



HEPCloud Project's Meeting

HEPCloud Leadership Team

Projects Meeting

Aug 2020

HEPCloud Notes



- HEP Cloud presented yesterday (12-Aug) at the Monthly HEP Cloud Stakeholders Meeting.
- Focus of stakeholders meeting was reporting on the plans for Theta integration and the progress that has been made
- Meet with Regetti Quantum computing yesterday (12-Aug)
- Working up plan for Glidein + HEP Cloud transition of front end functionality

Highlights:

- **Presented the Madison idea for condor within HPC**
- **Regetti project to move forward with AWS based prototype**
- **Frontier and CVMFS both are working at Theta**
- **Wilson cluster prototype continues to progress**
- **TACC has agreed to us standing up a new Condor-CE based on our security plan**

Phase 4: Framework

- Project focus continues to be on framework development
- **First quarterly release on July 23**
 - Release of the **new configuration system**
 - https://hepcloud.github.io/decisionengine/release_notes/release_notes_1.3.html
 - https://hepcloud.github.io/decisionengine_modules/release_notes/release_notes_1.3.html
- First milestone/release, latest release (1.3.0) still not fully tested
 - **Waiting on python3 glideinWMS**
- Work continues on logging, design of error handling and produces/consumes redesign

Krista M.

Monitoring



- HEP Cloud monitoring requirements document is on the HEP Cloud Facility sharepoint:
 - <https://fermipoint.fnal.gov/project/fnalhcf/Documents/HEPCloud%20Monitoring%20Requirements.docx?d=wce0c43a64aa347f685ac6fc439c8d4d4&csf=1&e=SSwtqx>
- We are reviewing these requirements

Krista M.



Theta Project



- A plan was created and documented a few weeks ago
- Proceeding with the “login node” design plan
 - So no reliance on Kubernetes@ANL
- Plan has 6 main steps:
 - Transfer work from CMS Global Pool to HEPCloud – 100%
 - Transfer and execute work on Wilson Cluster for “proof of concept” – 100%
 - Replicate Wilson Cluster work on Theta – 3%
 - Perform end-to-end test on Theta – 0%
 - Execute production CMS workflow at Theta – 0%
 - Post mortem work – 0%
- Work is currently 3 weeks ahead of schedule

Krista M.



Meeting w/ Madison



- Had a "special" meeting with Miron and Brian B. on the 27th.
- Intent was to do a deep dive on our theta plans
 - Notes and minutes are on indico:
 - <https://indico.fnal.gov/event/44661/>
- There was a focus on communications and project goals
 - Our intent is to setup a template for how we move forward more generally with other HPC environments
 - We want to establish a *standing **point of contact*** with each **organization** and establish a regular cadence for technical communications.
 - Example: For HEP Cloud/Condor the regular Friday meetings are a first step (Need more regular contact for real work).
 - **Let's discuss how we want to structure this (in a minute)**

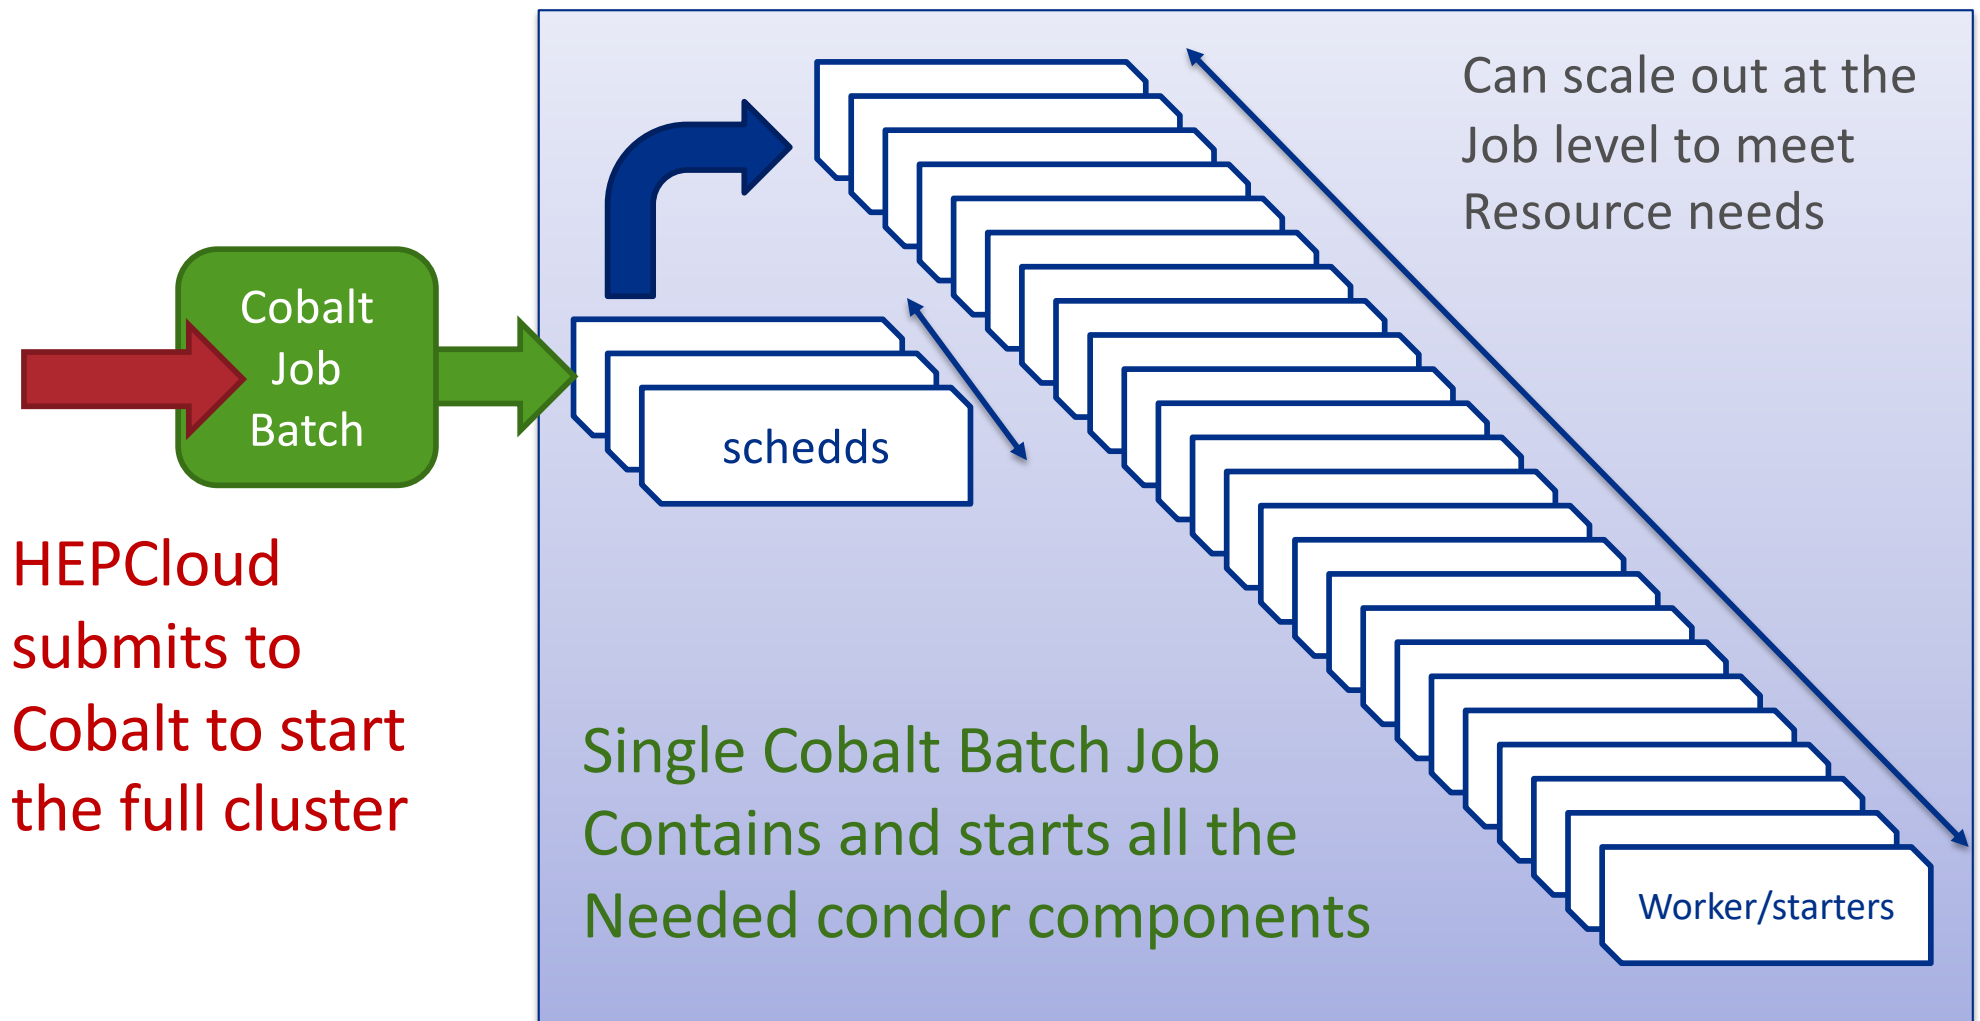
Meeting w/ Madison



- 2nd part of our meeting was technically focused on current plans for HPC integration on network isolated sites
- Big take aways:
 - Our original computing model of trickling jobs to sites is not well matched to this HPC model.
 - Instead we need to think in terms of aggregated campaigns
 - This would let us think about scheduling the ensembles as a whole
 - One proposal is placing the schedd's within the compute environment (not on the edge) and allowing them to scale out as needed
 - Essentially this takes something like “Rank 0” of the campaign being run and places the schedd there and the starters on the other ranks

All within Compute Domain

- My interpretation of the Madison plan



Meeting w/ Madison



- This design would be work that Madison does
- Fits nicely with further integration of HPC environments
 - Reduces complexity of some edge service needs and scaling
- Need to map out what this would look like (time wise) with our current plans to get Theta working with the CMS workflows

HEPCloud + Rigetti



- Met with Rigetti Quantum computing
 - This is a restart of the work that we initially began in Fall '19 that went on hiatus while the quantum centers and statements of work were being established
 - This is back on now
- Baseline design is to provide access to their QPUs as “accelerators” (so classical compute jobs that call out to the QPU for a specific computation)
- There are two infrastructure setups that are proposed
 - #1 uses AWS for hosting the classical compute and batch infrastructure
 - #2 uses computing resources housed in an independent data center “near” Rigetti’s QPUs to reduce latency

