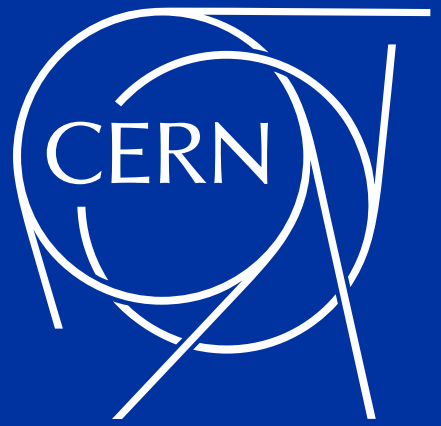

Rucio deployment

Radu Carpa

17 October 2023

About myself

- **Rucio core developer**
  - Mostly working on transfer and deletion workflows
- **Also, in charge of the ATLAS Rucio Kubernetes installation**

Ways to deploy Rucio

- Directly via `pip install`
- Containers (provided by the Rucio core team)
- Helm charts on Kubernetes (recommended way for production deployments)

ATLAS Rucio installation

- **Running in Kubernetes: 1 (small) integration + 3 production clusters (50 nodes in total)**
  - Required capacity: ~ ½ of that. The rest is for comfortable rolling re-installs of clusters
- **Configuration management via Flux2**
- **Self-managed load-balancer (haproxy 2.7.10) on puppet-managed VMs**
  - Still hoping to get rid of them

Why multiple clusters?

To increase our agility and reduce impact of risky changes:

- **Upgrading Kubernetes clusters**
- **Testing new versions of dependencies**
- **Less likely to be impacted by issues on CERN IT side (load balancer problems; removal of clusters)**

GitOps for ATLAS Rucio Operation

- **Many layers of templating engines***
  - Flux2
  - Sops (to store encrypted secrets in git)
  - Kustomize (required to use sops in flux2, but also used for rucio hot-patching)
  - Helm (managing the Rucio installation and containers)

* we are afraid Leonardo DiCaprio will have to come and save us from the Template**Inception**

