Blue-sky and the cloud: Container Orchestration for eResearch

Mark Endrei

Nishanthi Dasanayaka

Research Computing Centre
The University of Queensland

Why containers?

For eResearch

- **Rapid deployment**
- **Isolation**
- **Portability**
- **Scaling**
- **Resource management**
- **Lifecycle management**
- **Reproducibility**

Why Nectar?

Containers help harness full potential of your cloud investment

Containers Orchestration Service

- **OpenStack COE (Magnum)**
- **Kubernetes**

Open source

- **Transparency**
- **Collaboration**

Community

- **ARCOS**

Container Orchestration Services

Using Kubernetes without running Kubernetes

OpenStack Magnum orchestrates

- **Nova compute**
- **Cinder volume storage**
- **Neutron network and load balancers**

Cluster management

- **Create and resize**
- **Node groups**
- **Templates**

Registry

- **Nectar Harbor image store, sign, and scan**

ARCOS support service

Case Study

XNAT medical imaging eResearch platform

Initial priorities for migration

To Containers Orchestration Service

eResearch application services

- **Seamless research data access**
- **Interactive, batch, user-demand workloads**
- **Project lifecycle support**

Infrastructure services

- **Backup**
- **Logging**
- **Monitoring**
- **Notifications**
- **Standard platform**

Data Access and Workload Distribution

Using Scheduling Constraints

Researchers/Collaborators and Clinical Sites

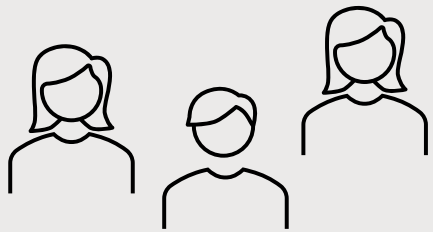

Kubernetes Cluster

Interactive

XNAT

Non-interactive

DICOM Gateway

XNAT Shadow

User Jobs

User Commands

GPU User Jobs

Jupyterhub

Notebooks

HPC Cluster

slurm workload manager

User Jobs

Research Data Manager

Researcher Data Collections

Deployment

Using Helm and Kustomize

Benefits

- **Quick deployments**
- **Public repos**
- **Re-use across environments and teams**

Challenges

- **Secret management**
- **Edge cases – use kustomize**
- **Check deployment manifests**

