Jun Huh
Product Manager, NeSI

Claire Rye
Product Manager, NeSI

Journey of provisioning HPC with Infrastructure-as-Code approach

eResearch Australasia 2024

Background

- NeSI is NZ’s national HPC provider
- Currently in the process of provisioning the third iteration of its HPC infrastructure.
- For the researchers
  - Continuity of their research
  - Improvements with the new generation of hardware
  - Adapting to the growing needs around data and AI

NeSI HPC Compute Refresh

- 10k new cores (AMD Genoa 9634 84-core 2.25-3.7GHz)
  - ~48 nodes standard HPC @ 2.2GB RAM per core
  - 1.5 - 2 x FLOPS of Mahuika Extension TBC
  - 8 nodes high memory @1.5TB RAM
- GPU nodes
  - New homes for PCIe A100s
  - New H100 GPU nodes
  - New L4 nodes GPU (ML focused)

NeSI HPC Compute Refresh

- WEKA cluster
  - 500TB flash with tiering to object storage
- 4PB on prem Ceph (object storage)
- Versity archival storage
  - SAN Array
    - Metadata volume 3.84 TB usable capacity
    - Data volume 50 TB usable capacity
  - Tape Library
    - 6x LTO9 Tape Drives (added to Spectra existing T950v at TDC)
    - Initial capacity of 1 PB

As-Code approach

- Learning from our partnership with AgResearch
  - AgResearch eResearch Infrastructure
  - HPC provisioned on bare metal with StackHPC Slurm Appliance
- Enabling iteration and collaboration across teams
- Using cloud VMs while physical set up is going on in parallel
- CI/CD

Slurm cluster ‘steel thread’

- Steel thread - end-to-end
- Using VMs to learn and break down the problems
- VMs → bare metal → MaaS (Metal-as-a-Service)
  - MaaS from Canonical
  - Rolling unattended updates of compute nodes
- Image building pipeline
  - Rocky image → Base HPC image → tailored images based on use cases
- End-to-end deployment on the cloud to enable isolated testing and iteration

Slurm cluster ‘steel thread’

