

2017-10-09 Sprint Planning

Mike, Ben, Kevin, Charles, Craig

Status:

- Kevin 25%
 - Done with CCSN-MRI-SIMS
 - Now on other projects (non NDS)
 - Interested in Kubernetes deployment in OpenStack
- Charles 25%
 - Normal schedule ahead
- Ben 50%
- Mike 25%

Current priorities:

- NCSA Industry conference demo on Wednesday
 - Workbench > Spark integration
 - Zeppelin > Spark via Livy
 - Spark > MongoDB using the NBI dataset
- SC17
 - Workbench > HPC integration
 - TERRA-REF image stitching on ROGER (TORQUE/PBS)
 - Jupyter > Agave API > ROGER
 - Jupyter > Agave API > Comet (Singularity)?
 - Comet access
 - XSEDE may be easiest
- Options:
 - Deployment in commercial cloud (multi-node cluster with storage)
 - No one is happy with the Kubernetes deployment process on OpenStack
 - kubeadm is promising
 - Shared storage
 - <https://www.minio.io/>
 - Review of OpenShift Origin
 - Kubernetes, Swarm, Mesos have more traction
 - Singularity and Shifter
 - Security
 - Authentication and authorization
 - OAuth support in Workbench
 - LDAP or OAuth scopes for authorization
 - Q. What does it mean to share authentication/authorization with
 - HPC cluster such as ROGER (NCSA LDAP)
 - Spark cluster? (Kerberos or nothing, what about Livy)
 - Container and filesystem users/permissions
 - Mapping UID/GIDs into running containers
 - Q. Root escalation in Docker
 - Data sharing and permissions
 - Controlling access to data on the filesystem but also in active database such as NBI mongo.
 - Same problem with Spark (pulling data from Mongo into Spark)
 - More sophisticated network configuration
 - Both for Kubernetes
 - Also with the cloud provider (e.g., OpenStack project)
 - Monitoring/LMA
 - Addons <https://kubernetes.io/docs/concepts/cluster-administration/addons/>
 - Prometheus?
 - Nagios/NRPE is a stopgap
 - Production maintenance
 - Upgrading the beta instance
 - Redeploying the beta?
 - /var/lib/docker volume scaling
 - Fixing inactiveTimeout for inactive accounts
 - ETK/earthcube instances
 - Cloud9 wily (and other Wily) upgrades; use Cloud9 all
 - Process for removing old specs if in use.
 - Clowder/Workbench plugin?
 - Evaluate FRDR
 - <https://portagenetwork.ca/frdr-dfdr/>
 - Other
 - Drop-in UI nonsense
 - Bower bug 410
 - Deploy tools using configmap?
 - Deploy tools using SMTP and standalone etcd
 - /var/lib/docker and kubelet mount issues (may need depends on)
 - Why are we using XFS for /media/storage?

Notes:

- Discussion of security in Spark
 - Kevin: focused on network access control
- Getting off of the demo treadmill
- Need to really understand OpenShift
 - Security model,
 - Application/deployment model
- Easy Kubernetes deploy
 - OpenStack (Nebula/SDSC)
 - HA?
 - Secure? Networking, TLS?
 - Then in AWS, Azure (Big Data Hubs), GCE
- Security is the biggest thing for now
 - Globus authentication
- TERRA-REF Use Case:
 - Auth into Workbench: ideally this would SSO with Clowder/BETYdb/ROGER – same user/password. In the end, this is LDAP/NCSA Identity
 - Container and filesystem permissions
 - RunAs me
 - Write files as me in my project
 - PAM/SSSD in container
 - Restrict access to some data to some users
 - See sample data, but not full set
 - SciServer – ACLs and data is only mounted in container if you are authorized
 - One cluster: Workbench + Extractors
 - Extractors need read-write access to core filesystems
 - Users can have RO to core filesystems (shared data)
 - TERRA has users directory on ROGER that I can mount via NFS
 - Replace GlusterFS
 - SSO via Oauth: need to do the work
 - Authorization: where it all gets hairy
 - Need ACLs
 - Handle UID/GID
 - Max needs to be able to run extractors
 - Today, he needs to ssh into master
 - Pile of extractor yaml files in admin repo
 - "terraref" namespaces
 - Extractors nodeSelector – nodes have RW access to the core data
 - Hardcoded UIDs into the container to run as filesystem owners
 - Force the "RunAsX" model of OpenShift
 - Kubernetes RBAC

Sprint 34 priorities:

- SC17 demo
- Production issues
- OpenShift eval
- Easy Kubernetes install on OpenStack?
- Oauth: things to do
 - LDAP authentication?
 - Change ApiServer
- Authorization model
 - TERRA-USE case (who has access to what data)
 - Who gets to control that (admin role)