

Sunet Drive

An Academic Toolbox with Jupyter
Notebooks

Micke Nordin

Richard Freitag (presenter)

Magnus Andersson

tnc24

RENDEZVOUS À RENNES
Rennes, France | **10-14 JUNE 2024**



Co-funded by
the European Union



Analyze and Publish



General – Sunet Drive and JupyterHub

SUNET Drive

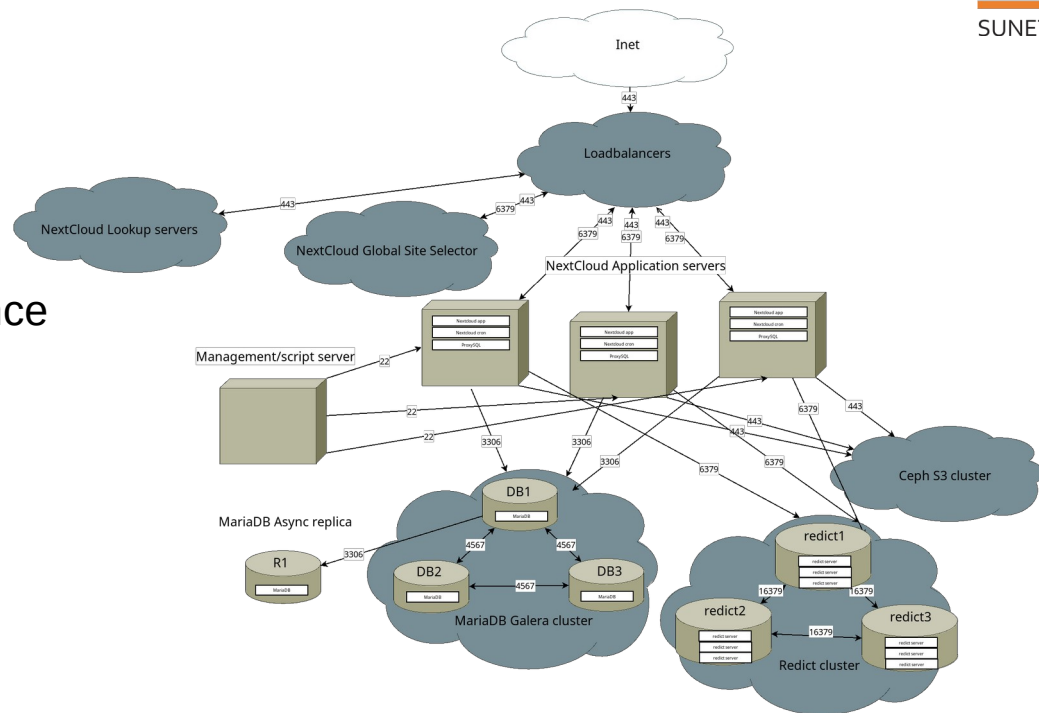
- Based on Nextcloud
- Storage in Ceph S3
- Currently stores ~450 TB of mostly research data
- 20 Universities – growing rapidly
- Upcoming features is:
 - Process research data – crunch the numbers
 - Publish research
 - MFA protected zones

JupyterHub

- a Hub (tornado process) that is the heart of JupyterHub
- a configurable http proxy (node-http-proxy) that receives the requests from the client's browser
- multiple single-user Jupyter notebook servers (Python/IPython/tornado) that are monitored by Spawners
- an authentication class that manages how users can access the system

Nextcloud Architecture

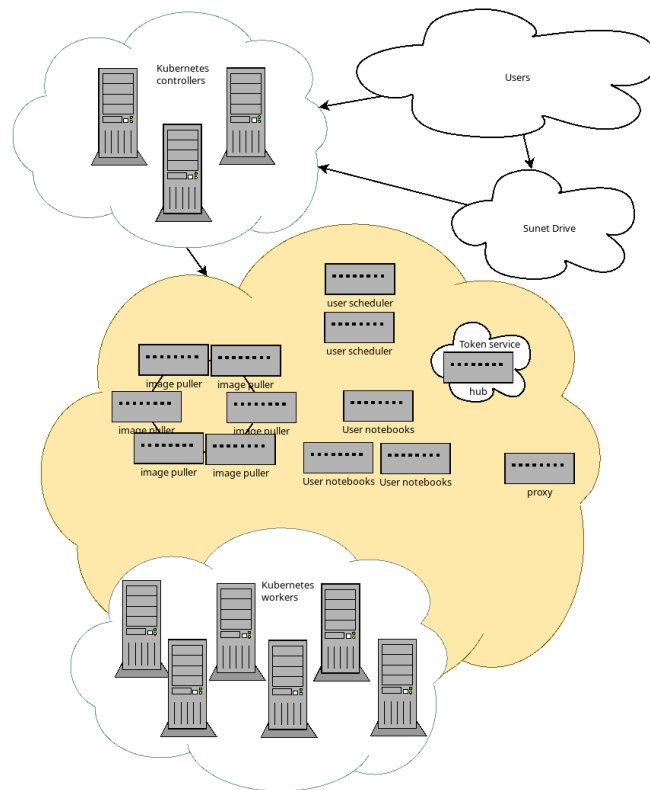
- Load-balanced Sunet Drive Frontend Servers*
- MariaDB Galera cluster as database
- Redis Cluster for improved performance
- All storage in Ceph S3 storage cluster



* <https://github.com/SUNET/sarimner-frontend>

Jupyter Architecture

- Load-balanced controllers*
- Scalable by adding more workers
- All storage synced to openstack cinder



* <https://github.com/SUNET/sarimner-frontend>

Jupyter Integration – How?

