

---

# HEPCloud

Project Outline

A.Norman, K. Majewski, T. Tiradani

## HEPCloud Status

---

- Phases 1-3 of original plan are complete
- Officially entered into Operations in Spring '19
  - Supported Resource sets:
    - Grid computing sites (FNAL, OSG)
    - Commercial Cloud (AWS, Google)
    - NERSC Cori
  - Supports "pressure" based resource allocation
  - Supports only "single node" piloting model (i.e. not MPI based workflows)
- In use by:
  - CMS (w/ global pool integration)
  - NOvA (demonstrators), DUNE (onboarding), Mu2e (demonstrators)

## Strategic Direction

---

- Shift efforts towards Leadership Computing Facilities (Argonne, Oak Ridge)
  - Involves special integration work at each facility to match their actual setup
  - Requires edge services or other gateway services to bridge boundary layers
- Build support for heterogeneous resource requests (i.e. CPU+GPUs)
  - Driven by ML applications and GPU enabled code bases
  - Overlaps hardware roadmaps for LCFs and NERSC.
- Build support for complex and multi-node workflows (i.e. MPI enabled topologies, pipelined topologies)
  - Driven by existing MPI applications & move towards data parallel processing models
  - Driven by filtering/analysis pipeline applications (image processing style)
    - May be needed to efficiently “feed” hybrid CPU+GPU workflows (DUNE)

## Technical Direction

---

- Shift efforts towards identified gaps and deficiencies in current decision engine framework
  - Primary deficiencies:
    - Unified configuration systems
    - Unified metrics reporting system and system diagnostics
    - Ingestion/Integration of external monitoring (a.k.a. Landscape)
- Refactoring of decision channels and modules
  - Develop coherent module design
  - Reimplement core decision paths with conformation to interface specifications
  - Develop decision channels for HPC workflows and sites under new model
    - Support for "transactional" provisioning instead of pressure based provisioning
    - Normalize model for HPC site integration to include edge service layers
      - Want sites to look "similar" if possible or use similar resource allocation models