kubectl kustomize via flux

Better multi-cluster handling

Store encrypted secrets in the git repository

Reduce repetition

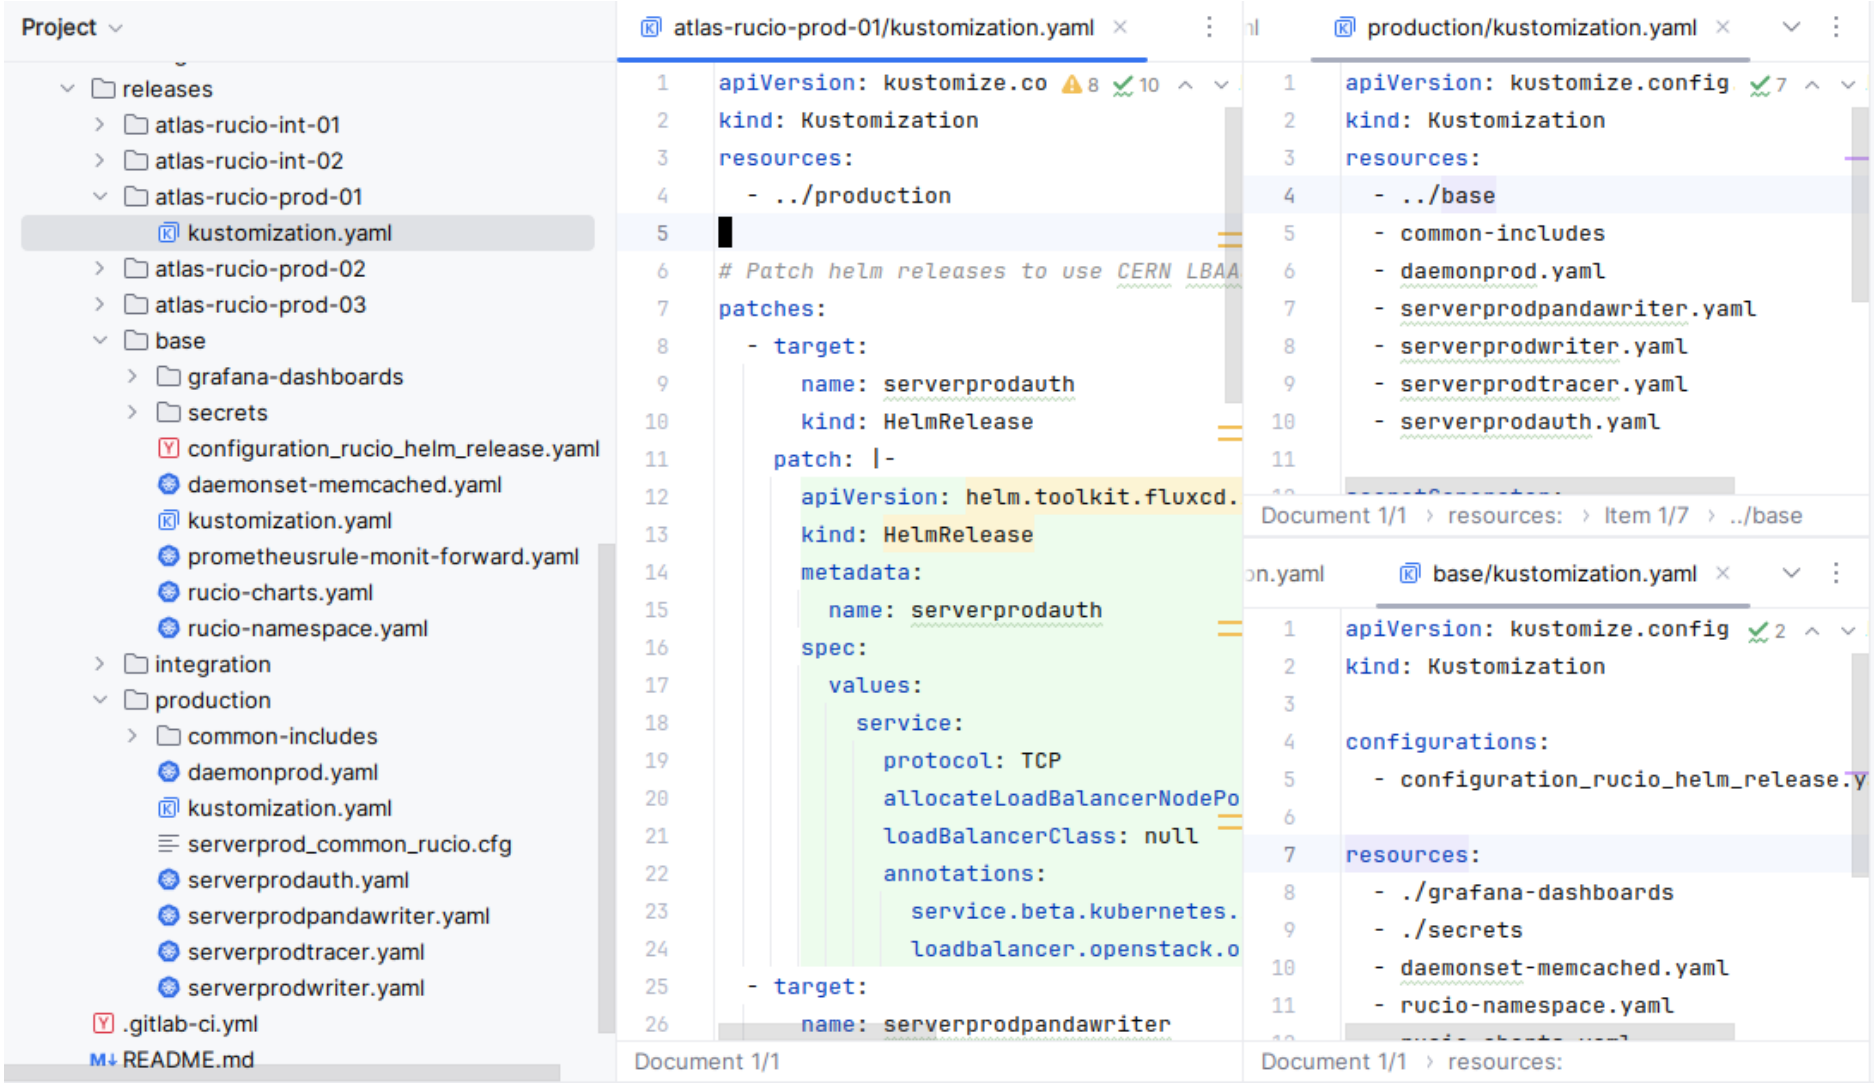

Helm via flux

Allows to manage any helm-based component (most of k8s world).

daemonprod.yaml

```yaml
apiVersion: helm.toolkit.fluxcd.io/v2be
kind: HelmRelease
metadata:
  name: daemonprod
  namespace: rucio
spec:
  releaseName: daemonprod
  interval: 5m
  chart:
    spec:
      sourceRef:
        kind: HelmRepository
        name: rucio-charts
      chart: rucio-daemons
      version: 32.0.0
  valuesFrom:
    - kind: Secret
      name: db-secret-daemons
  values:
    automatixCount: 1
    conveyorTransferSubmitterCount: 1
    conveyorPollerCount: 1

    automatix:
      sleepTime: 180
      threads: 1
```

rucio-charts.yaml

```yaml
apiVersion: source.toolkit.fluxcd.io/v1beta1
kind: HelmRepository
metadata:
  name: rucio-charts
  namespace: rucio
spec:
  interval: 10m
  url: https://rucio.github.io/helm-charts/
```