HEPCloud Progress



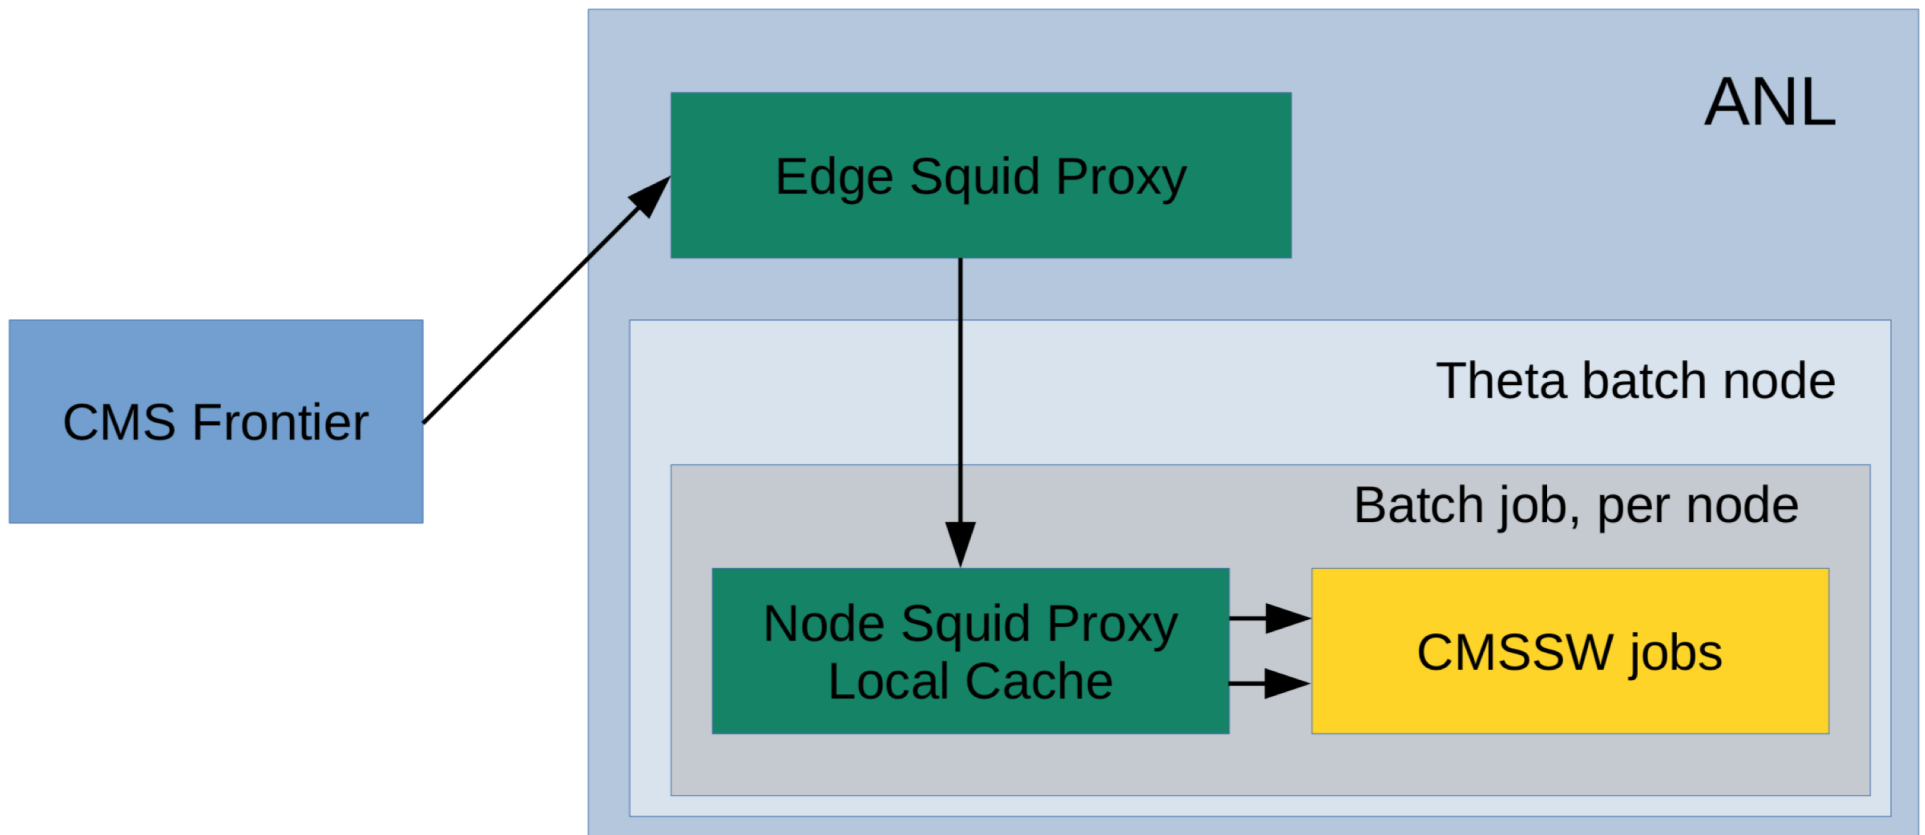
- Software development continues along on the decision engine
- There is a real bottleneck in our python 3 upgrade deployment
 - Dependency on Glideinwms. We need the Glidein upgrade pushed and have been stuck waiting for them for almost 2 months now.
- Kevin R. has gathered requirements on monitoring from experiments for improved monitoring
 - Document should be out ~next week
 - Separately Elizabeth Petit-Bois (GEM summer intern) has been working with Mike Kirby, Ken Herner and Andrew N. on monitoring interfaces (with a focus on DUNE but applicable too all of HEPCloud) and has been making significant progress
- Operations continues to be solid.
 - Mu2e was onboarded and completed a 5M hour campaign with minimal intervention
 - New record for concurrent claimed cores for CMS (@NERSC) was hit
- Plan for Integration of Theta with HEPCloud is solidifying