- JupyterHub in Kubernetes*
- Nextcloud app for easy access**
- Integration via OpenID Connect**
- Implemented sync via OAuth tokens**

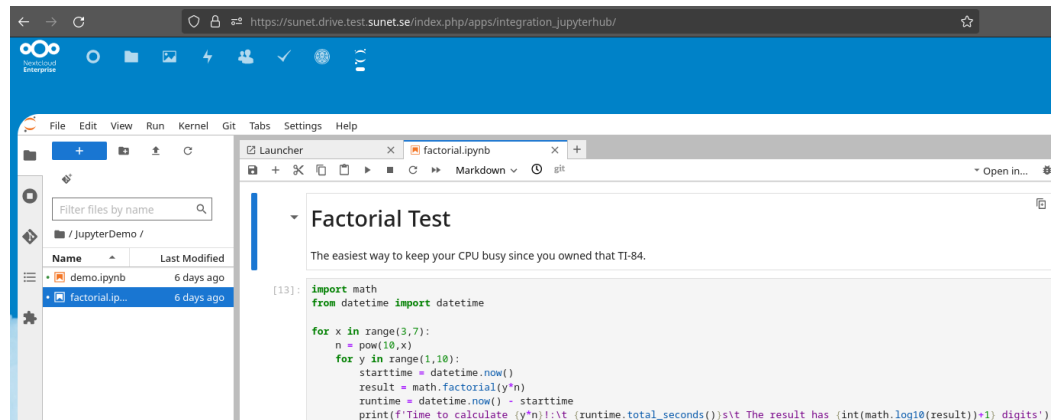


* <https://z2jh.jupyter.org/en/stable/>

** <https://github.com/sunet/nextcloud-jupyter>

JupyterHub Integration in Nextcloud?

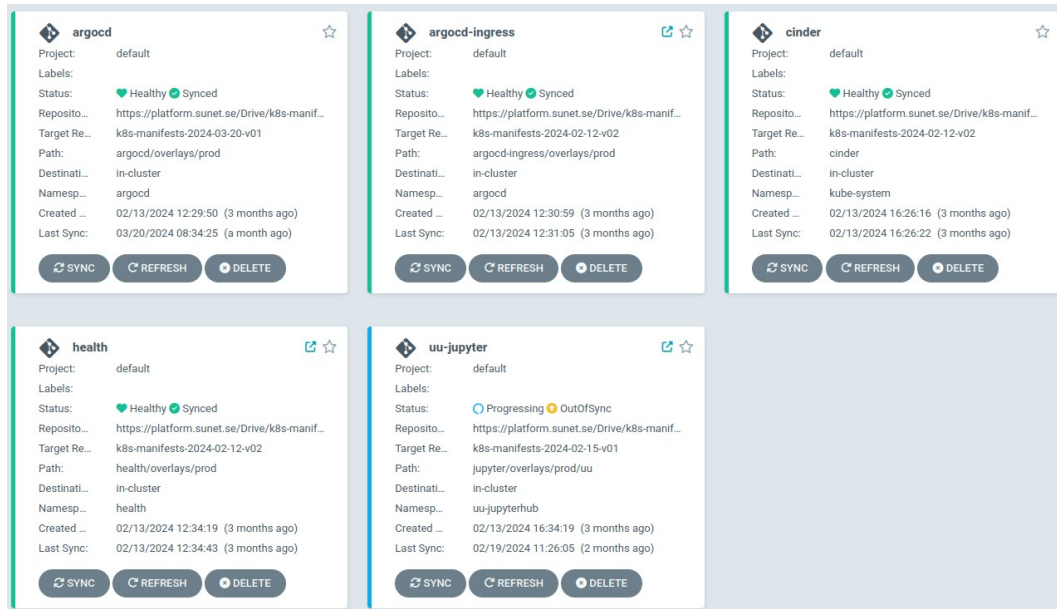
- Integrated in Sunet Drive
- Users can individually execute notebooks
- Users get 10GB “local” Jupyter Storage
- Storage is continuously synchronized between Sunet Drive and JupyterHub
- Other use-cases are being evaluated



Jupyter Deployment and Scalability

60 concurrent users

- 6GB of RAM --> 100MB per user
- 20-25% CPU load --> 1-2% CPU per user per vCPU
- CPU overprovisioning show no significant scheduling penalties
- 16vCPU/32GB could host 100-150 idling jupyter-servers, with 30-40 concurrent users per kubernetes-worker
- Memory seems to be the biggest concern



Deployment Name	Status	Namespace	Created	Last Sync
argocd	Healthy Synced	argocd	02/13/2024 12:29:50 (3 months ago)	03/20/2024 08:34:25 (a month ago)
argocd-ingress	Healthy Synced	argocd	02/13/2024 12:30:59 (3 months ago)	02/13/2024 12:31:05 (3 months ago)
cinder	Healthy Synced	kube-system	02/13/2024 16:26:16 (3 months ago)	02/13/2024 16:26:22 (3 months ago)
health	Healthy Synced	health	02/13/2024 12:34:19 (3 months ago)	02/13/2024 12:34:43 (3 months ago)
uu-jupyter	Progressing OutOfSync	uu-jupyterhub	02/13/2024 16:34:19 (3 months ago)	02/19/2024 11:26:05 (2 months ago)



Thank you

Any questions?



SUNET

kano@sunset.se

freitag@sunset.se (presenter)

mandersson@sunset.se

tnc24

RENDEZVOUS À RENNES

Rennes, France | **10-14 JUNE 2024**



Co-funded by
the European Union

