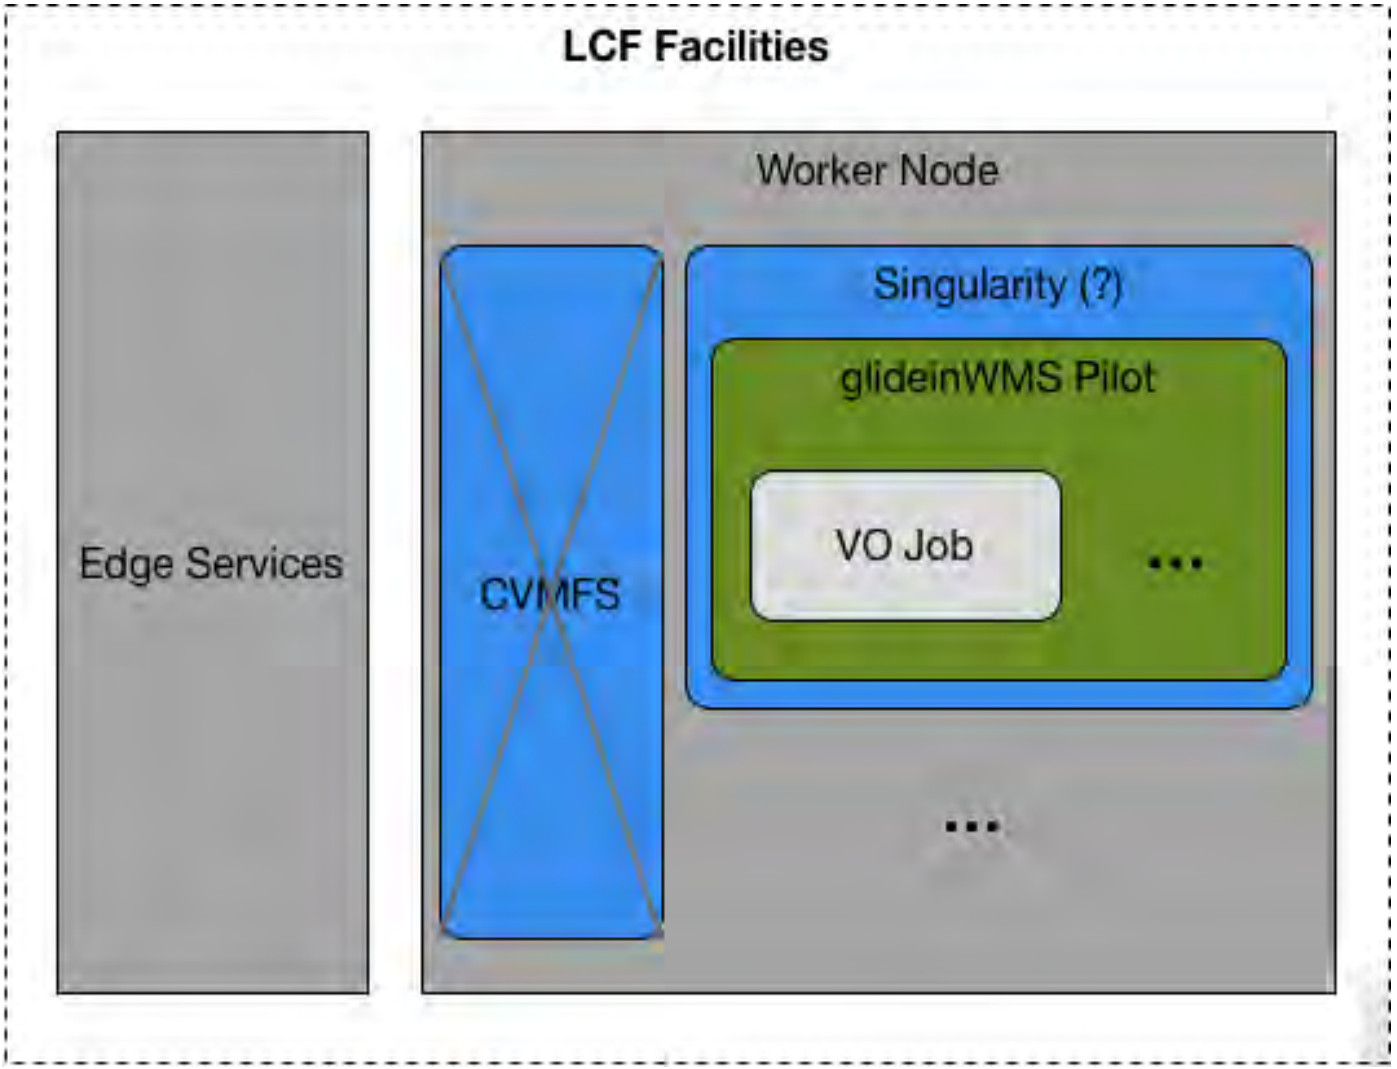

Implementation needs moving forward

- Different “resource providers” use different technologies and different configurations for the same technology
  - Rule uniformity is desirable or tools/services to convert between technologies transparently with minimum overhead
- HEPCloud still needs to fully develop
  - Build system services (user level)
  - Distribution service (internal to HEPCloud)
  - Administration/Configuration services (for the provisioning layer)
- Investigating tools/services provided by the different container technologies to help achieve this goal.

