2017-09-29 Planning meeting

Adhoc whiteboard meeting with Craig and Mike

Authentication

- Mike has a proof-of-concept using nginx ILB external auth feature for ☑ **NDS-813** - Implement workbench authentication module `RESOLVED`
- Discussed priorities going forward
    - Refactor UI to use token generated by new login module
    - Refactor API server to use SSO instead of basic auth
        - Auth into API server via CLI
        - Creation of secrets/ingress rules
        - At this point, we should be able to test end-to-end SSO with custom auth instead of HTTP-basic (no more double login!)
    - Handle authorization
        - Need some way to prevent user1 from accessing user2's stuff
        - Likely need to add new API endpoint to check_access given hostname and namespace/token
    - Oauth2 prototype
    - Refactor UI look/feel
- At this point, we still don't know enough about auth to go forward with OIDC/Oauth support
    - For TERRA-REF case, LDAP integration is more appropriate than Oauth, since ROGER uses LDAP and all users must have an NCSA LDAP account
        - Jim Basney says that CILogon can provide LDAP groups in Oauth scopes, but this must be configured per-client
        - Oauth scopes are not standardized
        - Globus Auth apparently has decent documentation https://docs.globus.org/api/auth/specification/#introduction
        - Talking with Nathan Tolbert, the usual practice inside NCSA is to use LDAP for authentication (delegates to Kerberos) and to use the groups for coarse-grained authorization. Applications (e.g., Atlassian suite) have local roles/ACL management interfaces, but support mapping to LDAP users and groups.
    - Still not clear when we would want to use CILogon for oauth
        - Maybe best example is public beta
        - Lose notion of groups/roles, so would map user to ACL directly.
        - Maybe OK for public beta, since we system end to end, as opposed to ROGER where we're working with existing infrastructure.

Priorities

- Einstein Toolkit tutorial server
    - Mostly done, in final testing now
    - Found bug (🔲 **NDS-1037** - API server creates namespace of unapproved accounts on restart `RESOLVED`)
    - Will need ongoing support
- BBD workshop
    - Demo/presentation
    - Slide to MBDH by Tuesday
    - Lock down Mongo, change collection name
    - Likely will have 5-10 users after workshop (deal with as needed)
    - Pilot description
    - Upgrade beta?
- Other issues:
    - Need further research on auth/authorization model
    - Need further investigation of root user /RunAsUser problem
    - Network bug

Industry demo

- Single node install industry.ndslabs.org, demo only
- Will have Zeppelin, Jupyter, RStudio, Cloud9 and Xpra containers
- Potential stories
    - User can run analysis on existing data via Spark
        - Thinkchicago 2018 solution
        - Hackathon
    - "Ellen" – Genomics developer
        - Today: local IDE git commit/push ssh spark git pull mvn build spark-submit
        - Today: local IDE scp spark ssh spark spark-submit
        - Demo
            - Cloud9Spark spark-submit
    - Choose your notebook
        - Jupyter+Spark? RStudio+Spark? Zeppelin+Spark? Cloud9+Spark – we don't care!