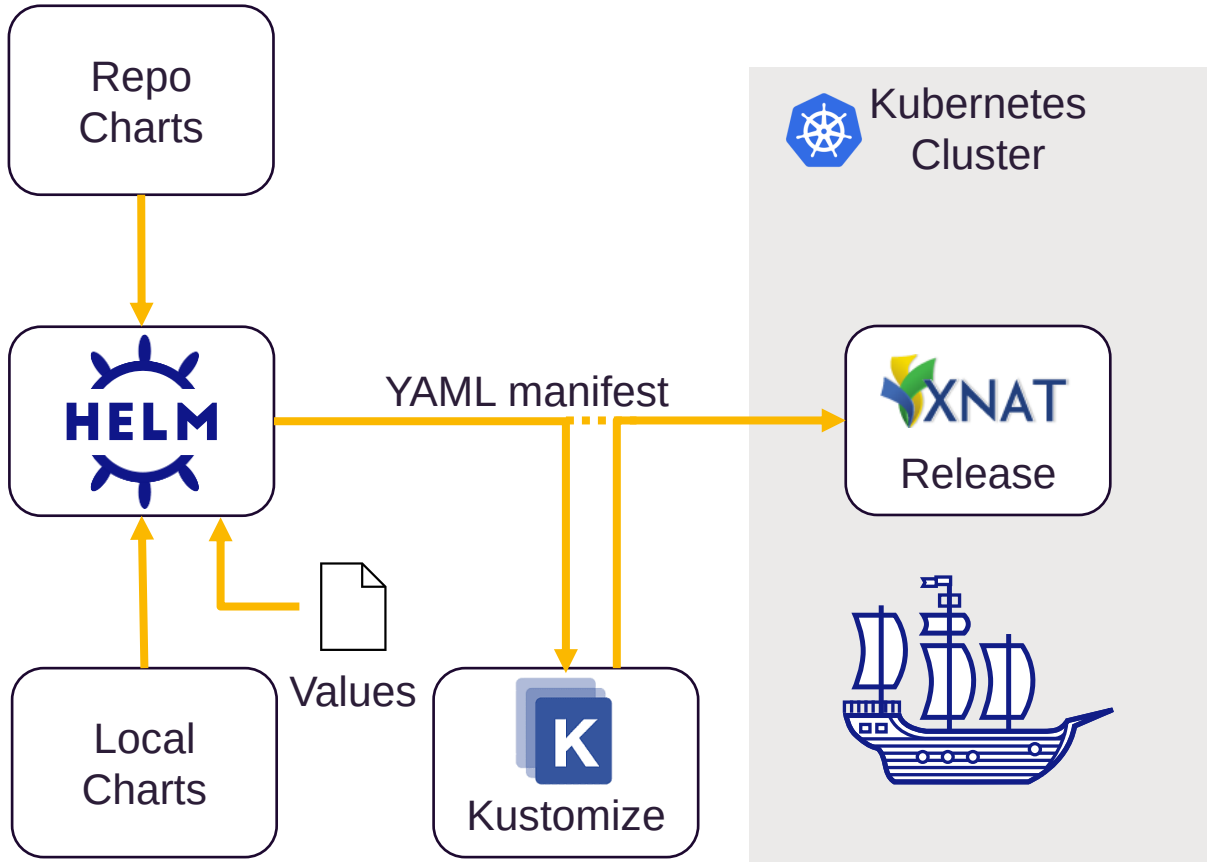

Backup

Using Duplicity

```
{
  "app": "xnat",
  "passphrase": " ... ",
  "inputs": [
    {
      "type": "pg_dump",
      "host": "xnat-postgres ... ",
      "db": " ... ",
      "user": " ... ",
      "password": " ... "
    },
    {
      "type": "fs",
      "folders": [
        "/mnt/xnat/prearchive",
        "/mnt/xnat/plugins"
      ]
    }
  ]
}
```