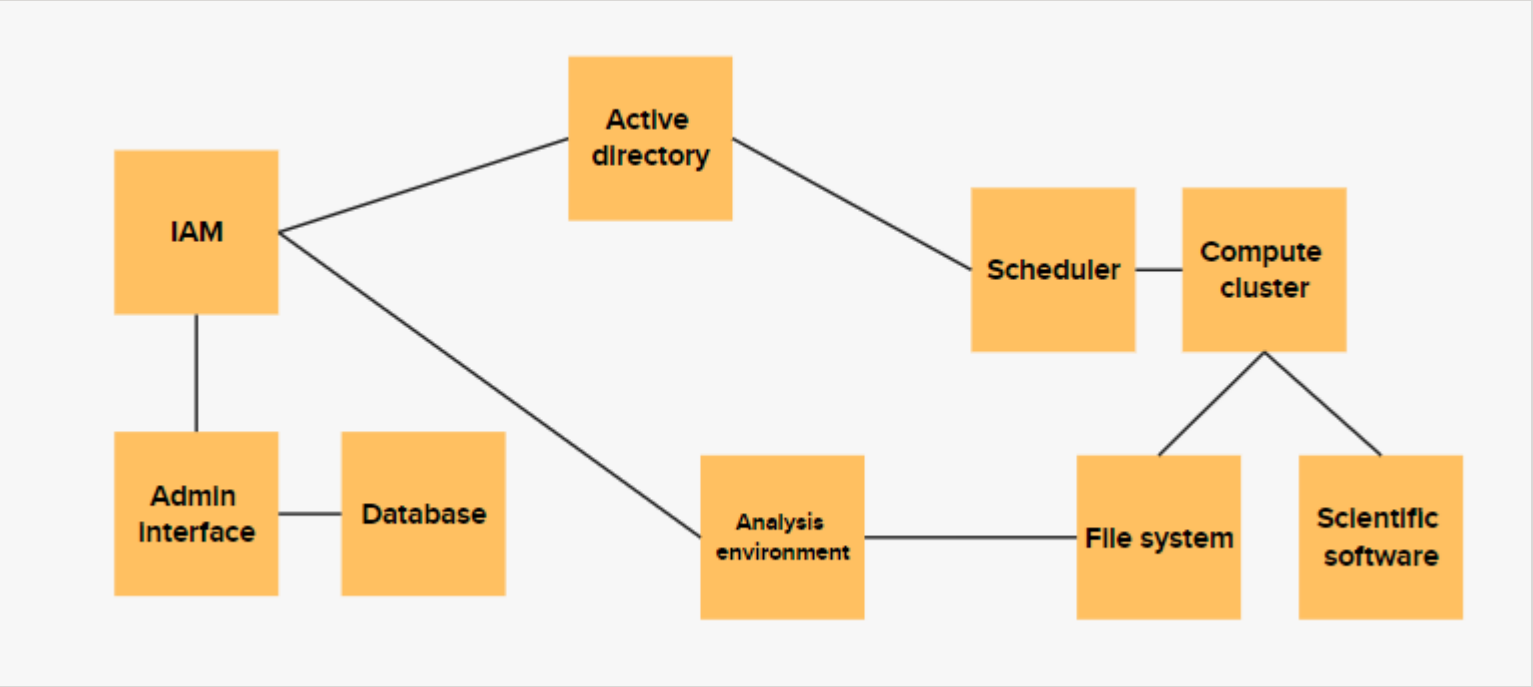

MaaS

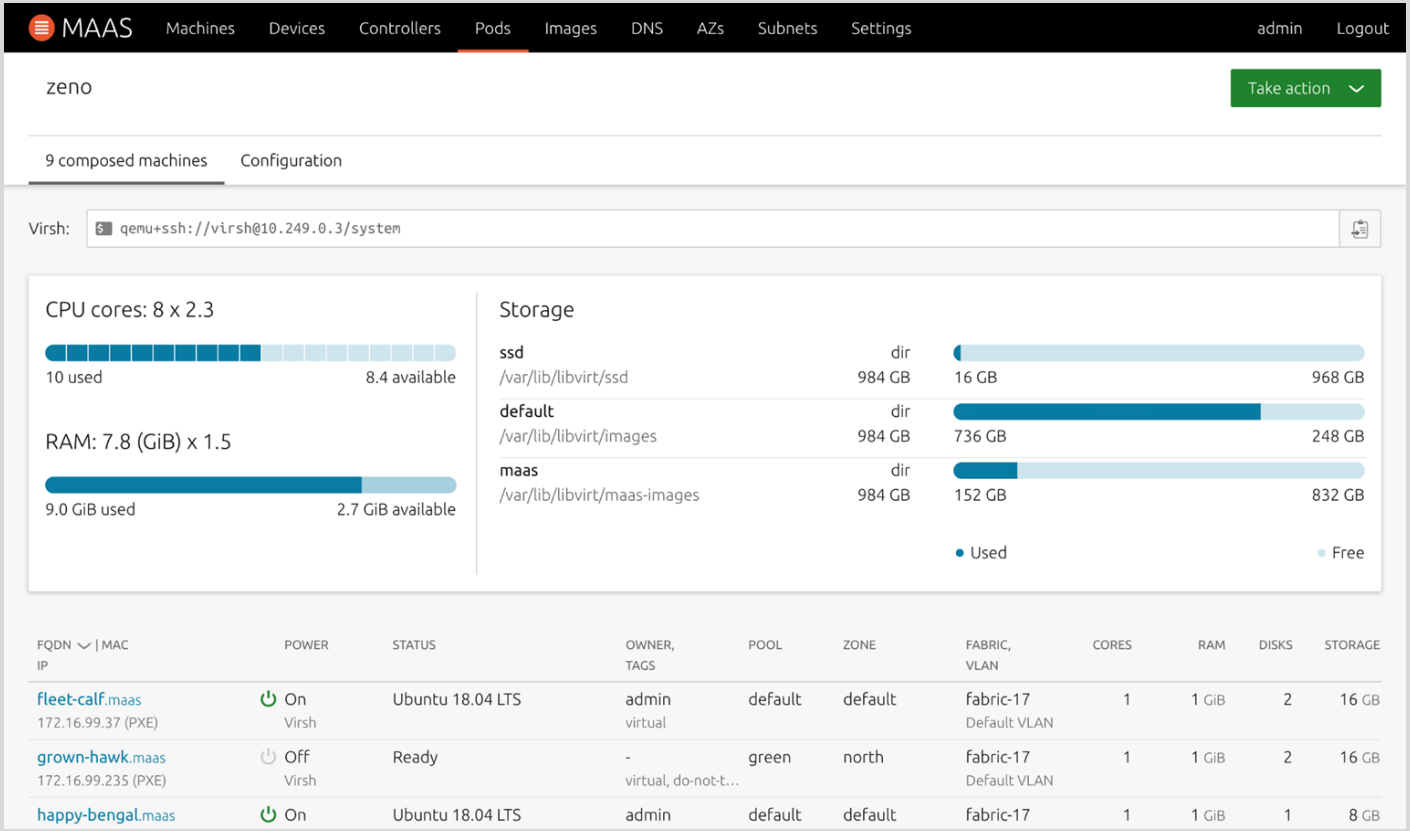

Image credit: https://maas.io/

NeSI OnDemand

- Replacing NeSI's JupyterHub environment for interactive applications
- Open OnDemand, Keycloak, NFS, FreeIPA
- Tooling
    - Terraform, Ansible, GitLab pipelines, K8s, Secret management with Hashicorp Vault (now Google Cloud Vault)

