

# RCS-SIS GitOps

Setup and use cases

By [Benjamin Bergia](#)

# RCS-SIS

*The CERN Scientific Information Service aims at efficiently managing, preserving and disseminating scientific information to make it openly accessible and reusable to CERN and the worldwide High-Energy Physics community.*


All about collaborations:

- Inside the CERN community ([CERN Analysis Preservation](#) and [CERN Academic Training](#)).
- With other institutions in the field ([InspireHEP](#) and [HEPData](#)).
- With the scholarly community as a whole ([SCOAP3](#), [SciPost](#) and [arXiv](#)).

# Requirements

- Codebases should be publicly available on Github.
- Docker images should also be publicly available.
- Some external, non-CERN, developers.
- Different release cycles.
- Production traffic 24x7.
- QA and Prod environment for each project.

# Overview



# Codebases

- Github public repository
- Currently python only
- Github actions:
  - Build docker images
  - Run test
  - Push to DockerHub or CERN Registry
  - Trigger events on the configuration repository

# Configurations

- Github private repository
- Flat YAML files + Kustomizations
- Github actions:
  - Test
  - Push to production branches
  - Update image tags
  - Call ArgoCD Webhook

# ArgoCD

- In-Cluster
- Different projects
- Pull from the Configurations repository
- Run Kustomize
- Apply result
- Auto Sync & Self Heal
- 61 Apps
- ApplicationSet

```
1  apiVersion: argoproj.io/v1alpha1
2  kind: ApplicationSet
3  metadata:
4    name: scoap3
5  spec:
6    generators:
7      - matrix:
8        generators:
9          - list:
10             elements:
11               - namespace: scoap3-qa
12                 targetRevision: master
13               - namespace: scoap3-prod
14                 targetRevision: scoap3-prod
15          - list:
16             elements:
17               - application: users
18               - application: scoap3
19  template:
20    metadata:
21      name: '{{ namespace }}-{{ application }}'
22    spec:
23      project: scoap3
24      source:
25        repoURL: https://github.com/cern-sis/kubernetes.git
26        targetRevision: '{{ targetRevision }}'
27        path: '{{ application }}/environments/{{ namespace }}'
28      destination:
29        server: https://kubernetes.default.svc
30        namespace: '{{ namespace }}'
31      syncPolicy:
32        automated:
33          prune: true
34          selfHeal: true
```

# Kustomize

- Flat YAML files
- Base resources
- Copy
- Transform
- You can layer transformations
- Doesn't enforce any file structure

Basically prototype-based (think JS)

## Pros:

- YAML all the way
- High Level (get a lot done quickly)
- Generators
- Remote resources

## Cons:

- New features
- Team responsiveness
- Arbitrary limitations (opinionated)

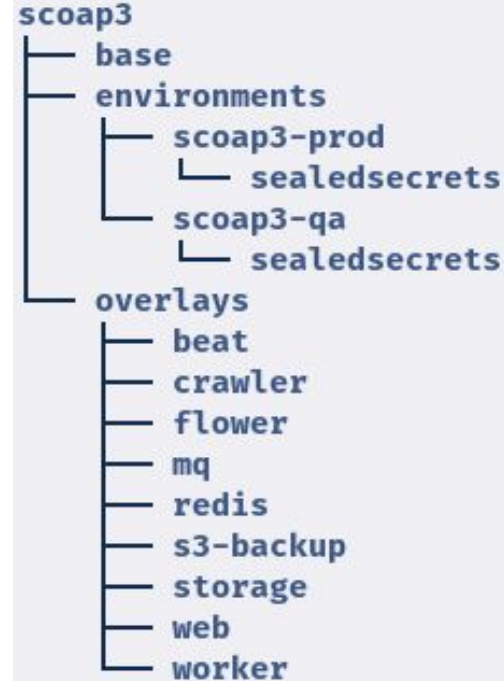


# Configurations structure


Base resources

Overlays that build on top of the base

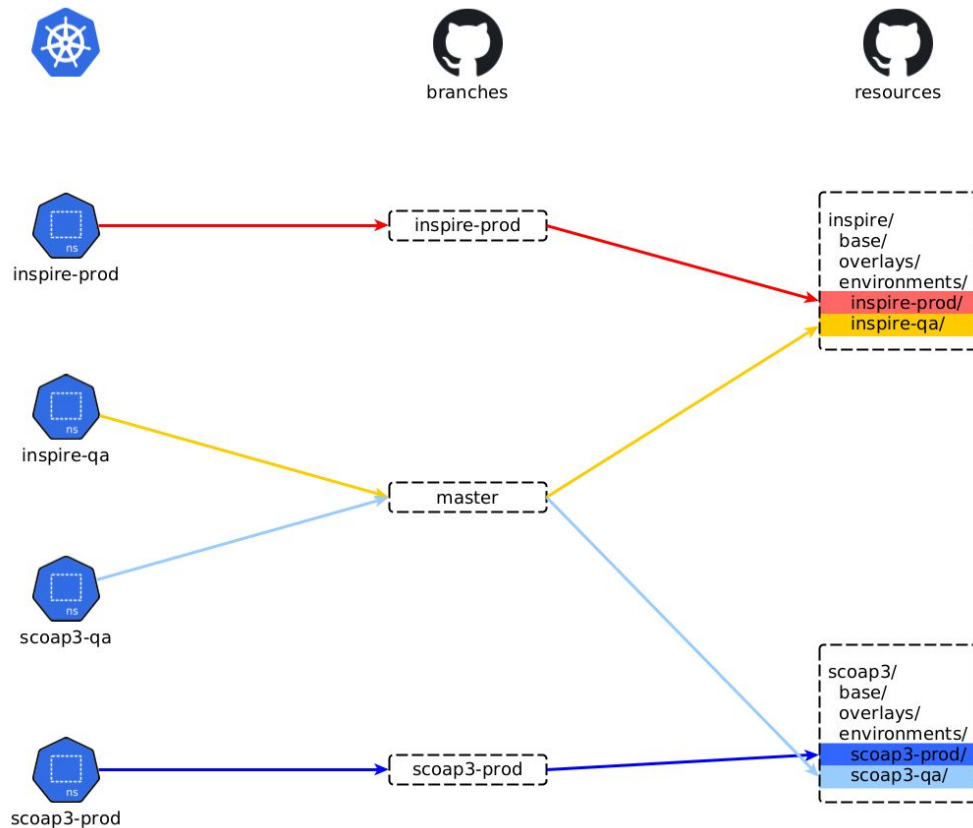
One environment for each namespace that include all the overlays needed.