Extras

HEPCloud Decision Engine

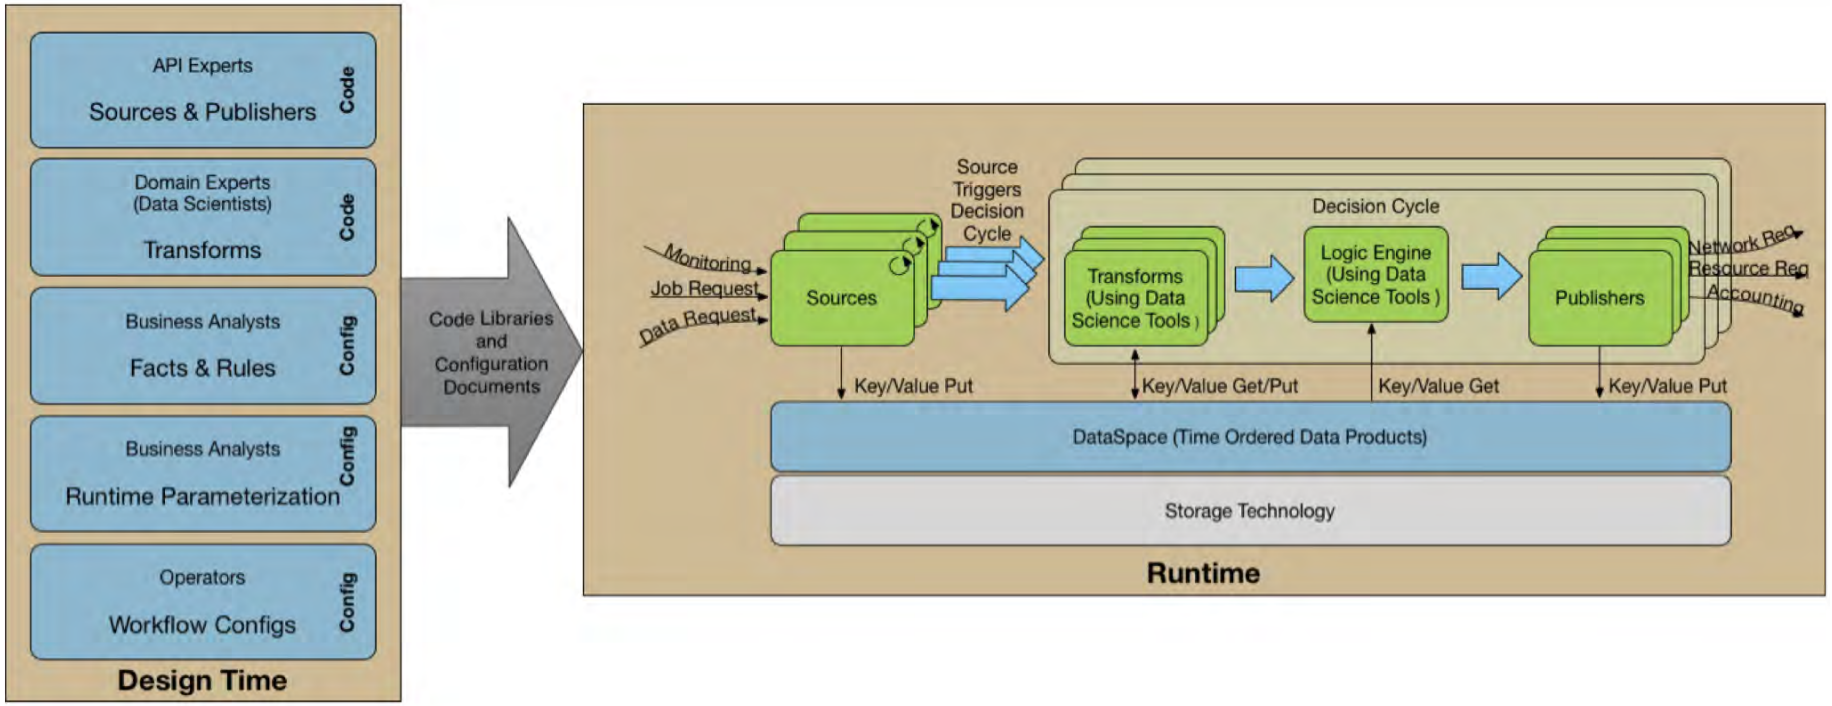

The decision engine: policy, rules, financial information, goals, availability (assisted by monitoring layer) provides "intelligent" elasticity to maximize timely scientific output, efficiency and cost effectiveness.

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov