Facility R&D Biweekly

WBS 2.3.5.3

Bi-weekly Thursdays, 10:00 am Central
Zoom:
https://uchicago.zoom.us/j/91426535627?pwd=UUExNlkvMGRpbEJudU5OZ05mV080dz09
Facilities Team R&D Google Drive Folder
Facilities Indico category

2024.10.31 - cancelled, virtual updates only

Present:
Apologies:

Agenda / Minutes:

- Quarterly Report (FY24Q4: July-Sep) - QR FY24Q4 - it was submitted.
-
-
-

2024.10.17

Present: Brian, Aidan, Armen, Eduardo, Fengping, John, Judith, Louis, Mark, Ofer, Qiulan, Fred
Apologies: Rob

Agenda / Minutes:

- Quarterly Report (FY24Q4: July-Sep) - Rob
  - Draft for 2.3.5.3 (Facility R&D): QR FY24Q4
  - Please add any additional updates as you see fit.
- Kuantifier news (Brian)
  - Multi-container support put on hold since our focus is on OSG 24 this month. Current state:
    - We need to use a new kubernetes scheduler metric, which seems to require changes at the site (control plane option and Prometheus options) as well as the other code changes

    - Plan to spin up our Prometheus also put on hold as we're working on k8s version upgrades and fixing our local Prometheus stack
        - Using NET2's Prometheus (remotely?) would be a great alternative but we're having issues authenticating remotely
- Updates from UTA (Armen)
    - All running fine, nothing new, busy with S&C Week and followup
    -
- Binderhub demo - impact of improvements on running instances (Fengping)
    - Image build service, polls the repos once a minute looking for updates, builds the image and pushes it to the registry.
    - A separate pod checks for new images in the registry and creates a daemonset to pull the new image
    - on NRP, some exceptions made
- Analysis Assistant - updates and more demo-ing (Ilija)
    - *Skipped this week - Ilija traveling*
    -
    -
- Netbird mesh update, deployment of binderhub on rp1.hl-lhc.io (Aidan)
    - All nodes from previous cluster are on the netbird mesh network
    - Exploring storage options utilizing older Dell R720 nodes
- AOB
    - Next meeting: October 31.

2024.10.03

Present: Fengping (chairing), Ilija, Armen, Brian Lin, Eduardo Bach, Aidan Rosberg, Horst Severini

Apologies: Ofer, Lincoln, Rob

Agenda / Minutes:

- Quarterly Report (FY24Q4: July-Sep) - Rob
    - From Suzanne:
        - Hello: The Jul-Sep 2024 report for the U.S. ATLAS Operations Program Quarterly Status Report will be due on **Friday, October 25, 2024**. Please notify your authors of this date and submit on-line to the U.S. ATLAS Reporting Center found here: https://atlasreporting.bnl.gov/. Also, please update the **Milestones** document for this same period here: https://docs.google.com/spreadsheets/d/1DsMH-16v7bJy6qEkvTEWdLCkfeAXC6VpL019rUpcET8/edit#gid=1636071618

- Draft for 2.3.5.3 (Facility R&D):
  - QR FY24Q4
  - I will go through the minutes and capture activities - please report below anything you would like added to the QR FY24Q4 document

- For **Milestones (last year's doc: FY24 Facility R&D Milestones) Review**:
  - Kuantifier and reporting will need an update. Can this be marked **Complete**?

335	Deploy monitoring, alerting and APEL accounting for UTA k8s cluster using Prometheus	Nov 2023	Sep 2024	Delayed

  - MS 342 was about adding CPU to a multi-site kubernetes resource. This will need to be "Delayed", as we have not yet demonstrated this:

342	Scaled K8S with over-pledge CPU in operation including PanDA backfill capabilities	Sep 2024	Sep 2024	On Schedule

  - MS 429 - Will mark "on schedule"

429	Develop and deploy AI assistant for analysis facilities	May 2025	

  - MS 338 - coming soon:

338	Decision and report on scheduling approach for the stretched K8S platform	Nov 2024	Nov 2024	On Schedule

- Kuantifier news (Brian)
  - NET2: what kind of Prometheus auth configuration?
    - Default used by OKD
    - Cert / token based
  - We are investigating using a kube scheduler metric for support of multi-container pods (including init containers)

        - https://github.com/rptaylor/kapel/issues/71#issuecomment-2384268967
        - Requires some minor option changes on the API server and potential changes to Prometheus scrape configs
        - For our documentation, what kind of Prometheus deployment methods does each site use? CHTC uses kube-prometheus
            - Chicago: kube-prometheus helm
            - NET2: installed via OKD (one per namespace and a central one)
            - SWT2: helm chart (prometheus-community)

- Updates from UTA (Armen)
    - Kuantifier works fine at SWT2
    - Checking accounting on the WLCG side looks correct
    - Looking for the URL for communication with Prometheus to replace IP, which works fine, but just as a more universal solution.
    - For Google queues need to be followed up with Fernando, to include in the OSG topology and then setup accounting. Right now, for the Google accounting, I am adding the numbers in APEL manually.

- Binderhub demo - impact of improvements on running instances (Fengping)
    -

- Analysis Assistant - updates and more demo-ing (Ilija)
    -

- Netbird mesh update, deployment of binderhub on rp1.hl-lhc.io (Aidan)
    - Binderhub now deployed on cluster that is using the netbird mesh network
        - Still need to move a few nodes from the old cluster
        - Very easy to add peer to the netbird mesh network
    - Missing some potentially desirable features
        - Cannot manually assign ips to peers or set subnet range without workaround by editing the database
            - Discussions are open in the github regarding these features

- FY25 - just a heads up to have a look where this work fits in the US ATLAS operations program:
    - 2024 WBS 2.3.5 CIOPS Scrubbing Slides
    - See these slides for new milestones in WBS 2.3.5 (some are specific to Facility R&D)
- AOB
    - Next meeting: October 17, https://indico.cern.ch/event/1445073/

2024.09.19

Present: Rob, Lincoln, Eduardo, Brian, Ofer, Ilija, Aidan, Armen, Judith, Fred, Fengping, Rob, Horst

Apologies:

Agenda / Minutes:

- Ops review close-out remarks relevant to Facility R&D (Rob)
    - The following are some **Comments** from the panel relevant to Facility R&D (and of course Analysis Facilities / Shared Tier3 - WBS 2.3.4):
        - *"US ATLAS is encouraged to proactively discuss with other partners, e.g. ATLAS international, WLCG, US CMS experts to investigate synergies on technical evolution of the facilities, e.g. storage architecture evolution such as ceph and CTA and infrastructure evolution such as Kubernetes, keeping in mind the constantly evolving industry standards"*
        - *Analysis Facilities at SLAC, BNL, and U. of Chicago are all different in setup because the concept is still being developed. Nevertheless, it makes sense to come up with a common minimal set of requirements that define an Analysis Facility within US ATLAS.*
        - *US ATLAS is encouraged to continue collaborating on R&D projects for HL-LHC software and computing with partners such as NSF IRIS-HEP and DOE HEP-CC."*
    - And **Recommendation**:
        - *"Provide a plan for an optimized Analysis Facility concept in the context of international ATLAS leveraging the experience from other LHC experiments, notably CMS. Present status reports regularly at LHC JOG meetings."*
    - Actionable things from this:
        - Determine a **common minimal set of requirements** for analysis facilities
        - Develop a **plan for an optimized Analysis Facility concept**
- AF Chatbot Demo (Ilija)
    - Note we have a milestone for next year on this topic:

429	**Develop and deploy AI assistant for analysis facilities**	**May 2025**

    - Have developed a prototype based on openai, https://af.uchicago.edu/chat
    - Slides are here.
- BinderHub updates (Fengping)
    - Discussion at iris-hep retreat
- NET2 integration/federation update (Lincoln)

    - Have access to create pods at NET2, working with Eduardo. Was able to stand up a simple pod and run iperf between MWT2/NET2
- Kuantifier update (Brian)
    - Encountered an issue at UNL with multi-container pods / pods with initContainers: https://github.com/rptaylor/kapel/issues/71
        - Our ideal solution will have us look at a different metric that is supposedly first-class in 1.27
        - What versions of k8s is everyone running?
            - MWT2
                - AF - v1.25
                - River - v1.29
            - NET2
            - SWT2 - v1.27
    - What auth method does the NET2 Prometheus use? We should be able to set up our own but want to develop against the right type of auth
- rp1 re-build/move (Lincoln)
    - This service will be moved to our new Netbird mesh over the next couple of weeks

2024.08.22

Present: Aidan, Lincoln, Louis, Eric, Judith, Brian, John, Ofer, Fengping, Horst, Armen, Eduardo

Apologies:

Agenda/Minutes:

- Some recent thinking about facility R&D in the context of analysis facilities, give it a read if you like:
    - Evolution toward Federated Analysis Facilities
    - 5 areas we have identified as needing work to get to a distributed analysis facility:
        - **Policy** → Working with sites to get users and work connected to resources and reduce friction
        - **Identity** → Providing users with a unified identity between sites, i.e. ATLAS IAM, to make it easier to float between
        - **Network** → High performance, low latency connectivity between sites, potentially adding encrypted overlay networks to simplify other parts of the stack

        - **Data** → Heavy use of caches, potentially working on things like a US-centric EOS service, quantifying performance for POSIX-like access over a WAN
        - **Compute** → Homogenizing the compute environment between sites in some way, distributing work between sites to maximize utilization
- **Netbird registration demo with IAM integrated (Aidan)**
    - Eduardo will discuss with Rafael about allowing us to create a pod in a namespace at NET2 to connect to Netbird
- **"s1" updates (Lincoln)**
    - Will be rebuilding on netbird
- **Accounting updates (Brian)**
    - Kuantifier 1.0 RC1 just released! Working with the UNL team to test
    - Includes nodeSelector configuration, breaking changes to the "config" block
    - Authorized Prometheus support not planned for 1.0
    - Armen will update the helm chart whenever Brian et al are ready
    - Matt working on documentation
- AOB
    - pam_oauth2_device plugin
        - Would allow for IAM-based SSH to nodes
        - Used on the UW OSPool and PATh Facility APs
        - (Lincoln) Can OSG ship it? Can it be in EPEL?
            - (Brian) Need to develop relationship with the developers etc but no blockers

2024.08.08

Present: Lincoln, Ofer, Brian, John, Fengping, Ilija, Aidan, Louis

Apologies: Rob, Armen, Horst

Agenda/Minutes:

- WireGuard (Lincoln)
    - https://docs.google.com/presentation/d/1_D02Pp4HagOYT5ZQuerENWcc9GD-k3ZIFO-0yypoUMA/edit#slide=id.p
- Netbird demo (Aidan)
- Kuantifier updates (Brian)
    - Ramping up to make 1.0 release happen by end of the month
- AOB

2024.07.25

Present: Lincoln, Judith, Fengping, Louis, Armen, Aidan, Ofer, John, Ilija, Philippe, Matt

Apologies: Rob, Horst, Eric, Brian

Agenda/Minutes:

- IU updates (Aidan)
    - Fixed network issue causing problems with exposed services on the cluster
    - Deployed MetalLB and Nginx Ingress controller to expose services as opposed to NodePort service
        - Added cert manager to handle X.509 SSL/TLS certificates
    - Modified cluster configuration to have 3 control planes as recommended by Eduardo
        - Can we document this? Somewhere.
- Kuantifier updates (Matt, Armen, Eduardo?)
    - Helm chart, being tested
    - Config that Armen sent is mostly good, just need to uncomment namespace and VO values
        - (Armen) Installation was successful in the sense that the cron job is in place and running - just missing some configuration
        - (Armen) Why is namespace/VO needed?
            - (Matt) Kuantifier only reports metrics for pods in a particular namespace. Default is null in the config but non-nullable in the Python script.
        - (Armen) What about SWT2 Google cluster?
            - (Matt) Will follow up
    - Working with Sam Albin at UNL to get an installation going there as well
    - Ongoing email thread with Eduardo at NET2
- Keycloak/Binder updates (Lincoln)
    - https://docs.google.com/presentation/d/1WOka2BgtbhbNdX45FILcQp9XNuj4YeebHxhhWOTjZIY/edit?usp=sharing
    -
- AOB
    -

2024.06.27

Present: Lincoln,Fengping,Aidan, Doug, Eduardo, Frederick, Horst, Judith, Louis, Ofer, Philippe, Armen

Apologies: Rob, John, Eric

Agenda / Minutes:

- REANA Updates (Fengping)
    - Need the Helm chart / code to be updated to support setting hostnames for ingress to access Jupyter.
        - Jupyter not currently functional
        - Developers are going to cut a release sometime in the summer but we need to use it in July
        - Who are the users that will use it?
- IU Updates (Aidan)
    - Created Kubernetes cluster with Kubespray
    - Deployed a cluster with Kubespray, 'hello world' type of application
    - 5 nodes with 2 control planes
        - Eduardo: 2 control planes. Do you know what algorithm you are using for election? May want to check it because it may cause faults if something fails
    - Older storage servers
    - Erasure coding?
        - We are not using it for AF
            - Would probably enable it if we did it again
        - Is it appropriate for ServiceX outputs?
            - Other options? XRootD S3? Seaweed FS?
- Accounting Updates
    - NET2: NTR
    - SWT2: Ongoing work with OSG to get this working. Hammering out the helm chart. Armen will follow up with Brian next week.
- K8S Comparison Paper
    - Completes the MS331 milestone
    - We are open to feedback
        - Ofer: What are the next steps? ATLAS internal note? Talk?
        - Doug: Who is the audience? What is the labor cost for sites deploying Kubernetes?
- K3S / DIRAC
    - John De Stefano sends that there was a DIRACX workshop in Lyon where k3s was used to deploy the development platform
        - https://github.com/DIRACGrid/diracx-charts/tree/master/k3s
- CVMFS - What is the best strategy these days?
    - What are folks using?
        - CVMFSExec (single-user)
        - osg-cvmfs-k8s
        - CERN CVMFS CSI
    - NET2 experience:
        - Separate CVMFS pod as a daemonset on the cluster

        - cvmfs-service container from CERN
        - Stuck pilots:
            - No defunct processes, no processes stuck in I/O wait e.g.
            - Restarting the CVMFS pod
        - CVMFS daemonset
            - Image comes from CERN but is *not* the CSI
        - Does not request cpu/memory/disk but doesn't seem to the problem
    - Doug: We should have a recommendations document, e.g. at BNL Jupyter will eventually be backed by OpenShift - how should they install CVMFS?

2024.06.13

Present: Lincoln, Rob, Shawn, Armen, Louis, Fengping, Judith, Ofer, John, Saroj, Tony, Eric, Philippe, Aidan, Horst
Apologies:

Agenda / Minutes:

- Focus this week is scrubbing.
- Milestone summary to date: Milestones Sheet and Facility R&D Milestones
    - (MS331) Decision on standardized K8S distribution for the US ATLAS Computing Facility
        - We have this as 'completed' in the milestone spreadsheet, but we really should finish this white paper or have a draft of it. Can we do it by scrubbing?
            - Rob: I will add a reference section today, and some research on multi-tenancy (for the OKD section)
            - We didn't find a one-size-fits-all solution for both on-prem and stretched clusters
                - Document to describe these differences and relative strengths/weaknesses
            - Need to make a statement to developers that they need to support the different distributions
        - https://www.overleaf.com/8766677924vgtxxtcqrtgw#51ec8d
    - (MS332) Implement bursting from shared Tier3s
        - Estimated time to completion (ETC): November '24
        - will go into FY25 milestones
        - Do we still want to do it? Yes!
        - How do we do it today? Moving machines from clusters into other clusters (e.g. MWT2 -> UC AF)
            - Not dynamic, not so easily reversible

        - Want to send jobs to T2 resources, but also potentially to Cloud
        - Potentially two demonstrators in FY25:
            - Scale out to a co-located T2 or T1
            - Scale to rented public cloud
            - Retire MS332 and create 2 new milestones?
            - This would greatly benefit from specificity
- (MS333) Build stretched k8s fabric across Tier2s with services (ServiceX, Dask)
    - Due Jul '24 - Complete
    - We have completed this for 2 US ATLAS sites (MWT2, AGLT2) but have not yet tried any services beyond BinderHub
        - ServiceXLite workers
    - What do we need to do to call this complete?
- (MS334) Server Side Demonstrator at scale (Ilija, Fengping)
    - Is this done? Can we mark it complete?
    - ServiceX over FAB
        - We can mark it complete
    - We may want to define a FY25 milestone using the (considerable) FAB cluster at CERN. We need to define an appropriate challenge metric (e.g. 200 Gb/s-like but *not* 200 Gb/s)
        - 700 cores, GPUs
        - Delivery out of EOS not fully exploited
        - New milestone:
            - Deliver hot datasets efficiently from EOS to the AFs via ServiceX
            - One group identified who would use this
- (MS335) Deploy monitoring, alerting and APEL accounting for UTA k8s cluster using Prometheus (Armen)
    - ETC: May '24 - **Delayed**
        - **New target:** TBD, follow up with Brian/Mat
    - Dependent on OSG-LHC development milestone for Kuantifier
        - Kuantifier: could be finished prior to scrubbing? will need to check w/ BrianL and Matt.
        - What happens next? Ready to "hand off" to the sites? (completed as an R&D, now to be deployed in "production" by Armen and Fernando for the SWT2-google siite)
    - To call this complete:
        - Kuantifier at UTA: MS335
            - Some issues with 'site name', renamed to 'cluster name' for OSG mode, some changes made in topology, etc.
            -
        - Kuantifier at NET2 / SWT2-Google: New milestones in 2.3.2.x ?
- (MS336) First resource tagging demonstrated
    - ETC: Dec '24
    - Can be moved into FY25

        - We have resources tagged in 's1' (labels such as 'site' and 'institution') but no demonstrator app
            - "sxl" label? (i.e. a transformer source)
            - On track - the labels exists, can be used for the decision making process but we need to build something that implements this.
- (MS337) Agreed I&A service in operation
    - ETC: May '24
        - Completed
    - We have Keycloak running and configured to allow ATLAS IAM auth
    - What more do we need to call this done?
    - Power users coming in to use K8S directly etc etc might be more of internal business
    - Today users can use ATLAS IAM and login to a resource and use things
- (MS338) Decision and report on scheduling approach for the stretched K8S platform
    - ETC: November '24
        - On track
    - Move into FY25
    - Related to our Kueue and Karmada work during the workshop
- (MS339) Demonstrated HPA for leading services
    - ETC: March '24 - Complete
    - Using in production on the Analysis Facility
- (MS340) Deployment of multi-site K8S across all Tier 2s
    - ETC: December '24
    - We need to figure out how to integrate resources at SWT2 and NET2 to call this complete.
        - use existing resources? e.g. Karmada
- (MS341) Multi-site continuous integration framework in operation
    - ETC: April '24 - **Delayed**
        - Let's call this done.
    - The stretched platform (s1) behind https://rp1.hl-lhc.io is using Flux hosted here
        - https://github.com/maniaclab/flux_admin/tree/main/infrastructure/s1
        - Private repo with other UChicago clusters
        - Has secrets (encrypted) etc so we should decide how to manage this going forward and decouple it from UC in particular
- (MS342) Scaled K8S with over-pledge CPU in operation including PanDA backfill capabilities
    - ETC: September '24
    - What is a good demonstrator?
    - Retired workers at UC went to AF
        - Add a PanDA queue?
    - On track

- New milestones?
    - Now is your opportunity to suggest new items!
    - The following FY24 milestones can be carried over into FY25:
        - MS332 - Bursting from T2s (November)
            - Break into two milestones, T3 → T2, T3 → Cloud
        - MS336 - Resource tagging (unless we want to call this complete) (December)
        - MS338 - Scheduling on the stretched K8S platform (November)
        - MS340 - Multi-site K8S across *all* T2s (December)
            - Potentially close this milestone
        - Namespace thing..
    - Questions from the WLCG meeting in Hamburg
        - Are there any AF-relevant facility questions we can turn into milestones?
- Scrubbing Highlights for Facility R&D
    - Built an experimental stretched Kubernetes R&D platform across 5 institutions (4 US, 1 Canada)
    - Held a 2.5 day Kubernetes hackathon, successfully building 20 clusters and exploring a number of interesting technologies
    - Worked with T2 sites and OSG-LHC to integrate tools for WLCG accounting for K8S clusters

2024.05.30 - CANCELED

Present:
Apologies:

Agenda / Minutes

2024.05.16

Present: Lincoln, Fengping, Brian, JD,
Apologies: Rob

Summary of WLCG IAM CERN Deployment Status and Plans –jd

- WLCG is migrating auth at CERN from VOMS in EL7 VMs to IAM in k8s containers, in coordination with EL7 EOL

- Initial IAM containers were deployed within CERN's OpenShift infrastructure
- CERN IAM team finds CERN's OpenShift implementation (not necessarily OpenShift itself) too restrictive due to security policies and other deployment decisions
- CERN IAM team plans to move containers to CERN's OpenStack infrastructure for more autonomous control over deployment, components (e.g., load balancing, DNS, performance tuning, logging, monitoring)
    - Again, this decision is orthogonal to whether OpenShift provides or permits these features natively; the CERN IAM team found that the CERN OpenShift team controls them too stringently in that deployment, and they prefer the flexibility provided by the CERN OpenStack deployment to manage these features themselves
- Side note on BNL plans: currently provide OKD to experiments, strongly considering moving to OpenShift as replacement for RHEV and other VM & container hosting technology
    - Also considering paid OpenShift training for key admin staff
- Additional details on CERN IAM services and deployment in talks by Berk and Maarten at WLCG/HSF Workshop

Agenda / Minutes

- VOMS client that OSG shifts does not work with IAM
    - e.g. 'voms-proxy-init' doesn't work with ATLAS IAM
    - Brian Lin would appreciate if someone could test the patched version
- WLCG IAM update ^
- K8s Tutorial Nodes
    - We will reclaim any leftovers at this point
- Harvester-K8s accounting update (Brian, Eduardo)
    - Helm chart PR in progress https://github.com/rptaylor/kapel/pull/53
    - Basic docs in progress
    - Build pipelines will probably be punted to post-May
    - UC team will look into resurrecting IRIS-HEP-UCHICAGO-K8S queue in CRIC
    - OSG-LHC team still waiting on Prometheus authz WIP branch from Eduardo
    - SWT2 k8s Prometheus should not have authz configured so it would be a good first candidate for installing the accounting helm chart
- Reana deployment at UC AF (Fengping, Eric)
    - Reana installed deployed on UC AF (thanks for collaboration from Tibor @ CERN): https://reana.af.uchicago.edu/
        - Seems to be working.
        - Installed via Flux
        - Need to create client in Keycloak for IAM authentication
    - Plan to work on a reana-dev instance for developments.
    - Next steps -
        - After IAM is added, get some canned demos running? - EL: Definitively.
- K8S Distribution Paper (MS331)
    - NTR from Lincoln

    - Anyone else have time to take a look?
    - https://www.overleaf.com/8766677924vgtxxtcqrtgw#51ec8d
- s1 / rp1 updates
    - IU nodes lost power, rebooted. Had to enable wireguard service at boot to get them to rejoin the cluster.
    - BinderHub service: https://rp1.hl-lhc.io/ (EL: not reachable from BNL)
    - Focus now is on getting user identities propagated through Binder into Jupyter (related to MS337)
    - Anything we can do to get NET2 / SWT2 on board?
        - Federation via Karmada?

2024.05.02

Present: Louis, John, Philippe, Fengping, Armen, Ofer, Rob, Lincoln, Brian, Eduardo, Shawn, Fred, Saroj, Judith
Apologies: Eric

Agenda & In-meeting notes:

- What's next following the K8s hackathon (Lincoln)
    - Kubernetes Hackathon Closeout
        - Retiring the 20 test clusters
            - Zach and Phillipe expressed interest in hanging onto them a little longer. Emailed to see if they're done with them
            - Farnaz has started reclaiming the rest
        - Login05 - probably will get rebuilt.
        - We have a number of items to follow up on (see closeout slides), among them:
            - Followup on Wireguard feasibility at BNL.
                - Will follow-up with Tommy
            - Try to mix ATLAS and OSG jobs via Kueue
                - Any interest from UTA or others?
                - Would be nice to backfill the resources during these periods of no available sim
            - Deploy ServiceX Lite on other K8S clusters? e.g. UTA, NET2.
                - no "heavy lifting" for a site. There are some instructions to do this.
            - Cloudflare geo-/latency- based DNS? Interesting things to try here.

        - e.g. "[xcache|varnish].usatlas.org" that resolves to the closest cache? Eliminate a lot of logic by pushing things to Cloudflare (for a nominal fee 💰)
    - NET2
        - interested in trying out Kueue (on the work stack - next few weeks)
        -
- FY24 Facility R&D Milestones
- Completing MS 331
    - Note - MS 331 in its original form last summer:

Evaluation of Kubernetes Distribution (OKD):

- Assess the software modifications required to run in an OKD environment. OKD is the preferred K8S distribution of the national labs given its enhanced security features. NET2 has also chosen it as the basis for their new Tier 2 Center substrate.
- Explore compatibility with Squid (collaborating with OSG), XCache, Varnish and ServiceX; ensure suitability for a stretched K8S.
- Plan the migration of Tier 3 resources to the selected K8S distribution if OKD is chosen.
- **MS 331**: Decision on standardized K8S distribution for the US ATLAS Computing Facility, Feb 2024 (Gardner 0.1; Bryant 0.1; Vartepetian 0.1)

- Since that time we decided to generalize it a bit since there are interesting pros/cons or suitability to understand between different distributions.
- https://www.overleaf.com/project/65e64767d65cefcf93c1595f
    - https://www.overleaf.com/8766677924vgtxxtcqrtgw#51ec8d editable link
- Will be collecting contributions to finish this doc - please do let us know if you'd like to contribute! (esp. those doing OpenShift training at BNL especially welcome)
    - Berk Balci talk mentioning migration from CERN OpenShift to OpenStack (link)
- Note the business about "ensure suitability for a stretched K8S".. well, after the hackathon we could provide a few words about adding an encrypted VPN mesh for the control plan -- something being investigated in MS 333.

Stretched K8S Multi Site Cluster:

- Investigate using IP-IP tunneling or VXLAN to extend the cluster over the WAN.
- Consult with colleagues from the (NSF funded) Pacific Research Platform and the NSF PATh OSDF which together have several years experience operating in this mode to gain insights into their respective approaches.
- **MS 333**: Deployment of multi-site K8S with at least 2 sites: July 2024 (Bryant 0.15; Vartepetian 0.15)
- **MS 340**: Deployment of multi-site K8S across all Tier 2s: December 2024 (Bryant 0.15; Vartepetian 0.15; New hire 0.15)

    - Work relating to wireguard is appropriate.
    - MS 333: delivered with s1.hl-lh.io.
    - MS 340:
        - MWT2 - done
        - AGLT2 - done
        - NET2 - lets discuss
            - has some configured for this.
            - Lincoln: could have OKD make a VM that could be attached to s1.
            - Also interested in this as a federated infrastructure
        - SWT2 - lets discuss
            - there a couple of options - probably option 3 - federated.
        -
    - nb. "New hire" is Aidan Rosberg (Indiana) - starts Monday.
- Harvester-K8s accounting update (Brian, Eduardo)
    - From Eduardo: "currently I am trying to have a configuration where an authentication to communicate with OKD's Prometheus is not required, so the Kapel helm chart can be the same for OKD and non OKD clusters. It will require that at least the kapel uses the same namespace than the one used for Atlas production. "
        - Brian: I think this will require some minor changes to the Prometheus queries
        - authenticated KAPEL?
    - At the CHTC, we're still battling through a deficiency in kube-state-metrics that can cause us to miss pod completion times: https://github.com/kubernetes/kube-state-metrics/issues/2118
        - This shouldn't be a problem for the way that ATLAS runs pods (Harvester submits k8s Jobs objects)
        - We still want to pursue this to provide a generic K8s gratia probe that is agnostic to deployment style
        - The solutions we're coming up with require larger changes to the existing implementation so we'll need to coordinate with Ryan
- Reana deployment at UC AF (Fengping, Eric)
    - Eric successfully deployed a reana instance on the one node kubernetes cluster he built during the workshop
    - We worked out a reana deployment manager role that allow an external user to manage the reana deployment cycle in a cluster.
    - Eric is trying deployment on river-dev. Had some connection issues with “reana-client ping”. Looks might be related to tls so we created a clusterissuer and configured tls with letsencrypt cert issuer. Waiting to see if that helps.
- 200g update (Ilija)
    - XCaches are now better distributed over racks/switches so direct xrdcp bandwidth tests show 360Gbps served to 460 concurrent clients. Other activities on AF cluster can decrease this rate as the limiting factor are upstream lags from top of rack switches.

    - Similar test but where distribution of clients is done by Dask reaches bandwidth between 200 and 260 Gbps.
    - ServiceX tests that just reads, uncompress, compresses and delivers data to S3 (not a normal ServiceX mode of operation) reaches ~80Gbps while deploying 1200 transformer pods on AF. Transformers are running at close to 100% efficiency (mostly compression/decompression). This test maximally stresses S3 as it is writing a 7.4TB resulting basket with 64k objects at 30-40Gbps. To run a full scale test (191TB, 200k files input), we had to increase ServiceX S3 user quota to 50TB. Number of ServiceX issues found and are being addressed.
      Next test is running a more realistic ServiceX test where cuts are applied by transformers so output to S3 is much smaller. This will reduce compression CPU cost and stress on S3.
    - Dask - number of changes in its configuration. Stable running with in excess of 2000 workers (1 process and 1 thread per worker).

2024.04.18

Present: Alexei, Lincoln, Rob, Brian, Louis, Fengping, John De Stefano, Ilija
Apologies: Eric, Horst, Ofer, Shawn,

Agenda & In-meeting notes:

- Harvester-K8s accounting (Brian, Eduardo)
    - Making some progress the Gratia container that will take Prometheus data from KAPEL and ship it into Gratia
    - Testing for the format of the gratia record
    - Will send PR to Ryan for inclusion in KAPEL
- OKD/OpenShift considerations/diagram, experience at NET2 (Eduardo)
    - Delayed for now
- s1 (Lincoln)
    - There was an issue with the UMich node fell out of the cluster
    - Loadbalancer up and running, seems to be working
        - Was able to create a service at MSU and proxy traffic to it through UC
    - Next week to add Indiana University
    - Going to try ServiceX Lite on s1, BinderHub
    - Wireguard reliability discussions
- 200g (Ilija)
    - Tested that UC dCache can deliver >200Gbps to xrdcp clients running on AF. Less than 400 clients was enough.
    - We have 8 xcaches - 4 on 50Gbps connections and 4 on 25Gbps. Due to suboptimal distribution over switches, it currently can’t give 200Gbps in

aggregate. Next week all the xcaches will be on 2x25Gbps and be better distributed.
    - Current tests using uproot and servicex are still far from 200Gbps. A lot of software bottlenecks to fix.
- dask-gateway (Fengping)
    - Dask Operator -> Power users using K8S to deploy
    - Dask Gateway -> Managed Dask, as a service, uses JupyterHub as auth service
    - Installed on AF and configured to work with binderhub on AF
    - Should be able to install one on a different cluster and make it available
- k8s hackathon next week (Lincoln)
    - 23 registered participants, 18 in-person.
    - getting rooms organized & catering
    - dinner Thursday evening: signup
    - Provisioned k8s hosts for the tutorial session
    - Provided SSL-River access to Eric and Matthew for reana deployment testing
- AOB
    - for fy25 need to add ARM eval to the program of work
    - Wisc team just got an ARM server, will work on OSG software stack

2024.04.04

Present: Lincoln, Rob, Shawn, Alexei, Eduardo, Fengping, Ilija, Eric, Vincent, Horst, Louis Pelosi, Armen, John DeStefano, Doug, Fred, Justin Spradley, Carlos Gamboa, Brian Lin, Tom Smith, Qiulan Huang, Robert Hancock, Ofer, Verena.
Apologies:

Agenda & In-meeting notes:

- R&D plan at BNL: REANA, RECAST on K8s (Eric Lancon)
    - See presentation on https://indico.cern.ch/event/1386696/
    - A reana testbed in the US ATLAS computing facility - a starting point for discussion
    - An integration with existing facilities - grid storage, e.g.
    - Exploratory - not all questions will be answered.
    - Questions:
        - Doug's comment: according the the Deputy spokesman - Stephane Willocq - "The search groups did establish at one point a requirement that each analysis needs to be implemented in Recast. However, this was rebuked in the past year by the Physics Coordinators. Essentially, Recast cannot be required but of course if analysis teams want to extend their analysis in this way it is welcome. The search groups are resisting it seems but it is NOT an ATLAS requirement. Recast is nice to have but not the highest priority for sure. There has been some discussion about using Reana to do some things like the CP

group calibration but I do not know that this is being followed up actively. CP groups have limited development capabilities."

        - Alexei: on the ATLAS papers slide, the papers are old. Had to do with Lukas' involvement. Timeline comments: setups, then benchmarking - is this new? Has some initial ideas. Why not at CERN now? Ans: compare the new platform with whats available at CERN. Concerns about Tibor Simko from CERN IT availability. Ans - some discussions on-going on the htcondor manager class; there is similar work elsewhere
        - Eduardo - on slide 14, the overflow to a cluster -- k8s, condor, slurm.
        - Alexei - how is this different from the cern open data portal? deploying portal to a us facility? New development? .. needs to be clarified. What is new from Lukas' last demonstration of this?
        - Doug - no sense of how much effort would be needed for the deliverables discussed on Slide 16.
        - Verena - in the past, there was a requirement for an analysis to run through reana/recast, and this required a lot of overhead. If there were ways to help users there might be more willingness.. interest from CP groups for calibration, flavor tagging group. Initial results were promising. There are several analysis attempting to use recast. A bit of a shame because of those challenges. Folks use ntuples. FTG reana has been effective.
            - Eric - need to reach out to the US community, via US IB.
        - Shawn - reproducibility is critical for science; its hard, there's a cost. Can we automate?A
        - =]
- Update on Harvester-K8s accounting (Brian, Eduardo)
    - getting access to NET2
    - KAPEL exploring
- ~~OKD/OpenShift considerations/diagram, experience at NET2 (Eduardo)~~
    - Eduardo has been busy with a downtime and has not yet had an opportunity to assemble the materials. We will defer to the next meeting.
- Stretched K8s cluster update (Lincoln)
    - Adding nodes at IU
    - Starting to add other things ontop of the vanilla K8S install - MetalLB, Flux, Ingress Controller, Cert Manager, etc.
- The IRIS-HEP 200g demonstrator - cobbling requirements (Rob)
    - Alex Held and Gordon are organizing this, in conjunction with IRIS-HEP's analysis grand challenge series.
    - They need
        - Approximately 200 TB of local storage is required.
        - Up to 4,000 CPU cores, depending on fraction of file read
        - Currently working on scripts and notebooks and a smallish dataset
        - They need to identify and transfer input datasets to UC.
    - Organizational meeting tomorrow: https://indico.cern.ch/event/1394124/
    - At UC we are organizing a "wide & fast" xcache

- Tutorial updates (Lincoln) for the k8s hackathon, https://indico.cern.ch/event/1384683/
    - We're working on making the tutorial now. I will send out some pre-instructions on the week of Apr 15
        - TL;DR - Get an account on the UChicago AF, we'll use that as a starting point for the tutorial infrastructure.
    - We will be focusing on two topics on the tutorial morning:
        - Kubespray cluster installation
        - GitOps with Flux
- AOB

2024.03.21

Present: Brian, Lincoln, Judith, Ilija, Armen
Apologies: Rob, Horst, Fred, Ofer

Agenda & In-meeting notes:

- Update on accounting (Brian, Matt, others)
    - Matt looking closer at KAPEL
    - Will follow up with Eduardo about credentials
- OKD/OpenShift considerations/diagram, experience at NET2 (Eduardo)
    - *Defer to next meeting*
- Topics for training session, and pre-requisites. (Brian, Lincoln)
    - Depends largely on the audience - not many people signed up so far. Some possible topics:
    - Concepts / introductory material (beginner)
        - Pods, Deployments, Ingress, LoadBalancers, etc
    - Continuous integration tools -GitOps / Flux (intermediate)
        - Managing objects deployed in Kubernetes with Git-based workflows
        - Useful additional tools like Sealed Secrets etc
    - Monitoring - Prometheus (intermediate)
        - Using the Prometheus Query Language (PromQL)
        - Making plots, etc
        - Adding metrics to existing applications
    - Logistics:
        - **DEADLINE for signup: No later than Apr 11 (2 weeks before)**
        - **We will send out another announcement to WBS2.3 mailing list**
- Stretched cluster over WireGuard (Lincoln)
    - Cluster assembled between UC, UMich, MSU and UVic
    - WireGuard used for control plane networking - all nodes join a WireGuard mesh with a private network
    - Still TBD how public-facing services will work in this configuration
- NRP policies for services

    - Open source - hosted on NRP Gitlab, Fengping has the links
    - Policies:
        - Everything is limited in time
        - Must request memory/cpu when creating an object
        - Automatic kill/cleanup for objects that aren't using their requests

2024.03.07

Present: Lincoln, Fengping, Rob, Brian Lin, Ilija, Armen, Jyothish Thomas, Ryan, Matt Westphall, Philippe, Ofer, Eduardo, Ivan, Fred
Apologies:

Agenda & In-meeting notes:

- Update from the accounting working group (Brian Lin)
    - Lots of traffic in email, and meeting last week with folks working on the problem. Welcome Matt Westphal from OSG Software team who will be developing on the OSG side.
    - May 31 is the target date
    - Will need to verify access to NET2, for dev purposes. Eg. will need pod metadata.
    - Related - k8s schedulers exploration? Operation on the "job object" -- is that always appropriate?
        - just want to make sure in the accounting design we can work with various schedulers.
        - assumption - how namespaces are used
        - Lincoln found that most tools consider things as managing job options.
        - Ryan - volcano - wouldn't be compatible with Harvestor, e.g. But kueue would seem more appropriate.
        - Eduardo - want to be compatible with API version; want to make sure we're in compliance. Another consideration.
        - n.b. there are already established schedulers out there - we should not want to require them to change schema significantly, ie. should be compatible with htcondor, harvester, etc.

- Outline whitepaper to meet **MS 331**: (Facility R&D FY24 Milestones) (Lincoln)

  Evaluation of Kubernetes Distribution (OKD):
    - Assess the software modifications required to run in an OKD environment. OKD is the preferred K8S distribution of the national labs given its enhanced security features. NET2 has also chosen it as the basis for their new Tier 2 Center substrate.

- Explore compatibility with Squid (collaborating with OSG), XCache, Varnish and ServiceX; ensure suitability for a stretched K8S.
- Plan the migration of Tier 3 resources to the selected K8S distribution if OKD is chosen.
- **MS 331**: Decision on standardized K8S distribution for the US ATLAS Computing Facility, Feb 2024 (Gardner 0.1; Bryant 0.1; Vartepetian 0.1)

    - Started a document https://www.overleaf.com/8766677924vgtxxtcqrtgw#51ec8d
        - Plan to compare kubeadm/kubespray, OKD, and Rancher solutions
            - Welcome additional perspectives and contributions esp from those of you who have used OKD or Rancher seriously
        - Goal is to allow future sites to make informed decisions about setting up Kubernetes
            - Not prescriptive - Most likely a 'one size fits all' solution is not possible
        - n.b. Ofer: Chris H has a document, not public.

- Getting started on the stretched platform - update (Lincoln)
    - Nodes:
        - Ryan Taylor has provisioned a node for us on the UVic Cloud with EL9 (ATLAS Canada)
        - Wenjing and Philippe setup login for DRAC to 2 nodes at AGLT2, one rebuilt with EL9 last night
        - 5 nodes identified for UChicago but not yet provisioned (this week)
    - Eduardo - SDN over L2 - interested in pursing
    - Investigated WireGuard and tools in the WireGuard ecosystem (e.g. for the **Provisioning layer,** not the data transfer layer, perhaps also the **control plane**)
        - Tunnel K8S cluster traffic through WireGuard VPN?
        - Possible advantages:
            - All stretched traffic goes over a virtual layer2
            - All low-level traffic secured with modern cryptography standards
        - Disadvantages:
            - Possibly complicate networking further
            - Not 100% clear if a mesh topology is required/desired, which would be difficult to scale without external tools like Tailscale (proprietary)
        - Need to understand if the juice is worth the squeeze on this one
        - Philippe: https://nebula.defined.net/docs/ interesting
        - Eduardo - we have a HA setup, experience with traffic between control planes (3, 5). Significant traffic during recovery from downtimes, e.g. during state reconciliation.
            - also, https://la11111.wordpress.com/2012/09/24/layer-2-vpns-using-ssh/

- Concrete plans for hackathon participation (Fengping)
    - Preliminary thoughts (Rob):
        - Goal will be to build a versatile "research platform" for facility R&D without an expectation of official production support. i.e. we want "applications" (analysis workflows) to benchmark and develop, but this shouldn't be conflated with our production analysis resources.
        - The platform, rp1.hl-lhc.io, will aggregate a combination of the stretched cluster (s1.hl-lhc.io), and other k8s clusters, including the IRIS-HEP SSL resources: river, river-dev, flatiron. And NRP, for its vast collection of GPU resources. Requirements would basically be API access / TLS proxy. Other clusters? :
        - Note resource will be suitable for the US ATLAS Summer training workshop, and the IRIS-HEP/ATLAS training event.
        - It will use ATLAS IAM for identity and auth
        - It will be open to all ATLAS members in IAM without approval (siimlar to grid jobs). ATLAS users not yet in Glance (undergrads) will addtionally be able to register.
        - It will use KeyCloak
        - It will provide a bindherhub interface
        - We will provide a storage component at UC to store notebooks and input/output data.
        - Will use IRIS-HEP AGC (Alex, Matthew), Z-->ee demo (Nikolai, Ilija, Yoav), and top quark group RNTuple (Tomas) demonstrators.
        - Will need to identify metrics; meaningful benchmarks; provide a dashboard.
    - Kubernetes installation manuals (or training modules)
    - Training curriculum on first morning (discuss)
    - Eduardo, Philippe, Brian, Ryan planning to attend in person.
- Kubernetes scheduling- kueue (or karmada?) (Fengping)
    - Karmada: open, multi-cloud, multi-cluster kubernetes orchestration
        - Use native api definitions for easy to integration with existing tools such as flux for ci/cd
        - Multicluster service discovery
        - Multicluster scheduling features
        - https://karmada.io/
        - Lincoln: these seem to be successors to kubfed, kubfed2,... Maybe would be good to survey these technologies: KubeAdmiral, Admiralty, Kueue - also doing multi-cluster

- Topics for next meeting (3/21)?
    - Update on accounting (Brian, Matt, others)
    - OKD/OpenShift considerations/diagram, experience at NET2 (Eduardo)
    - Topics for training session, and pre-requisites.

2024.02.22

Present: Rob, Lincoln, Horst, Armen, Ivan, Kevin, Ilija, Fengping, Jyothish, Ofer, Ryan, Brian, Eduardo, Fred
Apologies: none

Agenda & In-meeting notes:

- In-person Hackathon, https://indico.cern.ch/event/1384683/overview
    - See slides from Lincoln, Kubernetes Hackathon Plans
    - Rob: use the stretched platform for future training events
    - Ivan: is remote possible. Training can be remote.
- Review of milestones and planned actions
    - Jan-Mar 2024 Milestones
    - Accounting
        - Blin - cred to access to ne kubernetes. Uta need more clarification on what access is needed.
        - Ryan - cron job per vo(cpu_request times wall time?)
        - Rob - coordinate to try to avoid having to reinvent the wheel.
- Next meeting March 7
    - Getting started on the stretched platform
    - Updated on accounting
    - Concrete about hackathon participation
    - Kubernetes installation manuals
    - Kubernetes scheduling- kueue

2024.01.25

Present: Lincoln, Rob, Fengping, Brian L, Ivan, Judith, Armen, Horst

Apologies:

Agenda & In-meeting notes:

- Milestone - any updates for tomorrow's deadline?

    - Milestones
- Stretched k8s fabric - Lincoln
    - US ATLAS Stretched Kubernetes Fabric
    - Brian L - what about security? We do the PATh facility.
    - Eduardo - preferred method to respect campus policies - create a project to deploy into.
        - Admiralty may play a role
    - Ofer - worries about federated operations responsibilities
    - Fred - worries about manpower requirements
    - Brian - outsiders view: who is responsible for a security issue; also wrestling with this in PATh.
    - Eduardo - comment on security. Ideas of responsibilities - virtual cluster admin should have responsibility. Reduce complexity of having expertise on site.
    - App - Binderhub on stretched k8s
- Next meeting:

2024.01.11

Present: Lincoln, Rob, Ivan, Armen, John, Fengping, Brian, Ofer

Apologies: Horst

- Milestone review
    - Milestones
        - UTA Monitoring/Accounting
            - APEL records have been/will be uploaded manually until a mechanism has been developed
            - We can checkoff the milestone given that monitoring is in place at UTA and APEL records can be uploaded manually for now
        - Standardizing K8S
            - Need to develop set of questions that we want to answer in order to finish the milestone
        - Stretched K8S
            - Come up with a questionnaire that can be sent to sites
            - Ask each T2 to contribute at least 1 server to the stretched K8S and ask how much we can directly control the node
                - BMC level (DRAC)
                - Host level (root)
        - Some of the milestones may need reordered - some things depend on stretched k8s before we have the stretched k8s
        -
- Ideas for topical meetings:
    - Prometheus/PromQL tutorial

        - Needed by many sites especially running K8S (but not limited to)
        - May be worth finding someone who can give a presentation on PromQL
    - FluxCD, ArgoCD presentation(s)
        - Flux:
            - Possibly Brian Lin or someone from his team?
            - Refer to Brian Bockelman talk at CHEP?
        - Argo
            - Well integrated into OpenShift
            - Ask Ricardo or someone from his team?
    - Identity management in OKD via
- Topics for S&C talk
    - Intro to new area in the US ATLAS ops program
    - Overview of activity goals
    - Synergy with IRIS-HEP SSL & Analysis Facilities (current planning)
    - Identify management explorations
    - NET2 approach & experience
    - UTA k8s explorations
    - The BNL experience
    - Resources
        - Copy of WBS 2.3.5.2 Facility R&D
        - Copy of Facility R&D (Lincoln's slides 2023.08.02)

2023.12.14

Present: Lincoln, Ofer, Fred, Eduardo, Brian, John, Judith,
Apologies: Horst, Rob

Identity provision use case considering user groups management.

In OKD, the way one can associate a group of users to a group of resources is via an association of a Certification Authority certificate and a project (namespace) confining the resources intended to be used by this group of users. After the CA cert is imported into the cluster, all users with certificates signed by that CA will generally be accepted for authentication. While Authentication it's a cluster-wide configuration, Authorization is controlled via RoleBindings, assigning roles to groups within a project. Groups are mapped from the Organization (O) fields from the user's certificate, and the username is mapped from the Common Name (CN) field. Concretely : (to be completed)

2023.11.30

Present: Lincoln, Ofer, Matt, Mat, Brian, Fengping, Judith, John De Stefano, Armen, Eduardo, Rafael, Fred
Apologies:

Agenda:

- Accounting discussion
    - Meeting between OSG and NET2 to talk about WLCG accounting
    - NET2 doing the manual report currently
        - Where does this live?
        - CRIC
            - https://wlcg-cric.cern.ch/wlcg/accdata/list/
        - How is the data extracted?
            - ATLAS Grafana
            - Prometheus
            - Compare both
        - KAPEL ?
            - Depends on Prometheus data
    - OKD Prometheus
        - Metrics not exactly the same as they are for vanilla K8S
        - Eduardo working on changes to KAPEL to support OKD
    - SWT2, NET2, UC AF all using Prometheus
    - How is ATLAS<->K8S integration done?
        - Same between NET2, SWT2
        - Harvester submitting directly to K8S via API access
        - Permissions in OKD are different - needs additional attention from the admin to make jobs run correctly.
            - e.g., mounting CVMFS needs elevated permissions on the pod
            - OKD needs additional work here
    - Gratia Probe equivalent to KAPEL
        - Just use KAPEL directly?
            - Adding GRACC support to KAPEL might be as big as KAPEL itself
            - Should we set up another meeting with Ryan Taylor to talk about it?
        - What is needed for the Gratia probe?
            - Prometheus seems to be ubiquitous
            - Do we need a Prometheus/PromQL hackathon?
                - Find some experts, train us all on how to use it 🙂
        - What deployment models are we forcing on sites?
            - Don't want to force any particular deployment strategy on the sites (Pods, ReplicaSets, DameonSets, Jobs …)
            - We are however suggesting that Prometheus