# Overview



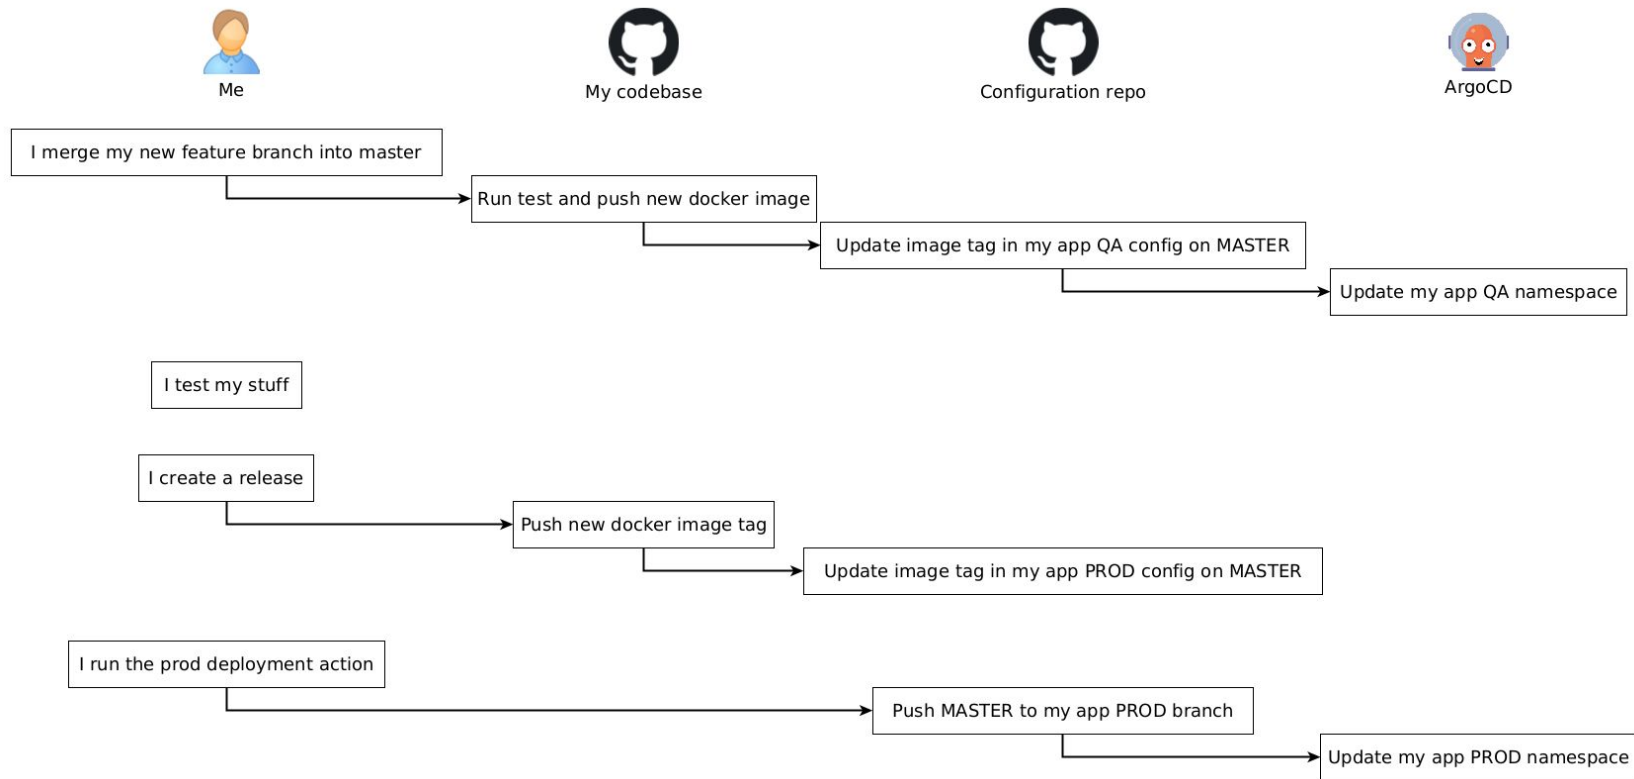
# Branching & environments



WHY?!

- Everything on master goes to all QAs
- Prod committed on master
- Each project advance its prod branch when needed

# Deployment process



# Why doing all of this

- Straightforward: what is on the repo is on the cluster.
- Simple rollbacks.
- Single source of truth.
- Releases don't impact other projects.
- The image used to run the codebase test is the one going on prod.

# Where we struggle

- Local Dev environments
- Manual interventions
- Git workflow on the Configuration repo
- Checking changes before deployment
- Testing before pushing to master

# Future improvements

- Store generated YAML on the Prod branches (Github Page like)
- Add more policies to conftest
- Branchless git workflow?

# Questions