HEPCloud

Fermilab

U.S. DEPARTMENT OF ENERGY | Office of Science

HEPCloud Facility Roadmap

HEPCloud Leadership Team

HEPCloud Stakeholders Meeting

June 2020

HEPCloud Road Map

HEPCloud

You will hear in detail today about:

- HEPCloud Operations
- HEPCloud Software Development
- Preliminary/Proposed Roadmap for HEPCloud features and integration activities

HEPCloud Roadmap

The HEPCloud operations and project has been working to plan out the major development and thrusts for activity that will occur in the near and projected future.

We want feedback on:

- HEPCloud Software Release Schedule
  - Impacts the feature sets that are available to the community
- Integration with new facilities
  - Tried to align with availability of new DoE computing resources
- Support of different workflow models
  - Determines if HEPCloud can be used for major experiment analysis campaigns

Release Schedules

- We want to make sure that HEPCloud is updated on a regular schedule and that the community has input into the features that are included in each release.
  - We also want to make sure that there is sufficient time to package, test, vet and deploy new releases.
  - We are proposing a quarterly release schedule
    - This is different from the current release methodology which is driven by feature completion
    - This will allow us to prioritize and schedule features that the community needs and provide a well defined environment for experiments to work against
    - Patch release and other “bug fixes” would still roll out on an as needed basis

HEPCloud Release RoadMap

- Operations Releases

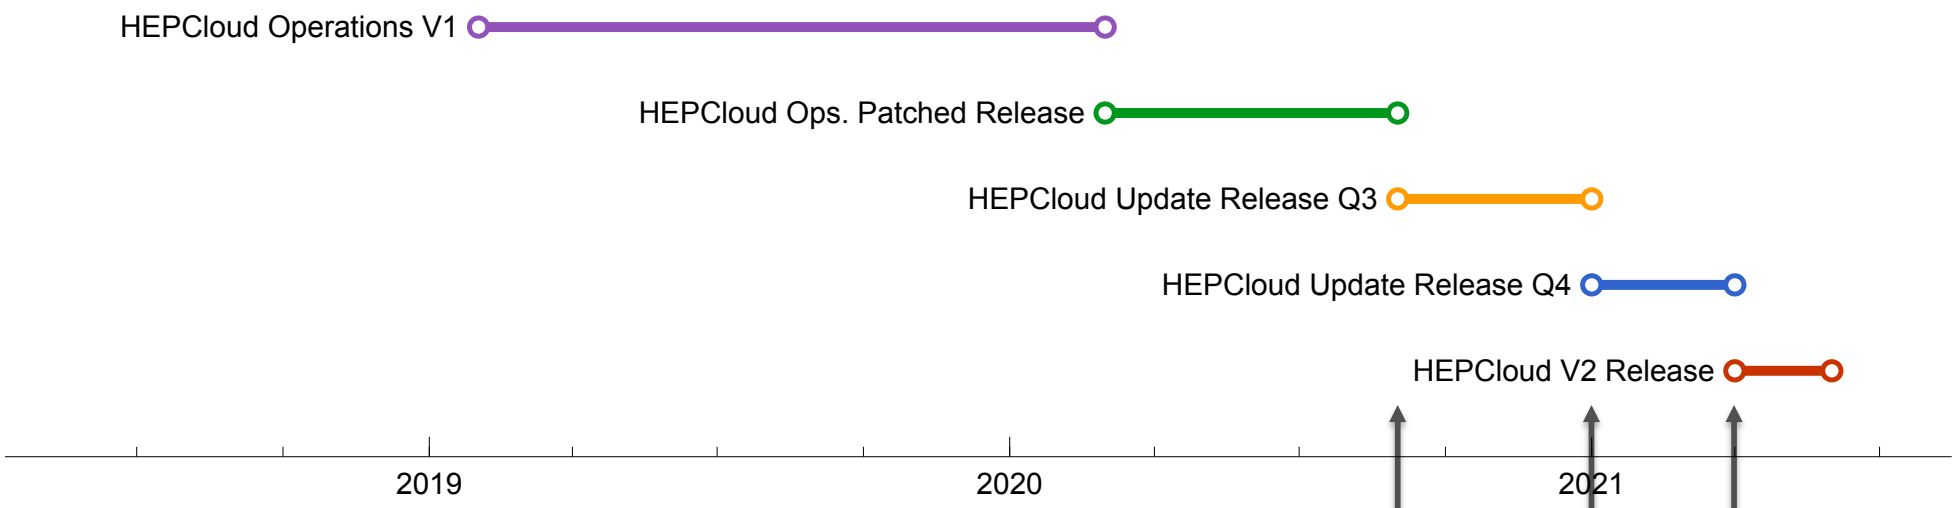

Updated decision engine and baseline channel logic for HTC. Configuration system and diagnostics reporting mechanisms