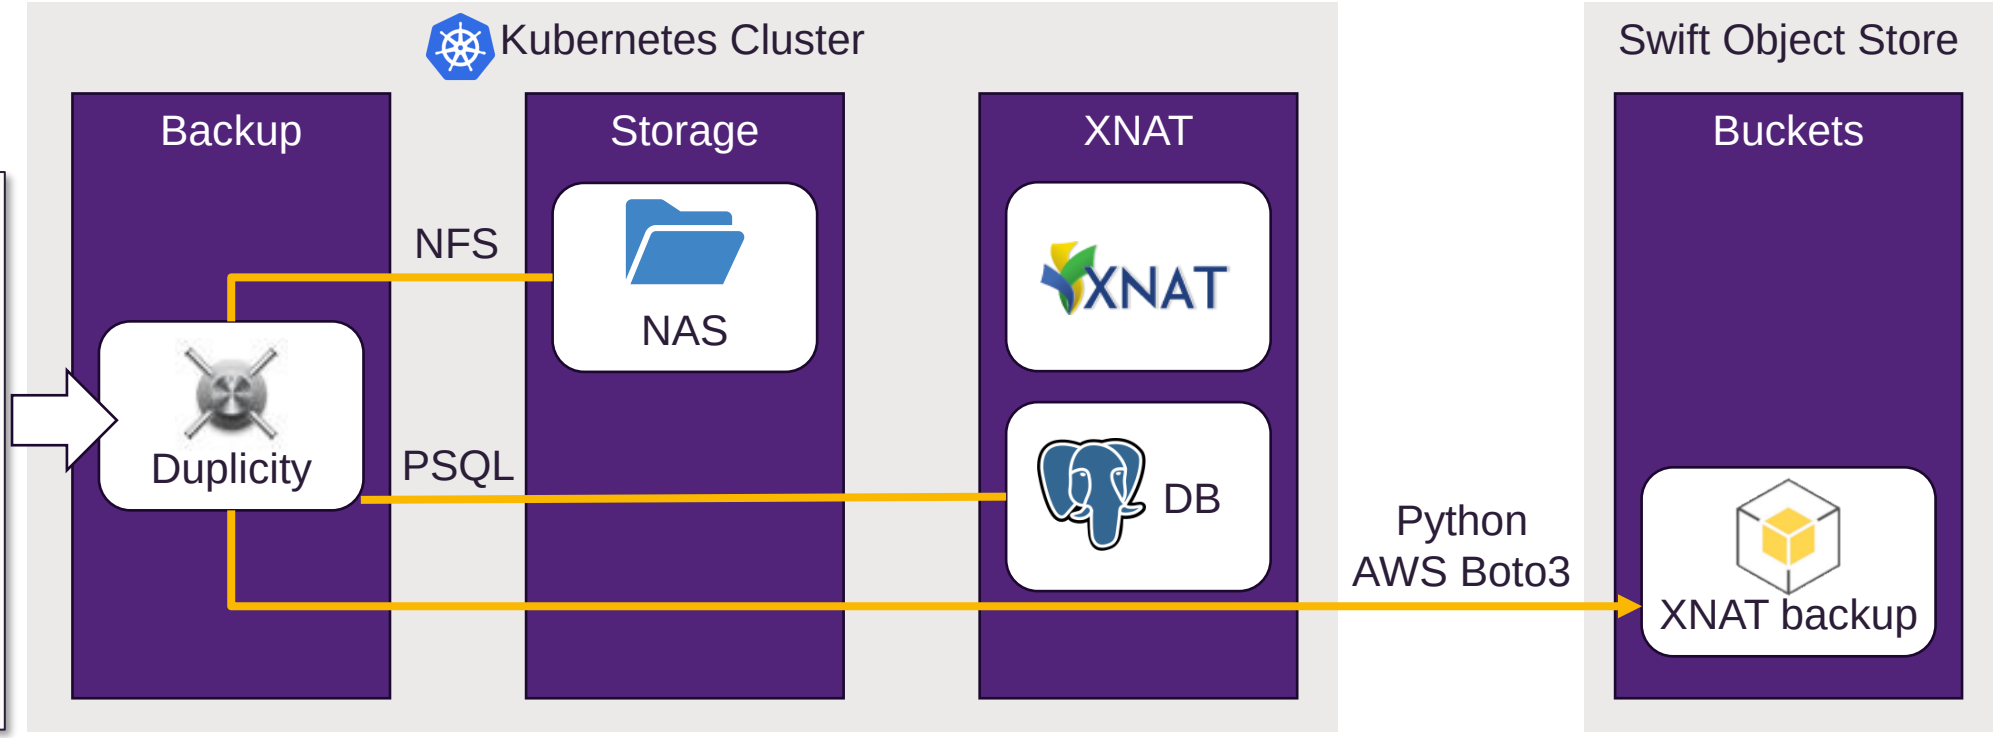

Logging

Using EFK (Elasticsearch, Fluent-bit, Kibana)