NeSIOnDemand Roadmap

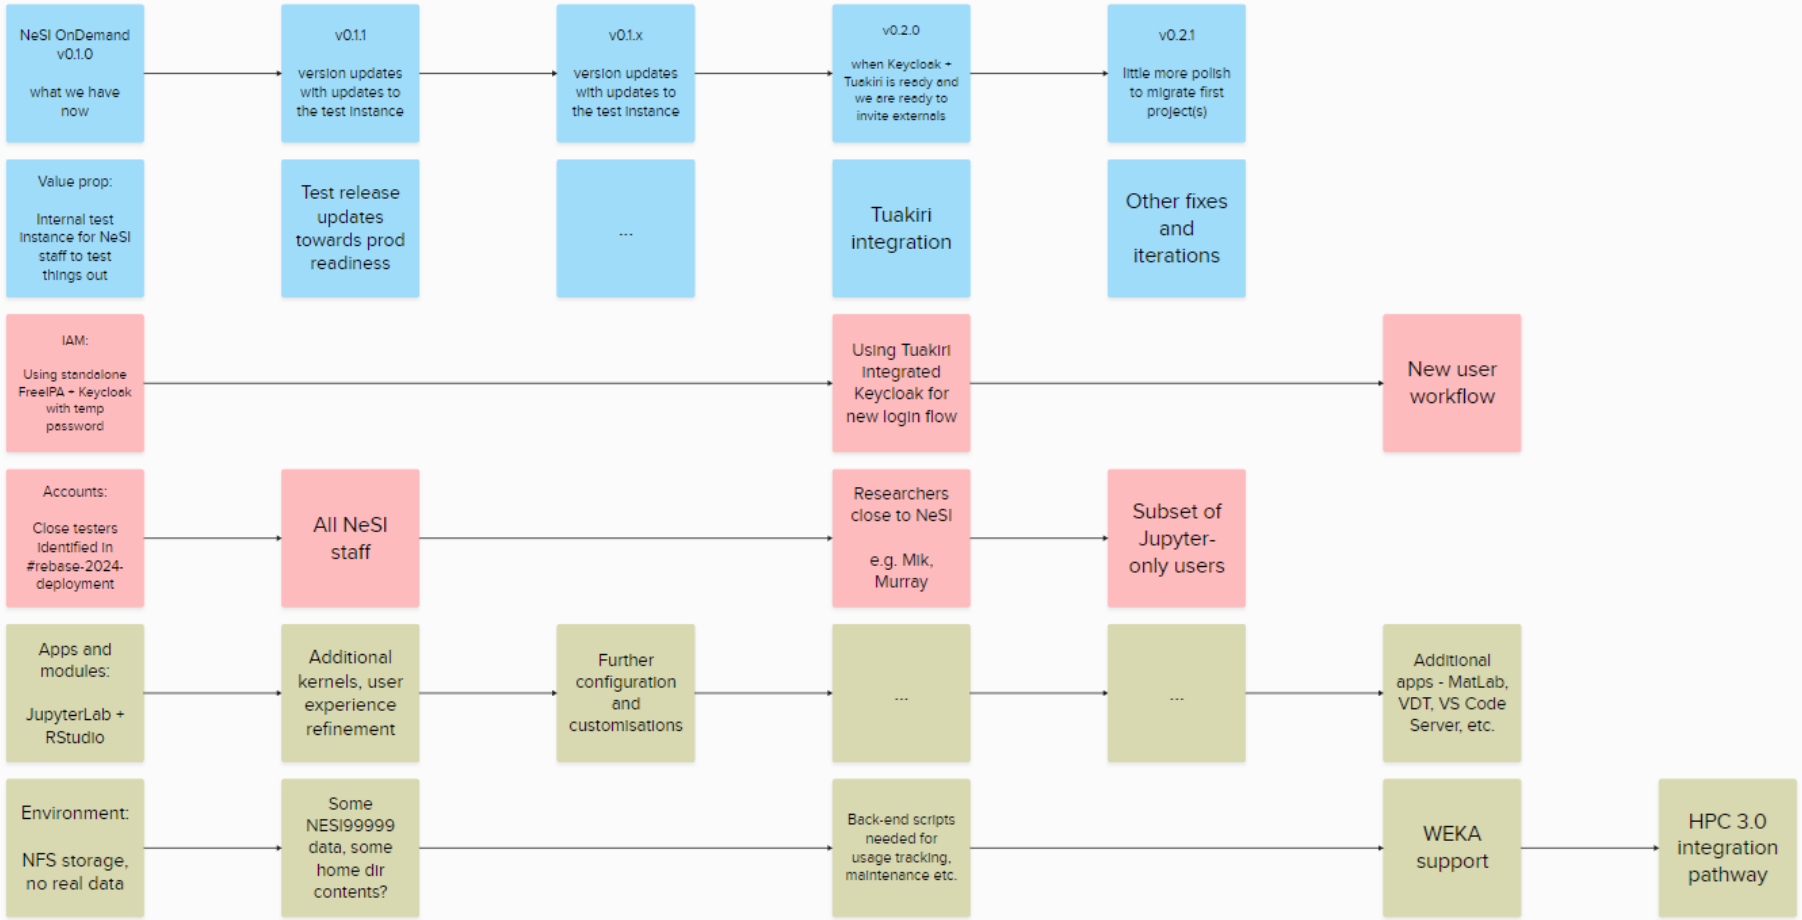

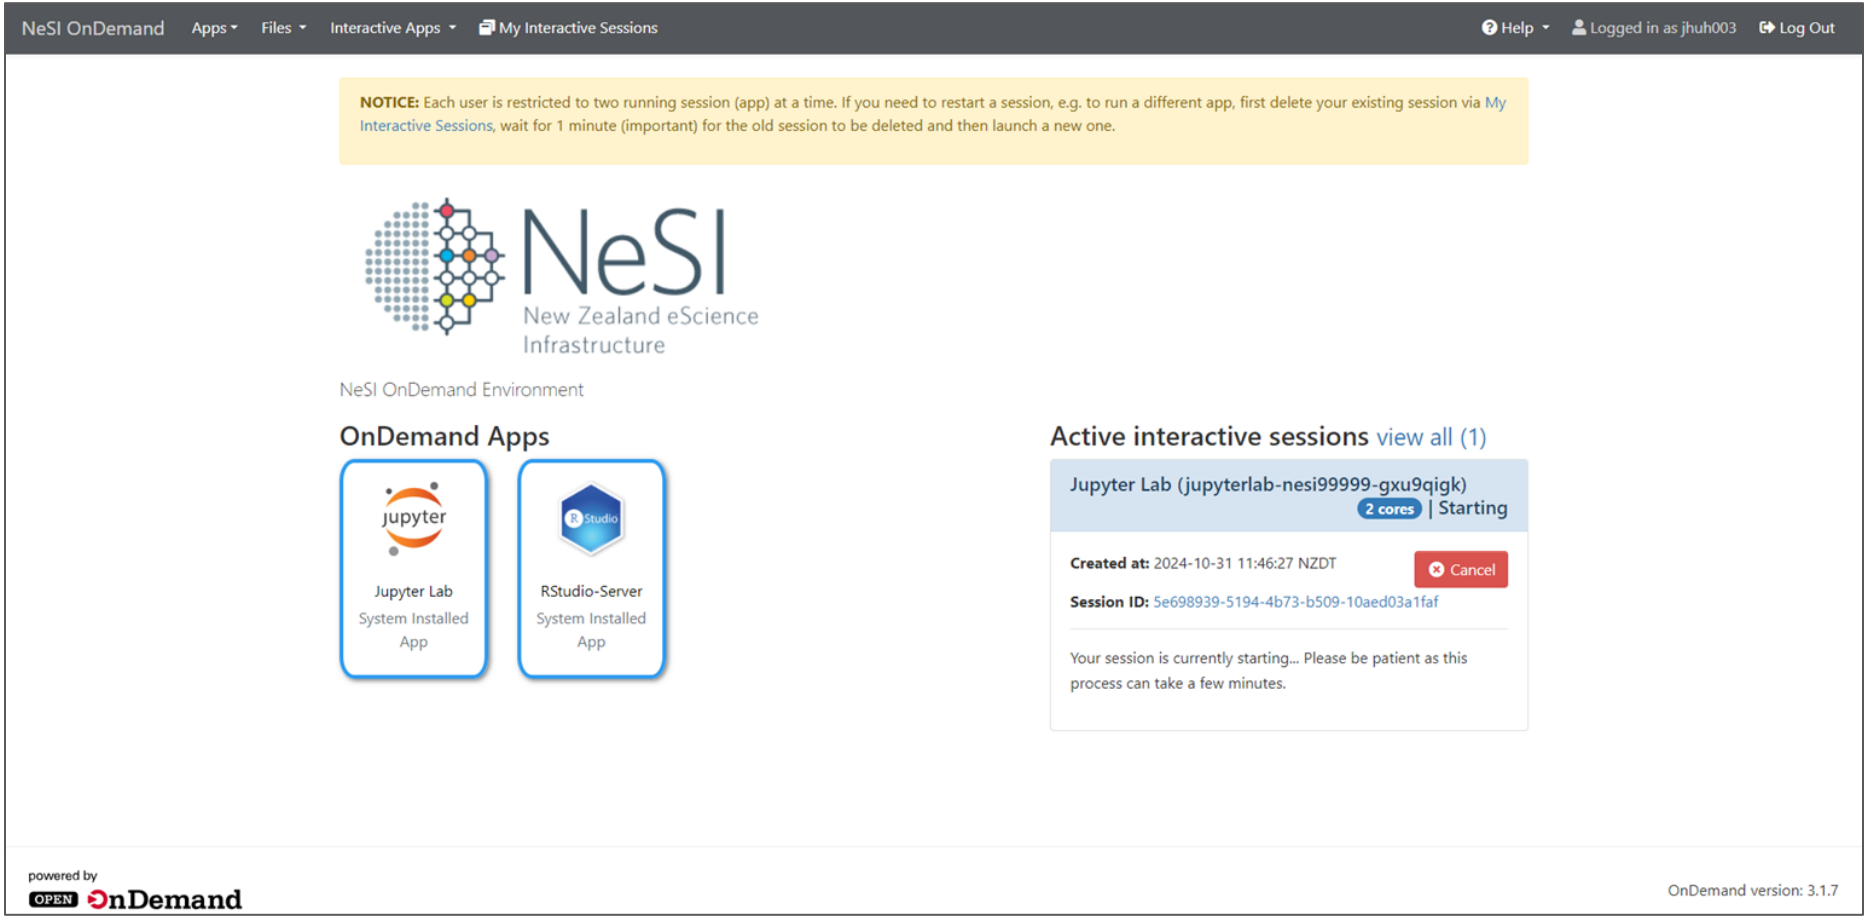
NeSI OnDemand
Apps
Files
Interactive Apps
My Interactive Sessions
Help
Logged in as jhuh003
Log Out
NOTICE: Each user is restricted to two running session (app) at a time. If you need to restart a session, e.g. to run a different app, first delete your existing session via My Interactive Sessions, wait for 1 minute (important) for the old session to be deleted and then launch a new one.
NeSI
New Zealand eScience Infrastructure
NeSI OnDemand Environment
OnDemand Apps
jupyter
Jupyter Lab
System Installed App
RStudio
RStudio-Server
System Installed App
Active interactive sessions view all (1)
Jupyter Lab (jupyterlab-nesi99999-gxu9qigk)
2 cores | Starting
Created at: 2024-10-31 11:46:27 NZDT
Cancel
Session ID: 5e698939-5194-4b73-b509-10aed03a1faf
Your session is currently starting... Please be patient as this process can take a few minutes.
powered by
OPEN OnDemand
OnDemand version: 3.1.7

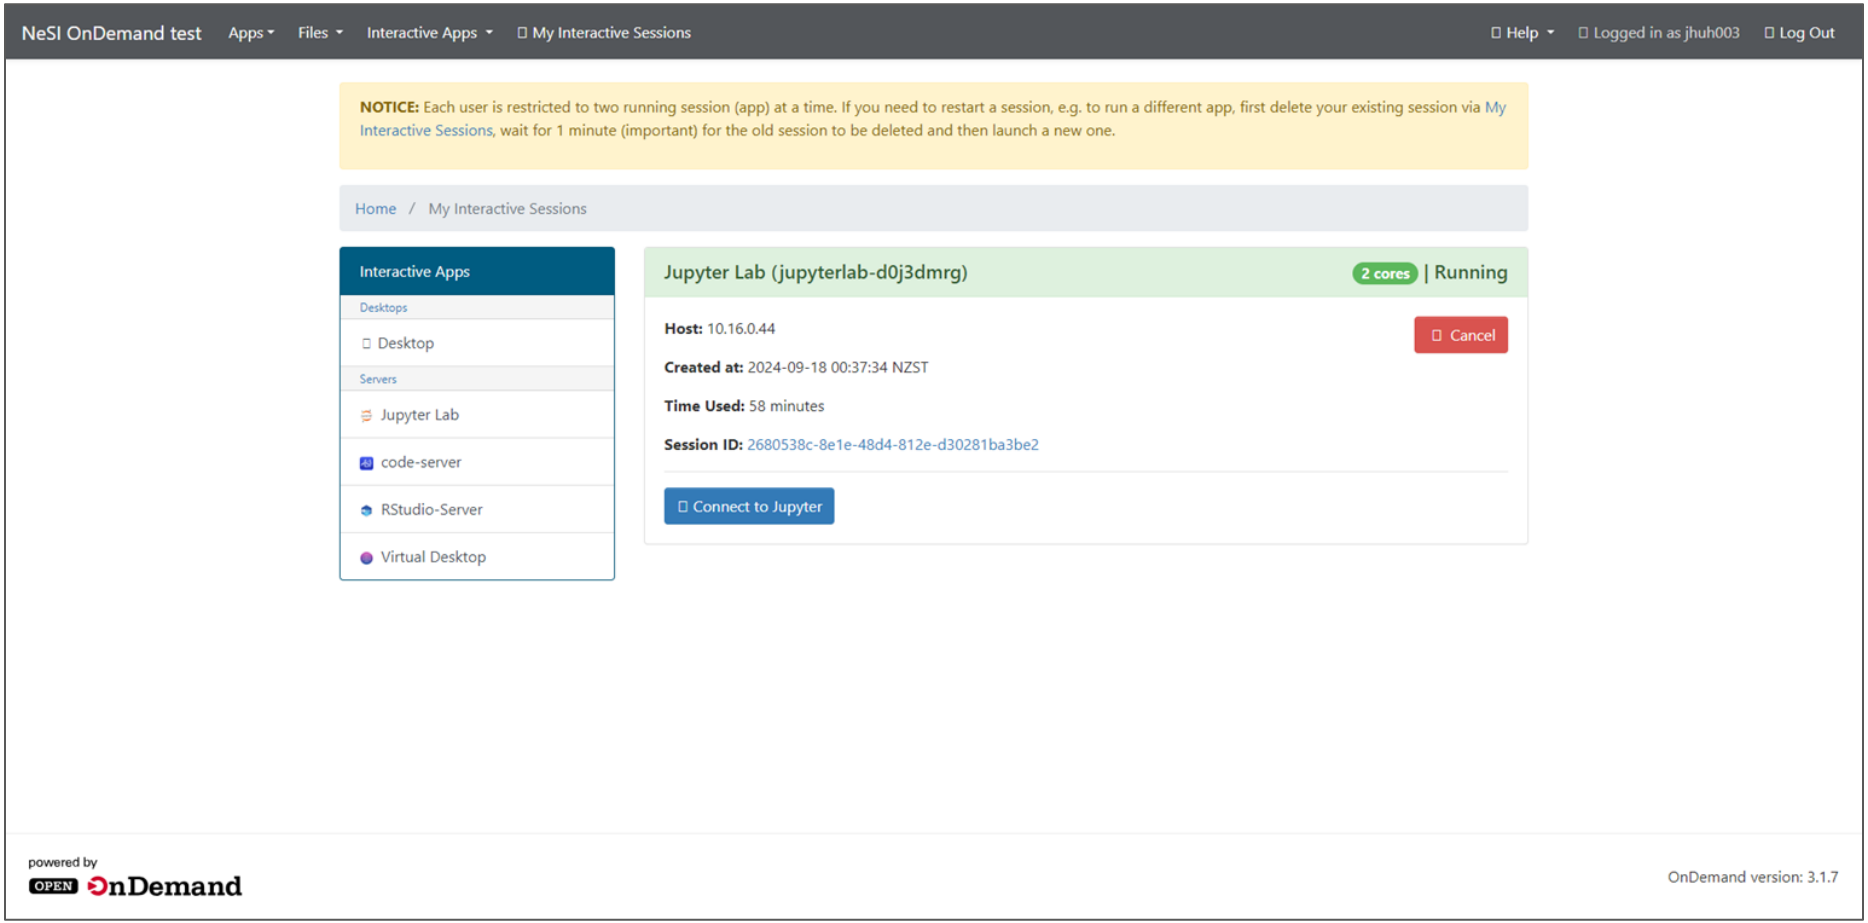
NeSI OnDemand test
Apps
Files
Interactive Apps
My Interactive Sessions
Help
Logged in as jhuh003
Log Out
NOTICE: Each user is restricted to two running session (app) at a time. If you need to restart a session, e.g. to run a different app, first delete your existing session via My Interactive Sessions, wait for 1 minute (important) for the old session to be deleted and then launch a new one.
Home / My Interactive Sessions
Interactive Apps
Desktops
Desktop
Servers
Jupyter Lab
code-server
RStudio-Server
Virtual Desktop
Jupyter Lab (jupyterlab-d0j3dmrg)
2 cores | Running
Host: 10.16.0.44
Cancel
Created at: 2024-09-18 00:37:34 NZST
Time Used: 58 minutes
Session ID: 2680538c-8e1e-48d4-812e-d30281ba3be2
Connect to Jupyter
powered by
OPEN OnDemand
OnDemand version: 3.1.7

File Edit View Run Kernel Git Tabs Settings Help

jhuh003@jupyterlab-d0j3dm | plotly.ipynb

Code | Open in... | Python 3.11.3 (gimkl-2022a)

Filter files by name

/

Name	Modified
00_nesi_projects	8m ago
ondemand	12d ago
plotly.ipynb	14s ago
Untitled.ipynb	2m ago

[1]:
```
import plotly.graph_objects as go
import numpy as np
t = np.linspace(-1, 1, 100)
x = t + t**2
y = t - t**2
fig = go.Figure(data=[go.Scatter(x=x, y=y, mode="lines", line=dict(width=2, color="blue"))],
                layout=go.Layout(width=600, height=600,
                                 xaxis=dict(zeroline=False),
                                 yaxis=dict(zeroline=False),
                                 title="Interactive Plotly Example"))
fig.show()
```

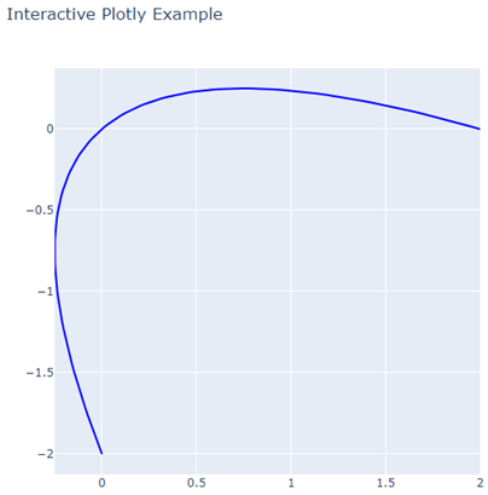

Simple | 1 | 2 | Python 3.11.3 (gimkl-2022a) | Idle | Mode: Command | Ln 11, Col 11 | plotly.ipynb | 0

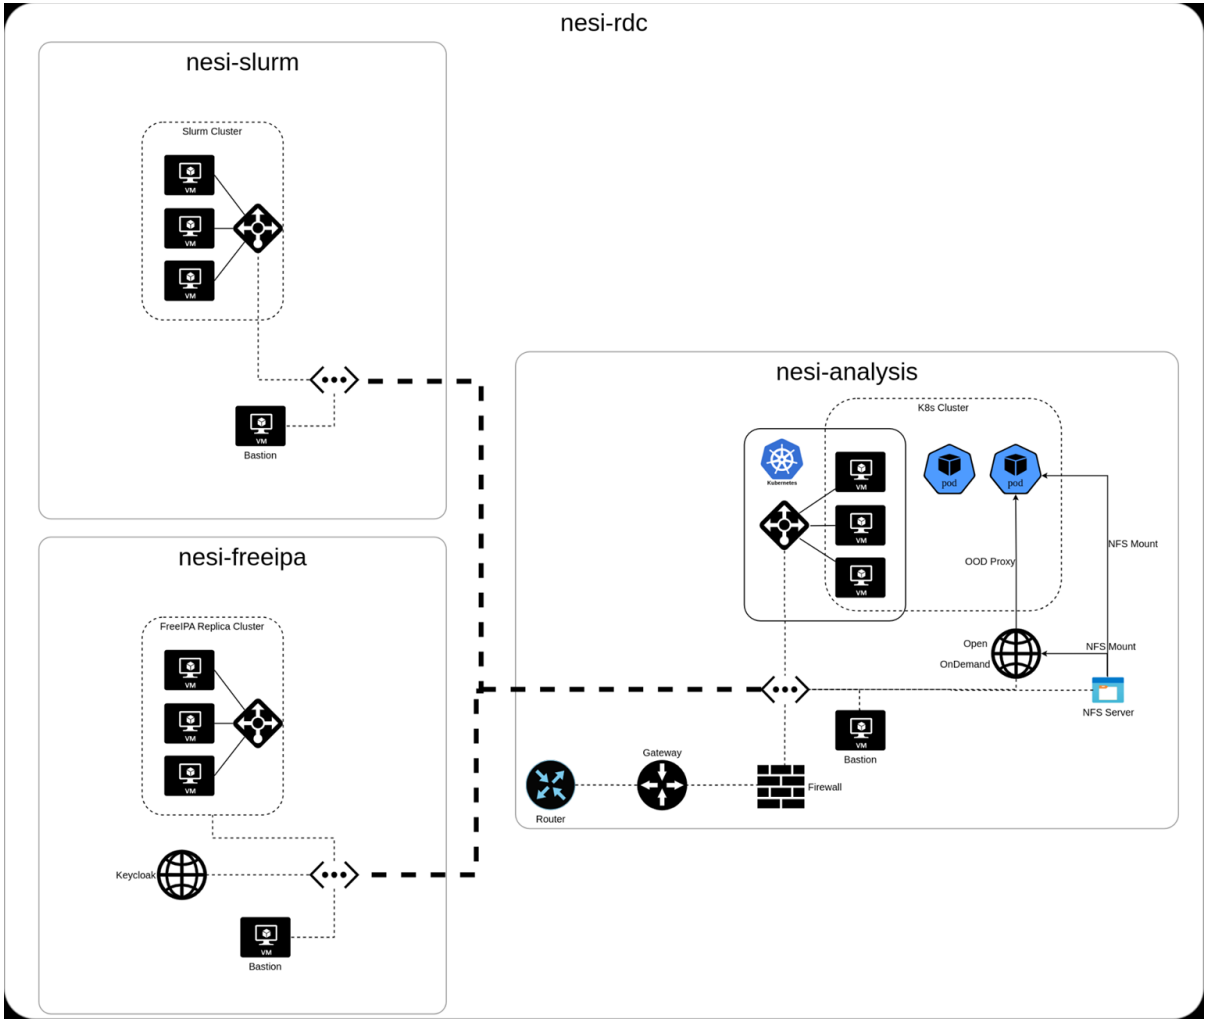
nesi-rdc
nesi-slurm
Slurm Cluster
VM
VM
VM
Bastion
nesi-freeipa
FreeIPA Replica Cluster
VM
VM
VM
Keycloak
VM
Bastion
nesi-analysis
K8s Cluster
Kubernetes
VM
VM
VM
pod
pod
OOD Proxy
NFS Mount
Open
OnDemand
NFS Mount
NFS Server
VM
Bastion
Gateway
Firewall
Router

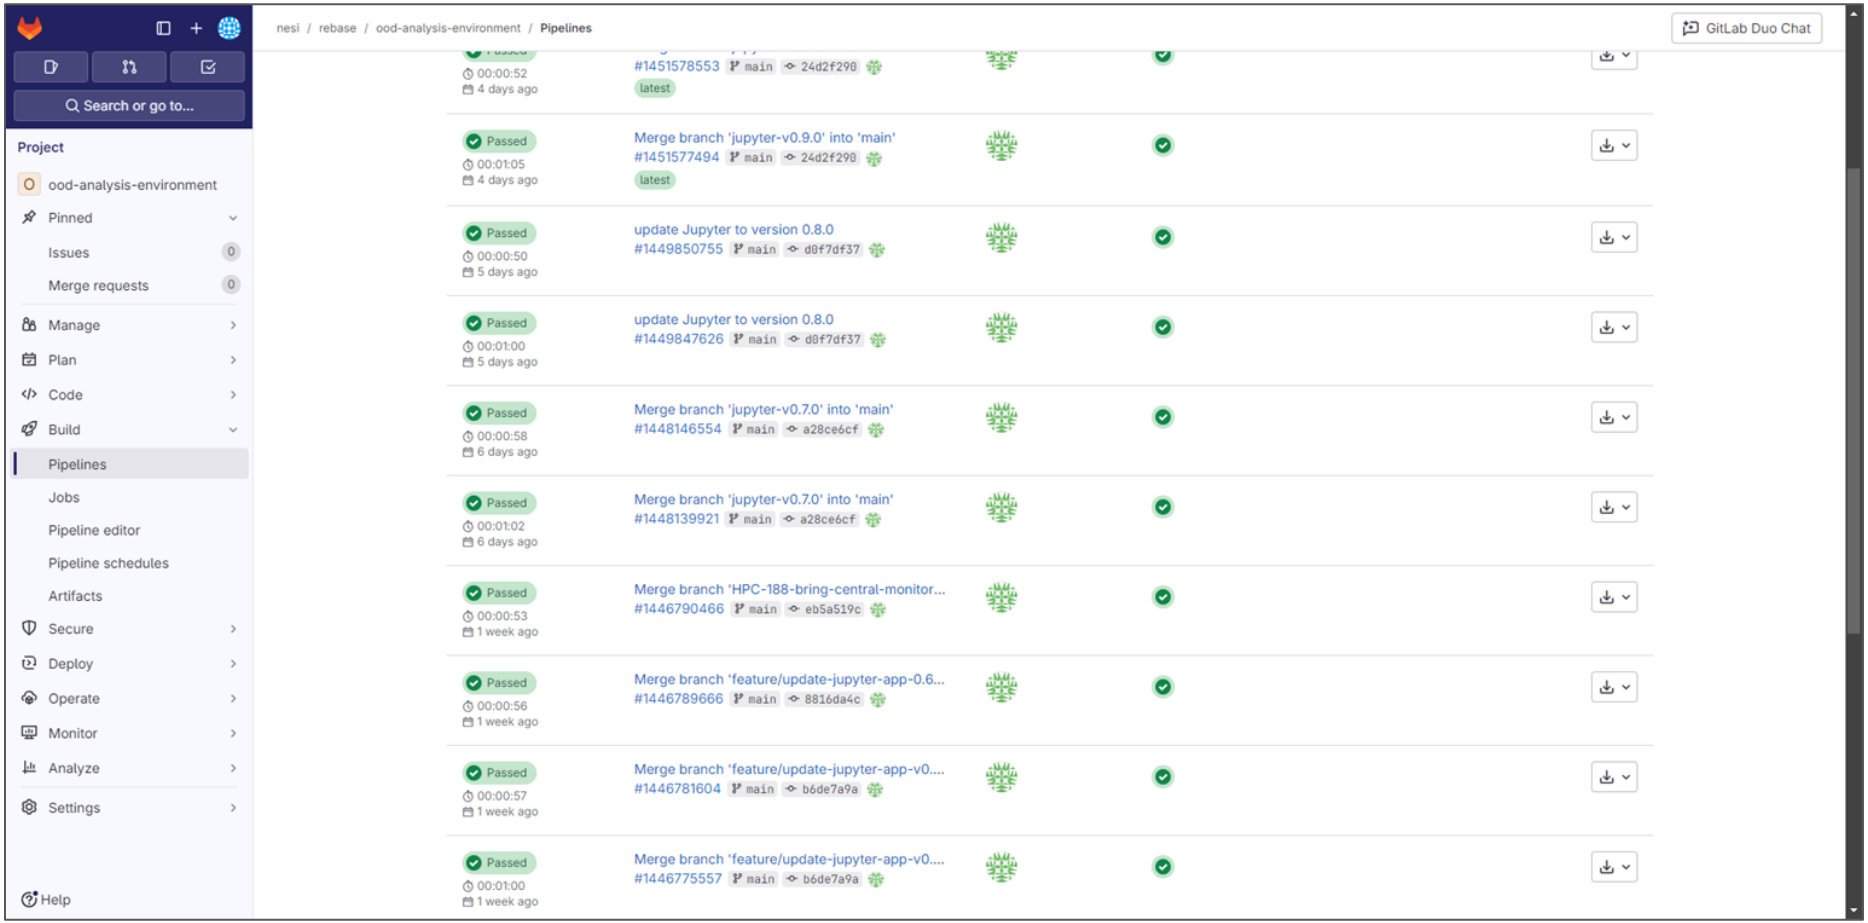
nesi / rebase / ood-analysis-environment / Pipelines
GitLab Duo Chat
Search or go to...
Project
ood-analysis-environment
Pinned
Issues 0
Merge requests 0
Manage
Plan
Code
Build
Pipelines
Jobs
Pipeline editor
Pipeline schedules
Artifacts
Secure
Deploy
Operate
Monitor
Analyze
Settings
Help
00:00:52
4 days ago
#1451578553 main 24d2f290
latest
Passed
00:01:05
4 days ago
Merge branch 'jupyter-v0.9.0' into 'main'
#1451577494 main 24d2f290
latest
Passed
00:00:50
5 days ago
update Jupyter to version 0.8.0
#1449850755 main d0f7df37
Passed
00:01:00
5 days ago
update Jupyter to version 0.8.0
#1449847626 main d0f7df37
Passed
00:00:58
6 days ago
Merge branch 'jupyter-v0.7.0' into 'main'
#1448146554 main a28ce6cf
Passed
00:01:02
6 days ago
Merge branch 'jupyter-v0.7.0' into 'main'
#1448139921 main a28ce6cf
Passed
00:00:53
1 week ago
Merge branch 'HPC-188-bring-central-monitor...
#1446790466 main eb5a519c
Passed
00:00:56
1 week ago
Merge branch 'feature/update-jupyter-app-0.6...
#1446789666 main 8816da4c
Passed
00:00:57
1 week ago
Merge branch 'feature/update-jupyter-app-v0....
#1446781604 main b6de7a9a
Passed
00:01:00
1 week ago
Merge branch 'feature/update-jupyter-app-v0....
#1446775557 main b6de7a9a