Transactional provisioning, channels for HPC workflows, expanded metric reporting

Accelerated/Advanced workflows support (i.e. pipeline toplogies with GPU provisioning)

Facility Integration

- Our highest priority is the integration of Argonne Theta to HEPCloud
  - This is a facility similar in scope to NERSC's Cori but with a radically different external networking model
  - Figuring out this environment and integrating is essential to onboarding the upcoming Aurora machine and other similarly isolated HPC machines

- TACC is fully integrated
  - You will hear more from Steve T. about that.

Facility Support Roadmap

- Supported Sites and Resources

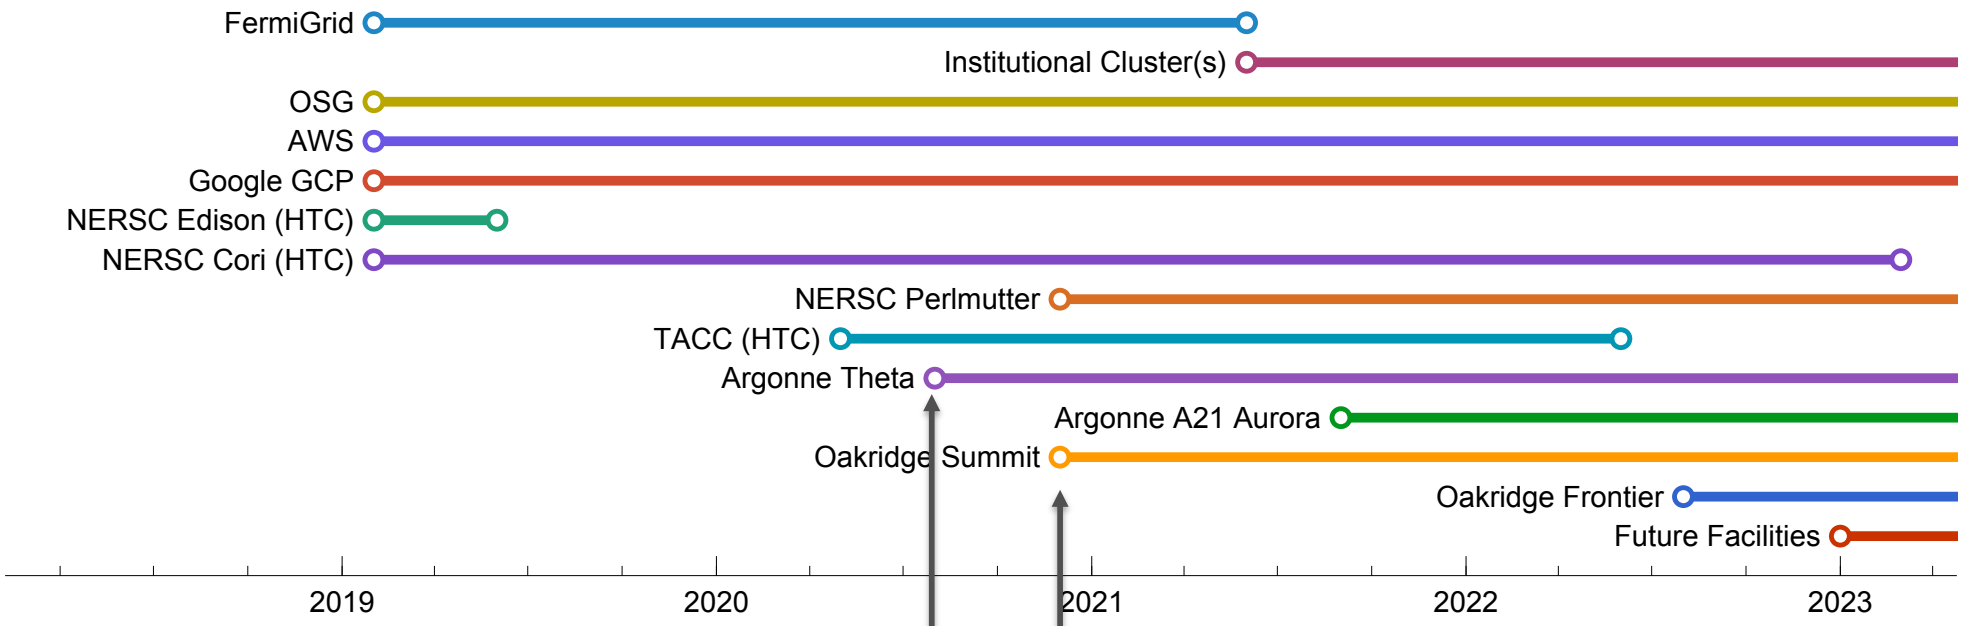

Requires ”edge services” implemented by ALCF to bridge network isoloation at Theta

Assumes Perlmutter access similar to Cori. Requires early access for integration (NESAP)

Critical Path Items

- Integration and use of Argonne's Theta system (KNL based) requires an "edge services" infrastructure be in place.
  - Required due to network isolation of the Theta machine
  - Communication in and out needs to be brokered
    - Prevents the use of "standard" Condor CE model
    - Projects underway through LHC experiment efforts to use filesystem as broker
  - HEPCloud requires some brokering system to allow:
    - provisioning through the Cobalt batch system
    - pilots to communicate out to Condor

- Argonne is working towards a Kubernetes infrastructure
  - This would allow for these types of edge services to run
  - Estimate from Argonne was Aug 2020
    - Recent estimate does not make this target

Workflow Topology Support

- Currently only the ”pressure” based provisioning model is in operations.
  - This has draw backs when combined with long latency provisioning times we see at facilities like NERSC
    - By the time the resource is provisioned, the jobs have been drained from the queue by another site
  - Move to a slightly different model that claims the jobs differently to deal with these latencies
- Neutrino science uses MPI for their analysis and fitting
  - Currently this is only possible through direct Slurm submissions
    - Working on adding support for this mode of operation prior to the next conference cycle
- Pipeline workflows
  - These are advanced ways of using GPU nodes. These are requested by Neutrino community to support LAr analysis