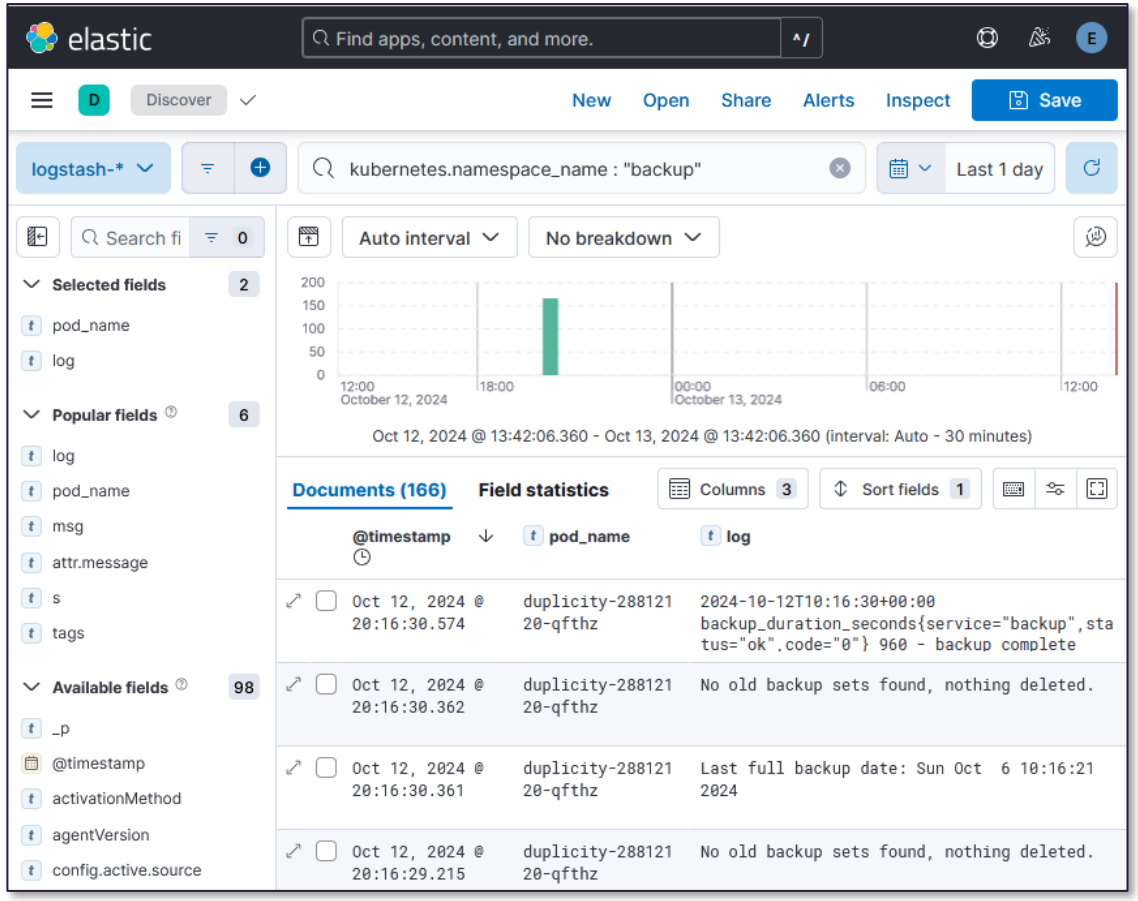

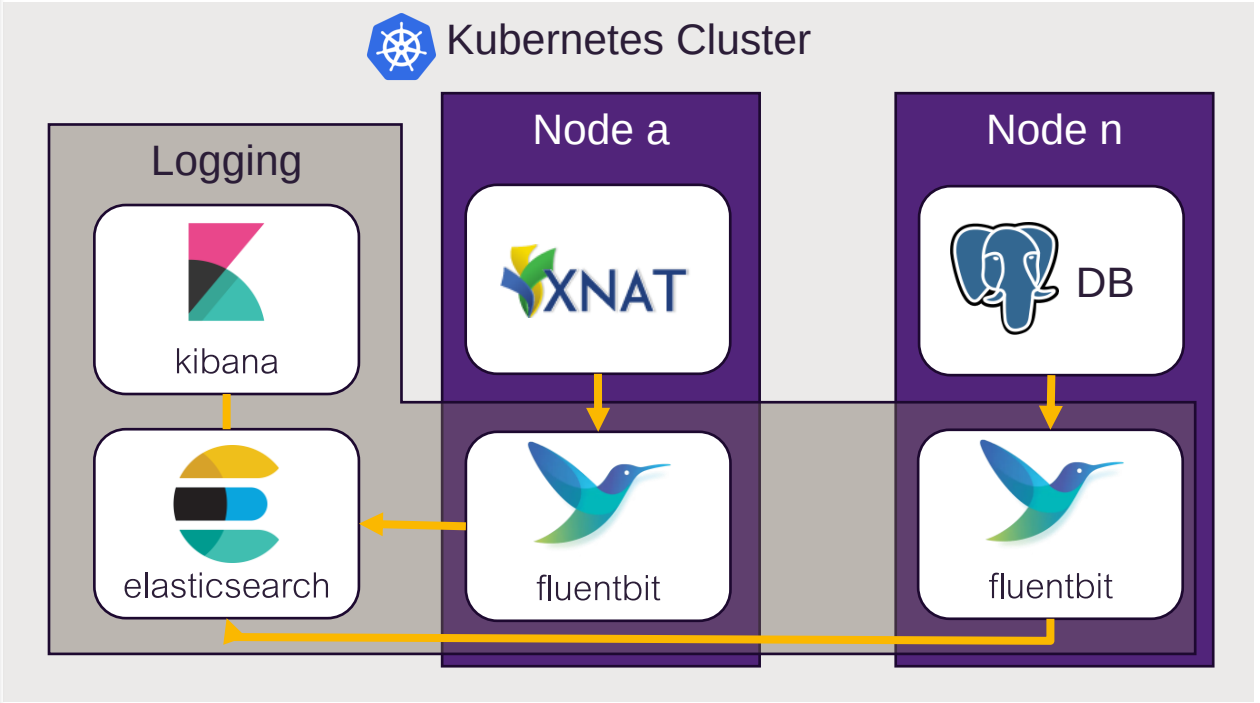

Monitoring

Using Prometheus and Grafana

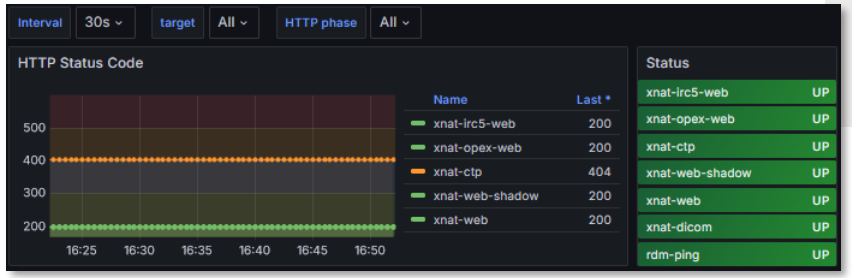

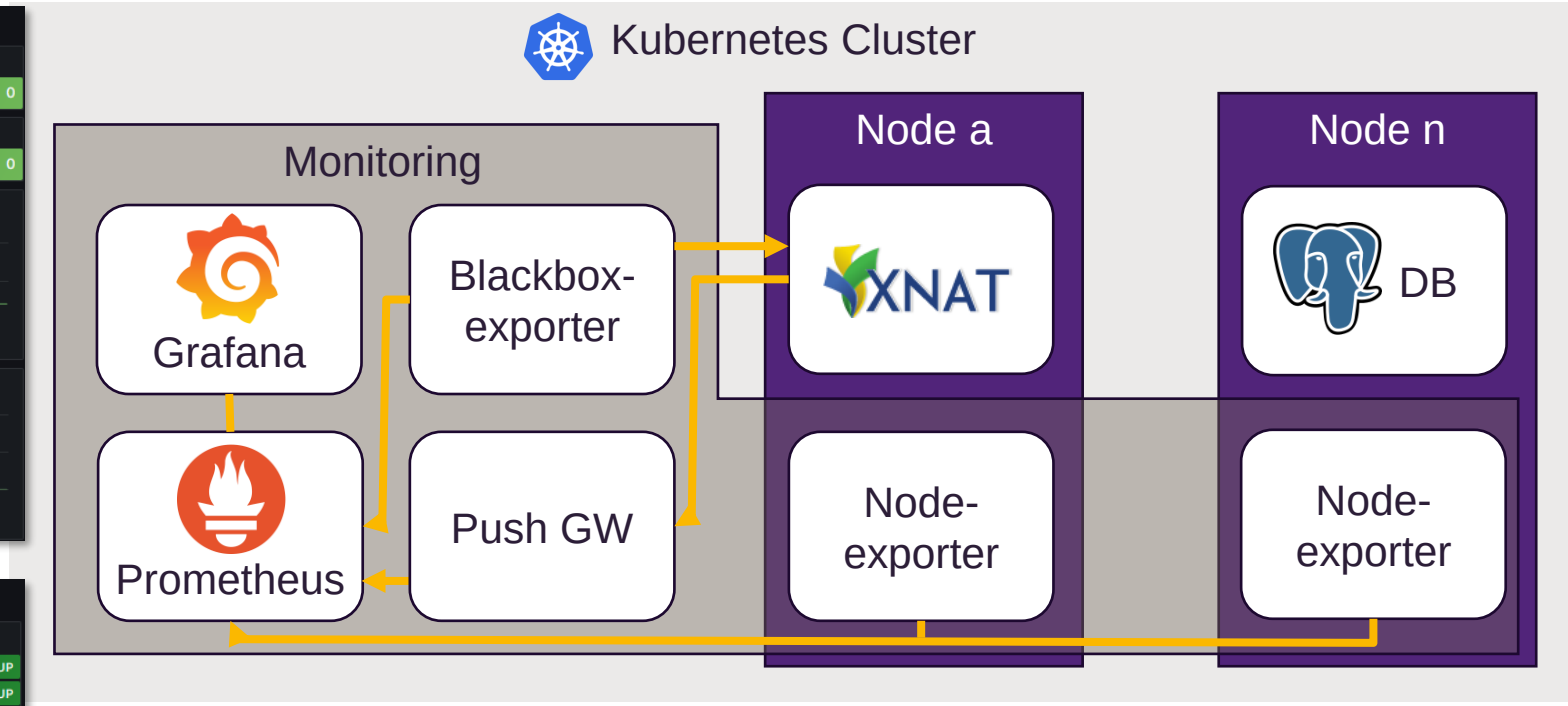

Monitoring

Using push gateway and alert manager

```