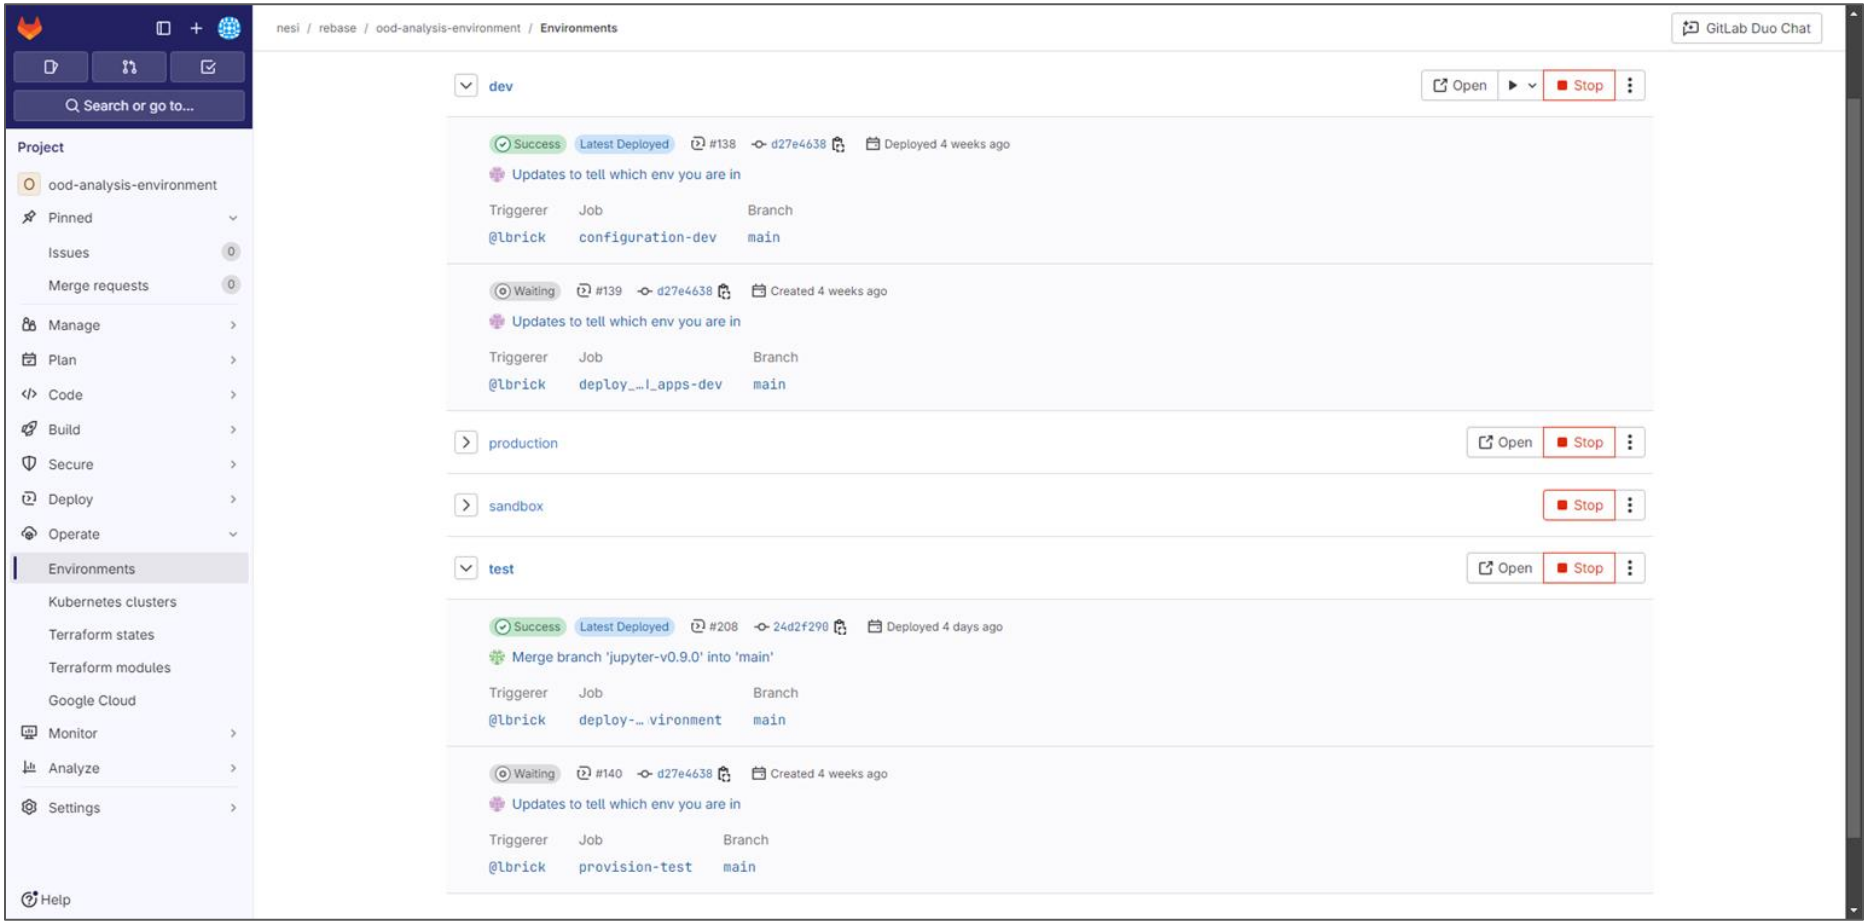
nesi / rebase / ood-analysis-environment / Environments
GitLab Duo Chat
Search or go to...
Project
ood-analysis-environment
Pinned
Issues 0
Merge requests 0
Manage
Plan
Code
Build
Secure
Deploy
Operate
Environments
Kubernetes clusters
Terraform states
Terraform modules
Google Cloud
Monitor
Analyze
Settings
Help
dev
Open
Stop
Success Latest Deployed #138 d27e4638 Deployed 4 weeks ago
Updates to tell which env you are in
Triggerer Job Branch
@lbrick configuration-dev main
Waiting #139 d27e4638 Created 4 weeks ago
Updates to tell which env you are in
Triggerer Job Branch
@lbrick deploy_...l_apps-dev main
production
Open
Stop
sandbox
Stop
test
Open
Stop
Success Latest Deployed #208 24d2f290 Deployed 4 days ago
Merge branch 'jupyter-v0.9.0' into 'main'
Triggerer Job Branch
@lbrick deploy-... vironment main
Waiting #140 d27e4638 Created 4 weeks ago
Updates to tell which env you are in
Triggerer Job Branch
@lbrick provision-test main