GitOps examples

1. **managing the Rucio installation**
2. **hot-patching Rucio** https://rucio.cern.ch/documentation/operator/administration/
3. **storing encrypted secrets in Git**
4. **applying a change only on one cluster**

Observability

Internal cluster metrics using kube-prometheus-stack

Cross-cluster aggregation via thanos

Alerts via alertmanager

Logs collected by filebeat and sent to CERN monit

Did you know about the following configuration option?

```
1  [common]
2  logjson = True
3
```

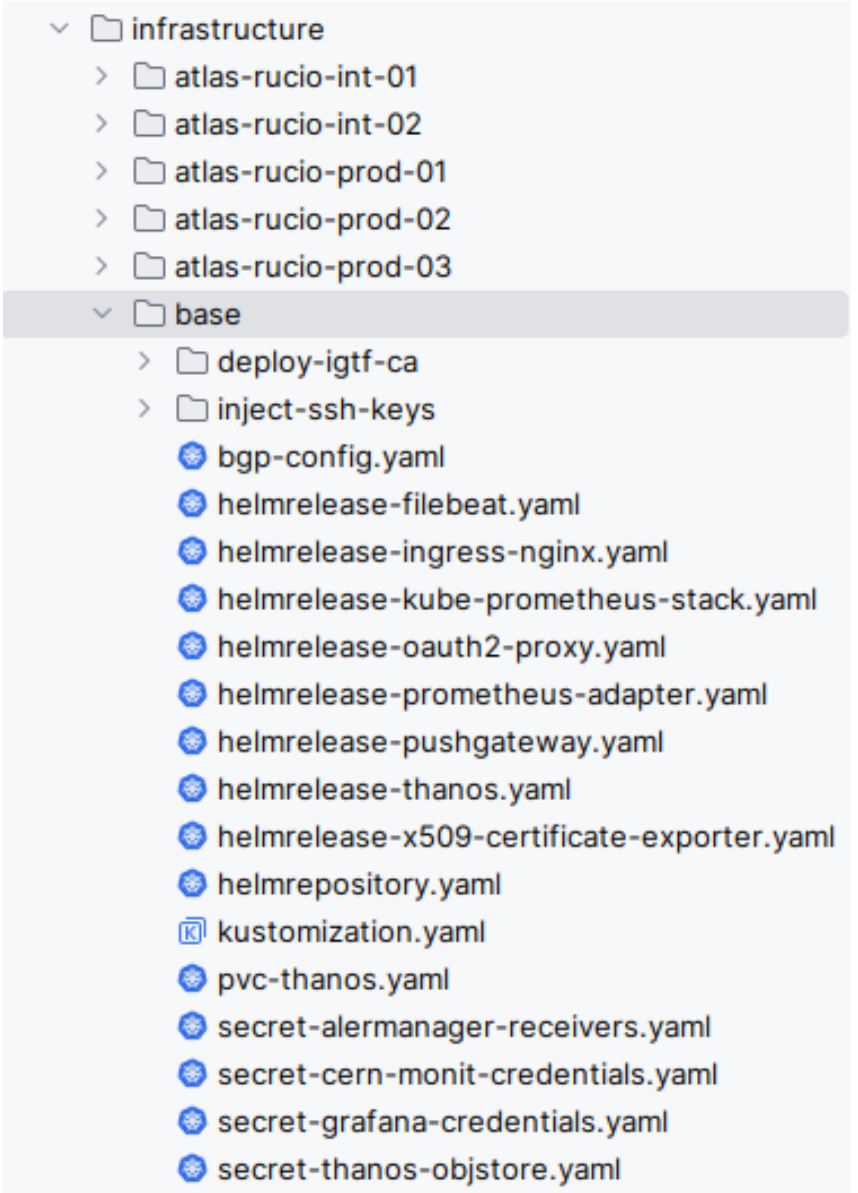

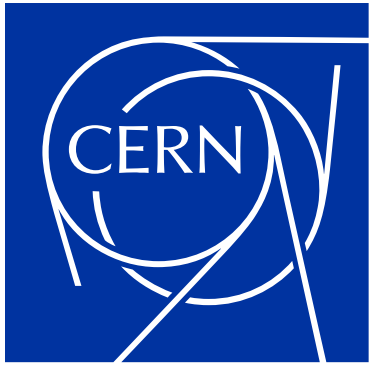

home.cern