# TYPE xnat_archive_quota_used_percent gauge
xnat_archive_quota_used_percent{code="0", job="xnat_web",
  service="xnat_archive_quota", status="ok", target="A1234"} 75
```

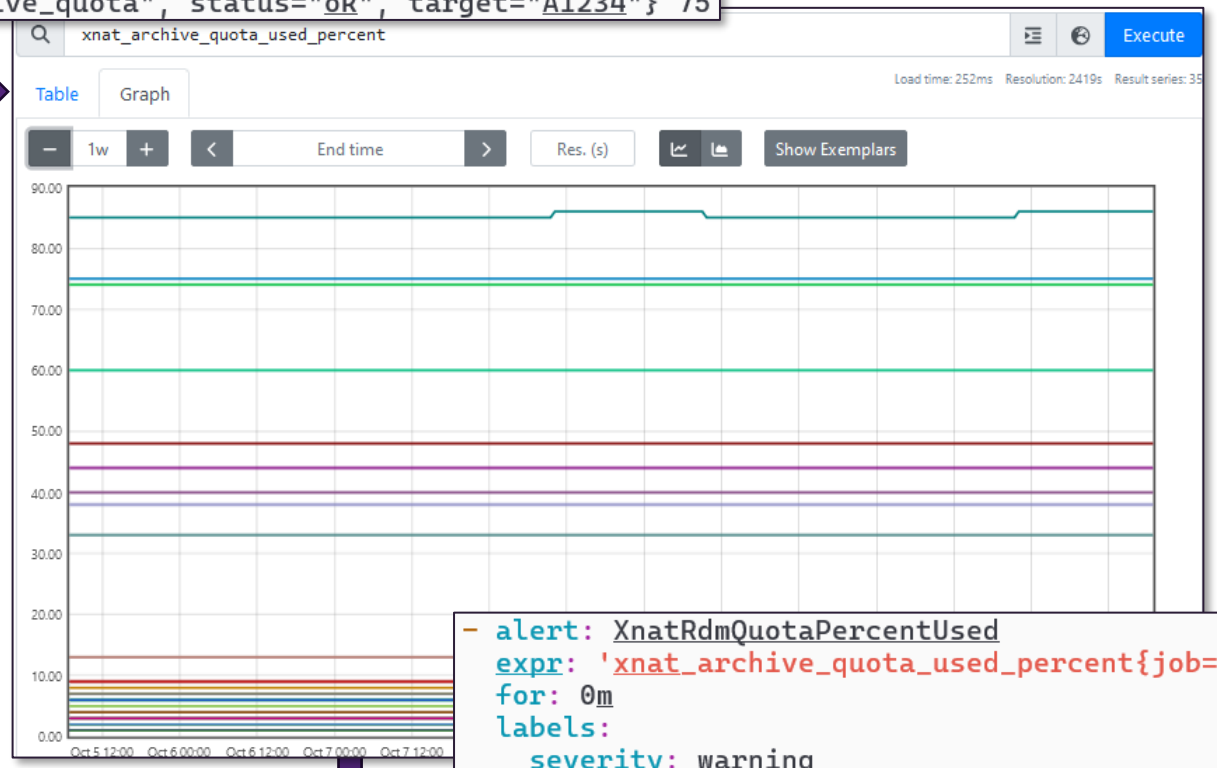


```
- alert: XnatRdmQuotaPercentUsed
  expr: 'xnat_archive_quota_used_percent{job="xnat_web"} > 85'
  for: 0m
  labels:
    severity: warning
  annotations:
    summary: 'XNAT RDM {{ $labels.target }} inode quota used over limit: {{ $value }}%'
    description: 'Advise collection owner to request more quota as per https://docs.xna
```