Storage updates

- Versity
- Using data temperature
  - Tape speed
  - Less disruption

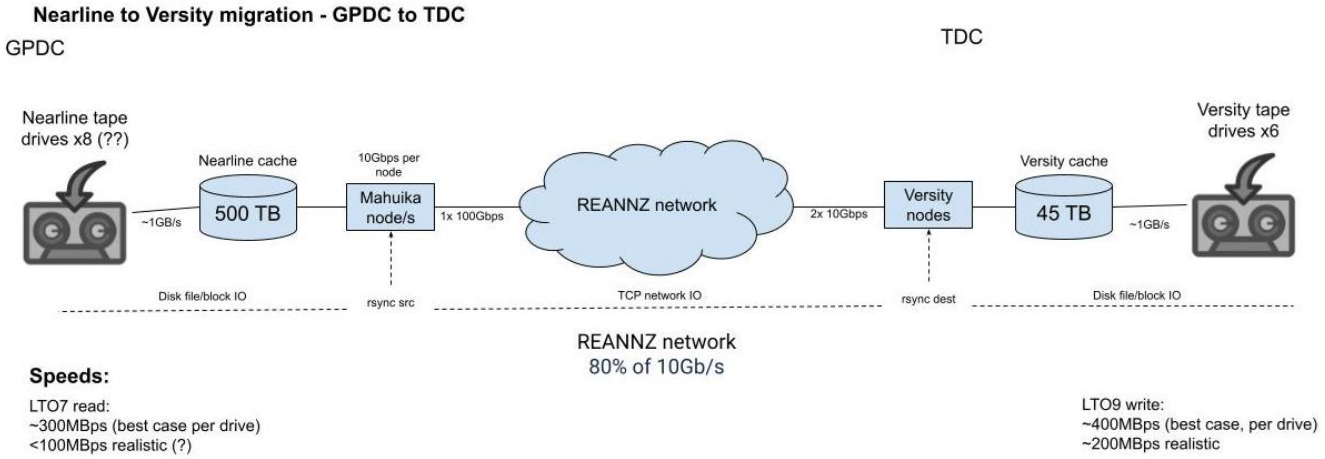

Wrapping up

- Maturity of internal practices
  - Cross team collab
  - Releases, versioning
- Sharing our journey
- What were you major challenges behind setting up your infrastructure?

NeSI
New Zealand eScience Infrastructure

REANNZ

XENON
High Performance Computing

WEKA

Versity

Hewlett Packard Enterprise

NVIDIA

StackHPC

Red Hat

sempre
Always Available

DELL Technologies