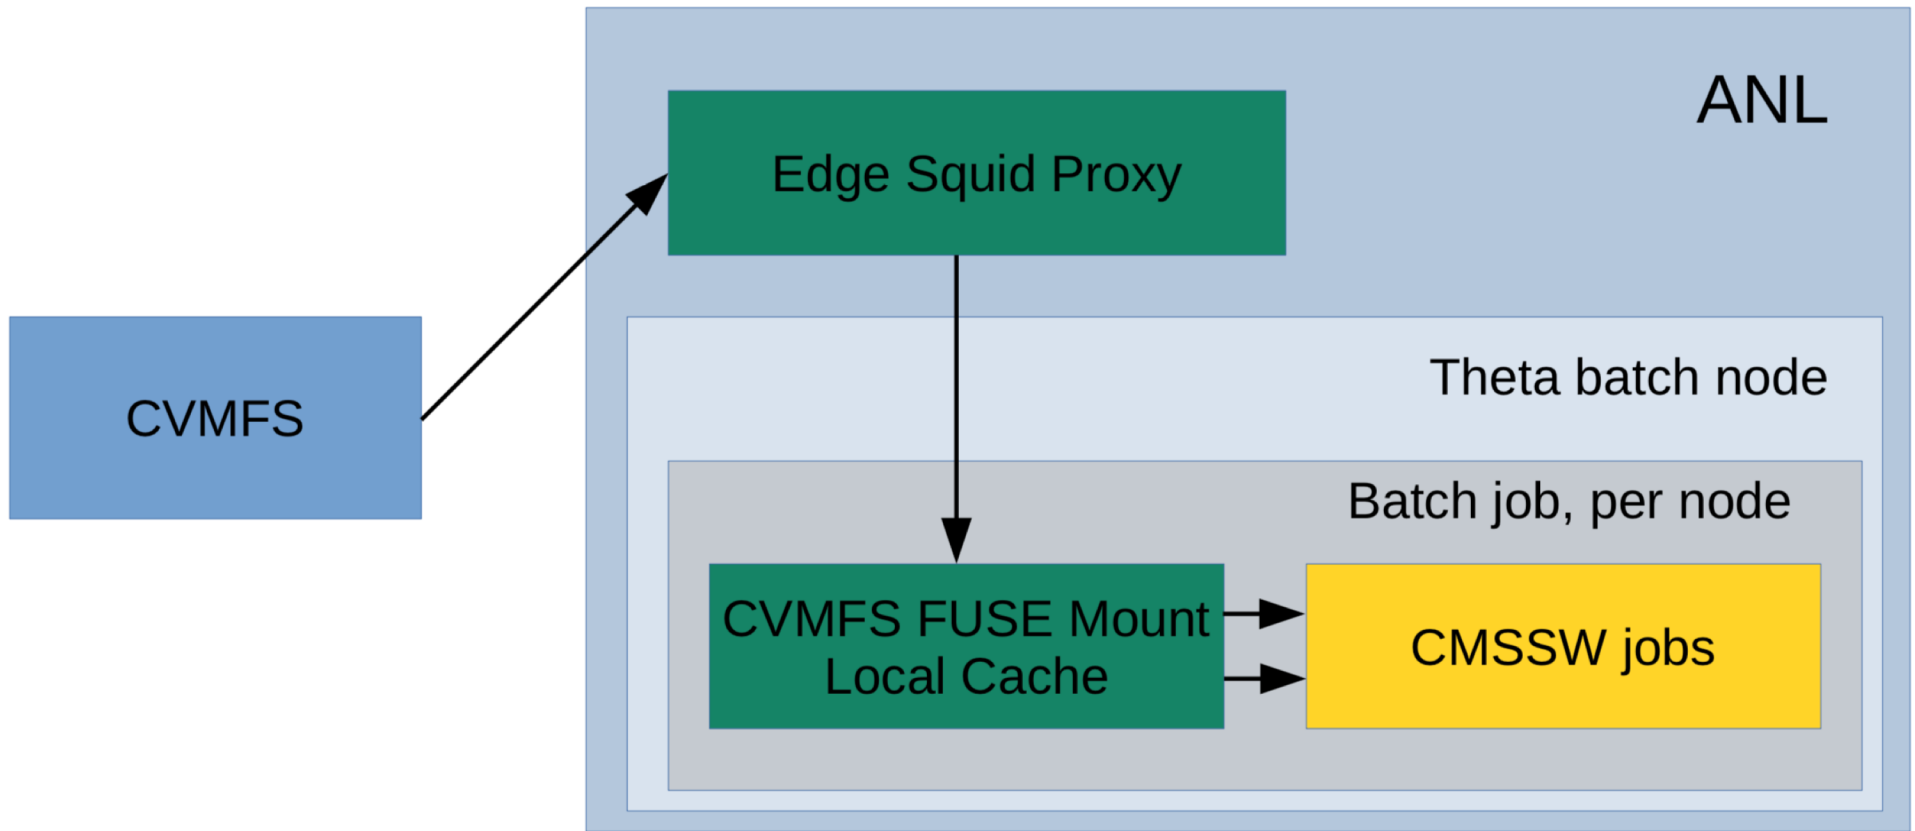
Frontier @ ANL



Frontier



CVMFS



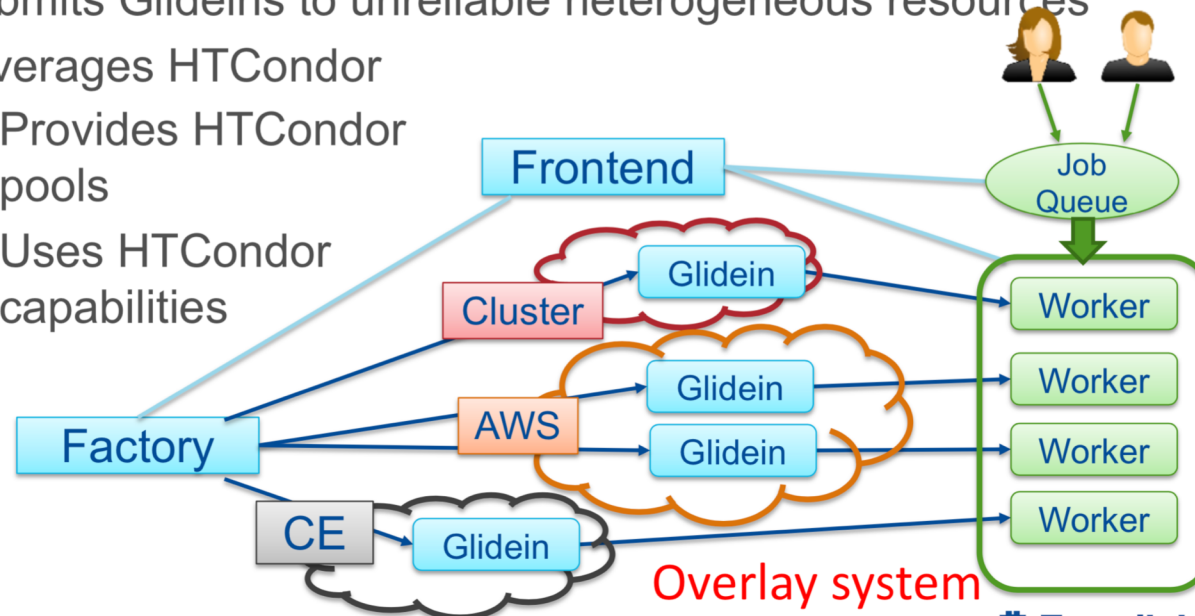
Glidein + HEPCloud

- Presented the plans we discussed last week at the Project Meeting.
- Planning to present at the OSG All hands.

GlideinWMS

GlideinWMS is a pilot based resource provisioning tool for distributed High Throughput Computing

- Provides reliable and uniform virtual clusters
- Submits Glideins to unreliable heterogeneous resources
- Leverages HTCondor
 - Provides HTCondor pools
 - Uses HTCondor capabilities



Dar Be Monster's Here (backup slides)