Workflow Support Roadmap

- Evolutions of Worflow Support

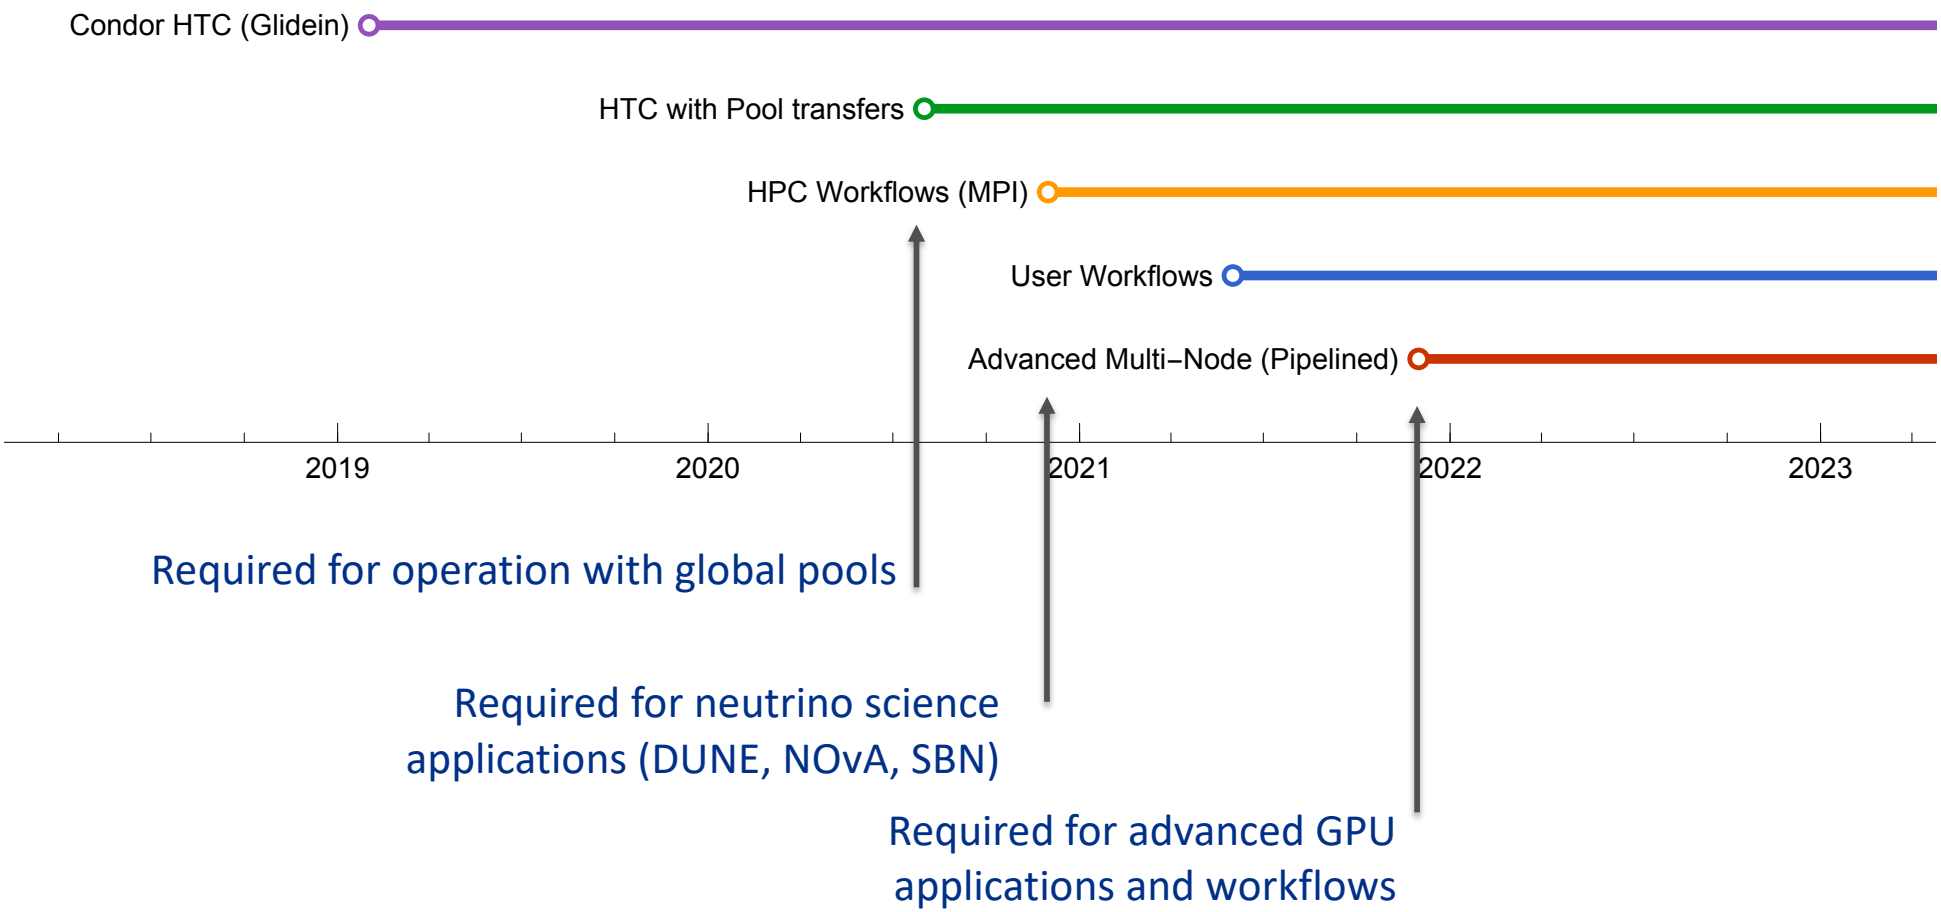

Off Road Items

There are number of items which are not included in the original development and integration plans

- Monitoring
  - Internal & Diagnostic Monitoring
  - External site monitoring
  - Summary Monitoring
- Deployment infrastructure and integration with payloads
- New security integration (i.e. token infrastructure?)
- Data storage integration & awareness for provisioning

Next

- We will hear details from Operations and Development

- Then we want to circle back and discuss a number of issues that have come up that we want to have talked about in a larger context

Part II

Requests and Issues

Request Discussion

- Let's here from you!
  - What are the items that you want on the roadmap?