Secrets management

Using Bitnami SealedSecrets

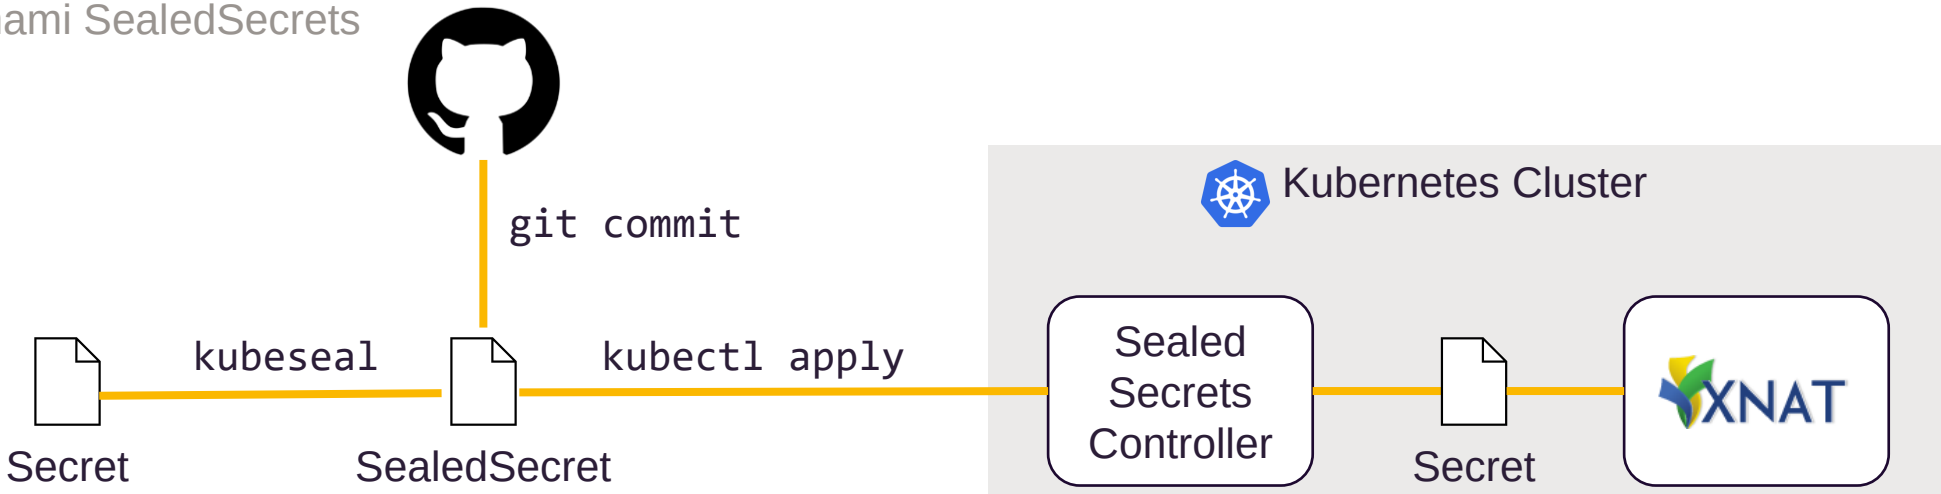

Using AWS Labs git-secrets

```
$ git secrets --add '\s*kind:\sSecret\s*'
$ git secrets -scan
xnat.yaml:373:kind: Secret

[ERROR] Matched one or more prohibited patterns

Possible mitigations:
...
- Use --no-verify if this is a one-time false positive
```

Challenges

And lessons learned

External NFS integration

Dealing with complexity

- **Start small**
- **Build on success**
- **Using library charts**

Future work

And priorities

CI/CD

Vault

Cluster autoscaling

Nectar Magnum

- **COE as a Service**
- **vGPU**
- **Network policies**

Nectar Swift

- **Object store encryption, versioning**

Thank you

UQ RCC

Mark Endrei
Nishanthi Dasanayaka
Jake Carroll

ARCOS

Jake Yip
Ping Chen