## Site Integration

---

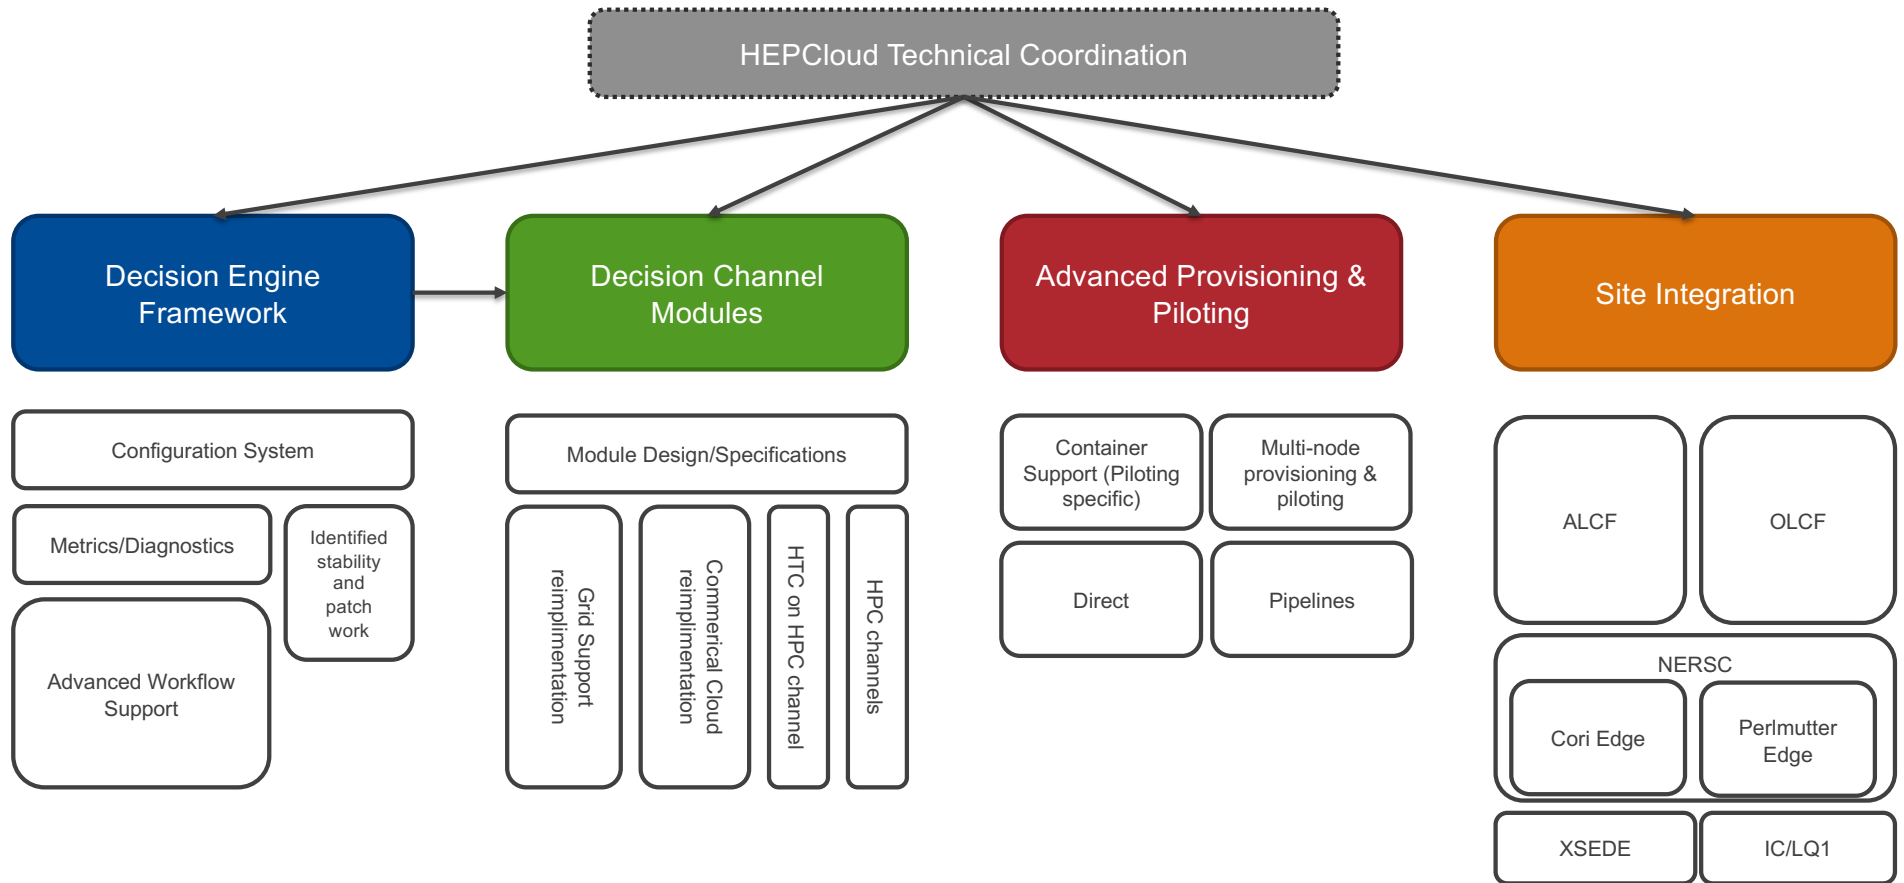
- Focused integration on Argonne (Theta)
- Exploratory work on integration w/ Oak Ridge (Summit)
- Development of edge services layer
  - Apply to both NERSC and ALCF
- Focused integration on NERSC-9 (Perlmutter) [Q3 2020]
- Includes exploratory work on integration of Institutional Cluster (LQ1 as first stage?)
- Provisioning infrastructure development/improvement (overlap w/ Condor)
- Piloting development/improvements (overlap w/ GlideinWMS)
  - Separation of frontend/decision engine from piloting infrastructure

## Operational Direction

---

- Refactor build/test/package/deployment methodology
  - Goal is to develop a maintainable product suite that can also be deployed external to core FNAL campus
  - Open development pathways to broader community
- Refactor configuration to be maintainable (see framework config task)
- Integrate metrics and diagnostics into operational footprint (feedback loop)

# View of Project



## Current State

---

- Signed charge (Sept '19).
  - *Available on the HEPCloud sharepoint (Phase-4 Documentation).*
  - Emphasizes focus shift to HPC/LCF resources.
    - Retains capabilities for cloud usage (reprioritizes ongoing integration w/ one provider)
  - Focuses development on design & refactor of the Decision Engine and channels
  - Focuses integration designs on HPC environments (i.e. restricted networking, etc...)
  - Expands capabilities to support multi-node workflows (emphasis on DUNE/LSST)
  - Expands support for heterogeneous resources (GPUs)
  - Maintains support near term for current operations release
- Draft Charter in Preparation (now)
  - WBS and Project execution plan to follow
  - Resource estimates when project plan and WBS established.