

# 2016-02-08 NDS Labs Planning

2016-02-08

Discussion of workflow:

- A ticket is owned by Mike (implementation + documentation)
- Create sub-task for unit testing, assign to Craig
- Once testing complete, marked as resolved, pull request is issued
- Create sub-task for pull-request
- Once pull-request merged, close tickets.

Prioritization of high-level tasks (will add tickets to JIRA with more detail, then provide estimates)

1. Kubernetes cluster
  - a. On Openstack
    - i. ☒ **NDS-98** - Easy deployment of Kubernetes cluster to Nebula **CLOSED** : Deployment of Kubernetes cluster on multiple VMs on OpenStack
      1. New Task: Testing (due 2/?)
      2. New Task: Update labs portal documentation (assigned to Mike, 2/?)
  - b. Development stack- NDS Dev (David + Mike ~?) (single node)
    - i. ☒ **NDS-82** - Clowder on Kubernetes **CLOSED** : Allows anyone to bring up Kubernetes
      1. New tasks: Testing (due 2/9- Craig), documentation (done - Mike)
2. NDS Ops services (the thing that let's you deploy #3)
  - a. ☒ **NDS-92** - Enable NDS Ops services **CLOSED** : Set of containers (logging, monitor, alert) = ELK
    - i. New task: Testing (needs to be done as of 2/8), documentation (done)
  - b. Ability to scale it up and down (David)
    - i. New task: Allows anyone (NDS Labs user) to scale up their cluster if they need to – adding more OpenStack resources. (Not for NDS5)
3. Demo #1 (TERRA)
  - a. ☒ **NDS-105** - Elasticsearch / terratoolsrv Service Integration **CLOSED** : Elasticsearch + Jupyter + Terratoolsrv
  - b. Scale (elasticity, sharding mongo, sharding rabbit) see how it fits
  - c. Script/documentation of what the demo is (done by Kenton in diagrams, but we need to review)
4. Demo #2 (DataVerse)
  - a. ☒ **NDS-103** - DataVerse support in NDSLabs **CLOSED** : Adding DataVerse support
  - b. Script/documentation of what the demo is (review Kenton's diagrams)
  - c. Open Question: How can we make this more compelling?
5. Backend APIs
  - a. Administrative API (provision resources, new namespace)
  - b. Project user API
  - c. ☒ **NDS-108** - API/CLI prototype **CLOSED** : API/CLI prototype
6. Deployment GUIs
  - a. ☒ **NDS-93** - NDS Labs: GUI Prototype **CLOSED** 1 day prototype/proof of concept (Mike)
  - b. Actual GUI (est TBD)
7. Do the TERRA demo validation (end-to-end walkthrough, Dev + OpenStack envs)
8. Do the DataVerse demo validation (end-to-end walkthrough)
9. Save it and make it a container, add it to the services (projects?)
10. Workshop preparation (presentations)

2016-01-21

Notes from discussion about tasks for upcoming NDS5 workshop.

Prioritization of high-level tasks (will add tickets to JIRA with more detail, then provide estimates):

1. Kubernetes cluster
  - a. on Openstack (99%) (David ~4hrs) (multi node? need to confirm)
    - i. Need to clean up existing tickets.
    - ii. New Task: Easy deployment of cluster on openstack. Allows anyone who wants to bring up a project to boot up a Kubernetes cluster on OpenStack (nebula only) (due 1/29)
    - iii. New Task: Testing (due 2/5)
    - iv. New Task: Update labs portal documentation (assigned to Mike, due 2/5)
  - b. Development stack- NDS Dev (David + Mike ~?) (single node? need to confirm)
    - i. NDS-82, 80%: Allows anyone to bring up Kubernetes
    - ii. Two parts:

1. getting our devstack sorted out, how people will use ndsdev – 1/29
2. Clowder under Kubernetes (95% done with the configuration) – will be done by 1/29. Outcome: kub config (yaml) that brings up an instance of vanilla Clowder that works end-to-end.
  - iii. New tasks: Testing (due 2/5- Craig), documentation (due 2/5 - Mike)
2. NDS Ops services (the thing that let's you deploy #3)
  - a. NDS-92: Set of containers (logging, monitor, alert) = ELK? (David + Craig) - (1/29)
    - i. New task: Testing, documentation
  - b. Ability to scale it up and down (David)
    - i. New task: Allows anyone (NDS Labs user) to scale up their cluster if they need to – adding more OpenStack resources. (Not for NDS5)
3. Demo #1 (TERRA)
  - a. Write additional pod specifications
    - i. identify which ones
  - b. Put those all together + glue
    - i. More containerizing for Clowder
    - ii. Scale (elasticity, sharding mongo, sharding rabbit) see how it fits
4. Deployment GUI
  - a. 1 day prototype/proof of concept (Mike)
  - b. Actual GUI (est TBD)
5. Do the TERRA demo validation (end-to-end walkthrough, Dev + OpenStack envs)
6. Save it and make it a container, add it to the services (projects?)
7. Demo #2 (MDF)
  - a. MDF
    - i. Workout with Globus
  - b. Metasearch
    - i. Flesh this out (Craig)
    - ii. Write the pod specifications
    - iii. Put those all together + glue
8. Workshop preparation (presentations)

There was a discussion of different services.

Question was asked whether the services list contains high-level component (e.g., clowder) or everything (e.g., clowder, rabbit, mongo, elastic). Answer: All.

Discussion of Kubernetes -- two steps: 1 deploy, 2 scale/replication,

Things to discuss later/

- Implementation that spans Nebula and Amazon? Once you get to #4 – all you need is the account. Idea: Terra is running on Nebula, migrates to Google Cloud?
- Is storage more important?