        - What needs to be done?
            - Queries need to be adapted for the Gratia codebase
            - Eduardo offers view access to NET2 to help develop against an OKD resource
        - Common layer - Pods
            - Jobs create Pods, ReplicaSets use Pods, etc
    - SWT2-K8S needs to do this as well - manual reporting to start.
        - Armen will email Julia for manual reporting for SWT2 K8S
    - Is there a way to correct the numbers on the GRACC level?
        - Contact support@osg-htc.org
    - Action items:
        - OSG interested in K8S scheduler work, what it means for multi-tenancy
        - OSG needs to take a deeper look into KAPEL, Prometheus and PromQL and start sketching the design
        - Eduardo will give OSG SW access to OKD at NET2
            - View account, vs being able to submit jobs etc. Eduardo can provide whatever is needed
        - Prometheus hackathon
            - Attach this to a K8S Meetup?
- AOB

Minutes:

2023.11.16

Present: Lincoln, Fengping, Ilija, Eduardo, Mat, Matt, Judith, Ofer, Armen, Brian, Horst, Will, Fred
Apologies: Rob, John

Agenda

- ServiceX Lite presentation
- HTCondor in K8S / Autoscaling presentation
- Recap on accounting efforts, next steps

Notes

- ServiceX Lite
    - What is the role of PostgreSQL?
        - All requests are stored in pgsql
    - Helm chart exists, but its still very complicated and always changing

    - Discussion about how to deploy ServiceX Lite on a cluster
        - Eduardo mentions giving access and letting Ilija operate the service, handle the upgrades, etc as long as there is nothing special required from the admin
        - Brian notes this is essentially the SLATE model
    - Question about cleanup:
        - What is cleaned up when a request is done?
            - Cleans everything concerning the request. Doesn't delete the entire ServiceX infrastructure on the 'slave' cluster (partner cluster? :) .
    - How does the ServiceX transformer know where to run transform?
        - Does the master intelligently dispatch the transformation request?
        - Ilija: Currently, if a request for 3K files is sent to Master ServiceX, then it will ramp up the number of transformers 'locally', then ServiceX lite will poll and spin up its own workers. No optimizations yet because the ServiceX LIte instance on River is very close to the master ServiceX
    - is it helpful to have an S3 cache at the sites as well?
        - The users should be near the S3 object stores. So if BNL has a lot of ServiceX users, they should talk to a nearby S3.
    - Would BNL deploy this?
        - Difficult to give access, but if they were deploying themselves it would be more tractable, especially if the software is simpler than a full-blown ServiceX deployment
    - What secrets are copied to a partner cluster?
        - RabbitMQ? Read/write? What the implications giving these secrets to untrustworthy people?
            - (To be followed up - Ilija had connection issues)
- HTCondor on UChicago AF
    - How does Kubernetes know if a HTCondor pod is idle?
        - HTCondor autoscale manager
            - Occupancy metrics fed into the HPA
    - Why are there long/short queues, why doesn't HTCondor just fair-share it?
        - Short queue has a limit of 4H, encourages job churn
        - Don't want HTCondor to be dominated by long-running jobs

2023.11.02

Present: Lincoln, Rob, Rafael, Will, Eduardo, Fred, Horst, Armen, Judith, Fengping, John, Ofer
Apologies: Brian Lin

Agenda

- - Review of milestones and plan
- - OKD setup at UChicago
- - Accounting
- - Open floor

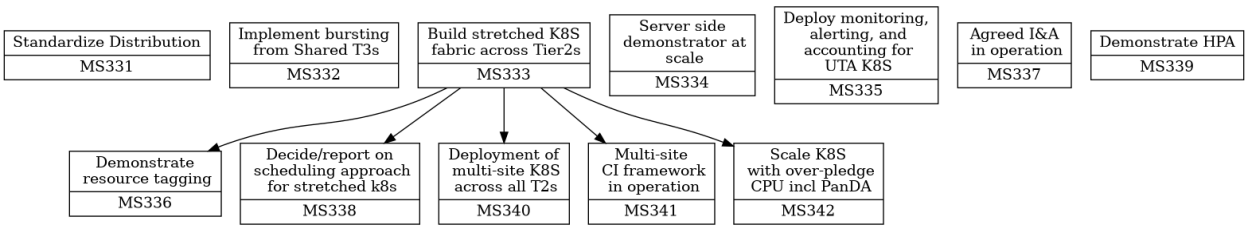

Notes

- Intro/overview by Rob
- Lincoln reviewing milestones and relationships
- Most are kubernetes-based
- E.g. one of the first issues is the k8s distribution - and OKD vs vanilla
- Ecuardo - OKD offers the ability to manage infrastructure as well.
- Need to make sure apps work not only on vanilla
- Opinionated - cannot choose the network controller; more restrictive security rules; does have functional considerations
- Lincoln - setting up OKD at UC since many orgs are choosing it, and we need to be able to support development and applications
- E.g. NERSC uses Rancher
- Ofer - vanilla as a subset of OKD? app compatibility - works on OKD should just work on vanilla? Eduardo - permissions differ, they will not work out of the box. Permissions are somewhat malleable. Not quite 100% containerization ready.
- Lincoln: can the security folks enumerate the security risks associated with vanilla? Would be good to know this, as motivation.
- Eduardo - need to understand better what we can support . And recommend to new sites. They will have different requirements.
- Implementing bursting from shared tier 3
    - Condor glideins to run on the tier 2s
    - Cloud resources
    - HPC site GPUs
    - Dask?
- Server side demonstrator: ServiceX
    - IRIS-HEP, FABRIC/FAB
    - Cluster at CERN, on the FABRIC network
    - Could re-use same technology for the T1s and T2s?

- UTA K8S Monitoring
    - Also for NET2
    - OSG is helping with this for getting information into GRACC
        - Should be separate milestone because of timelines
        - Split monitoring (todo, will take time for OSG) from alerting and accounting (mostly done)
        - Problems with using KAPEL due to authentication
        - Armen is collecting the metrics now, but needs the mechanism to submit to OSG
        - To discuss more with Brian Lin
        - **Set up separate dedicated meeting with Brian Lin**
- Stretched K8s across T2s
    - Single interface/control plane to cluster, workers are spread across multiple sites
    - Already provided by harvester? No, not necessarily just panda jobs. Similar, but more K8S focused (labels, documented for users/site admins, etc.)
    - Scheduling is simple, no fair-share scheduling by default. Develop multi-tenancy support more similar to our existing batch systems
    - PanDA backfill to keep cluster full
    - Supplement analysis facilities (with no hardware funding) with additional resources
    - Using old hardware at NET2 might be an issue due to to datacenter requirements. Might need to figure out an alternative location for them
- Organize future meetings to deep dive into individual topics