ANZ Plus

Coming Full Circle: Immutable Clusters in the Era of Managed Kubernetes

08.10.2024

Michael Fornaro
Lead Platform Engineer

Olga Mirensky
Platform Engineer

01 Introduction

02 Kubernetes Early Days

03 Rise of Managed Kubernetes

04 Pets vs Cattle

05 Demo

Michael Fornaro

2016	***ANZ***, Joined and helped migrate to a Container
2017	**ANZ**, large scale adoption of **RedHat OpenShift**.
2019	***Raspbernetes***, OSS Kubernetes project hosted on Raspberry Pi(s)
2019	**ANZ Plus**, Predominantly working on GKE
2023	***Google Next***, Presenter on Fungible GKE clusters

Olga Mirensky

1.9 – 1.16	***Iflix***, Video streaming platform, kOps AWS on EC2 Immutable cluster upgrades for the lack of other options
1.17 – 1.21	***RedHat***, Azure Red Hat OpenShift – Develop Azure Resource Provider and customer support. In-place upgrades for customers' clusters
1.22 – 1.31	***ANZ Plus***, Predominantly working on GKE

ANZ Plus

Managed Kubernetes Evolution

Early Days

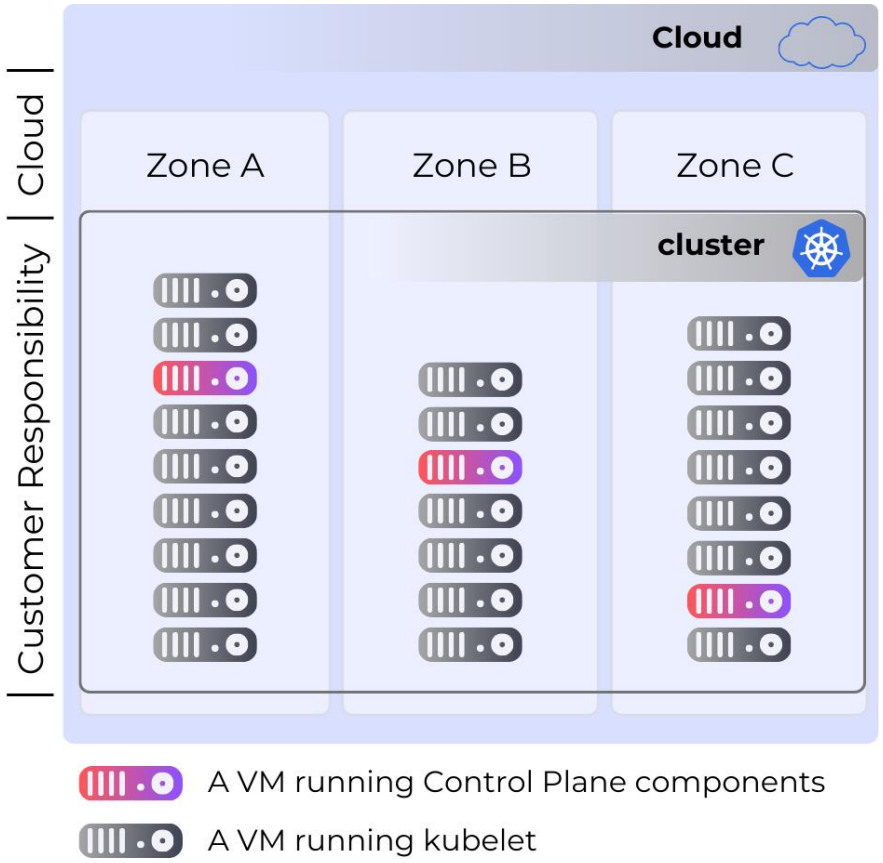

Orchestrating Blue/Green upgrade is much safer option than rotating control plane in production

1. Create new cluster
2. Smoke test
3. Deploy services (not jobs)
4. Automatic weighted DNS to shift traffic
5. Soak
6. Cut-over
7. Scale up jobs nodepools in new cluster

From DIY

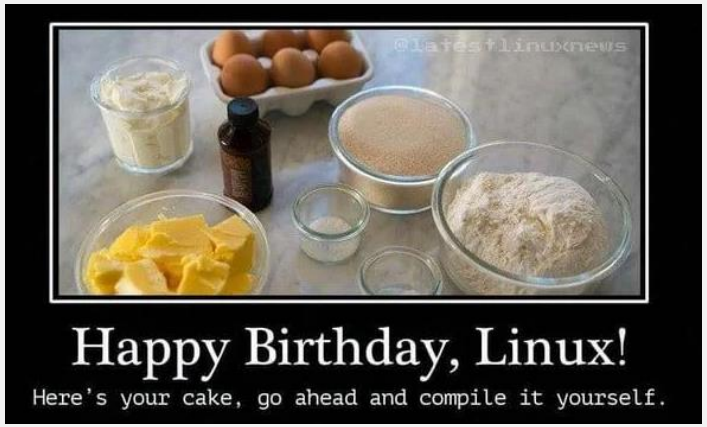

- CNI and network policies
- CSI and Persistent storage
- IAM for cloud resources
- Policy enforcement
- …

to feet-on-the-desk experience

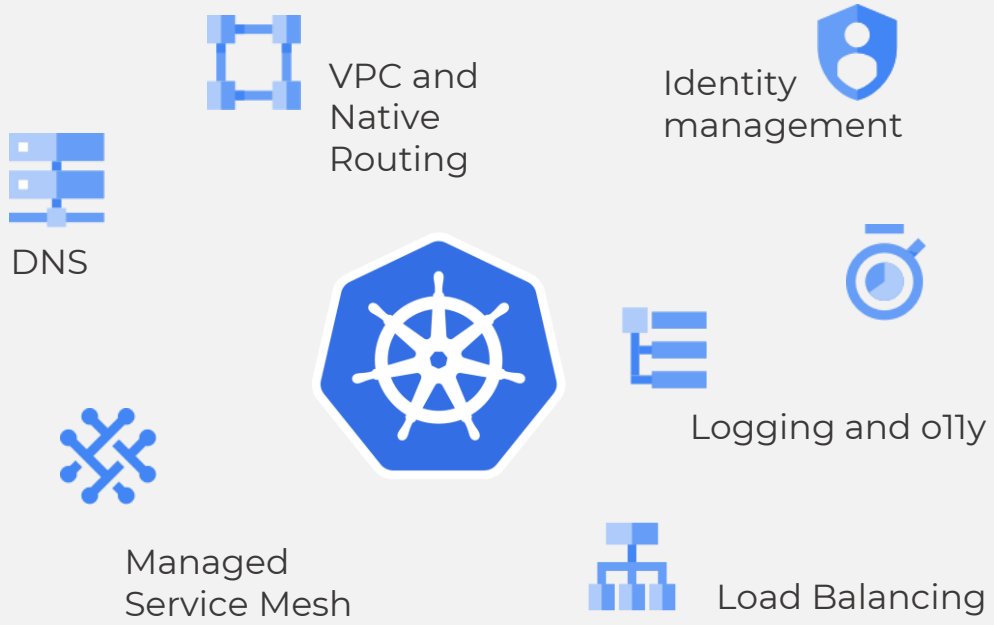

- Managed Control Plane
- Managed in-place Upgrades

Second law of thermodynamics

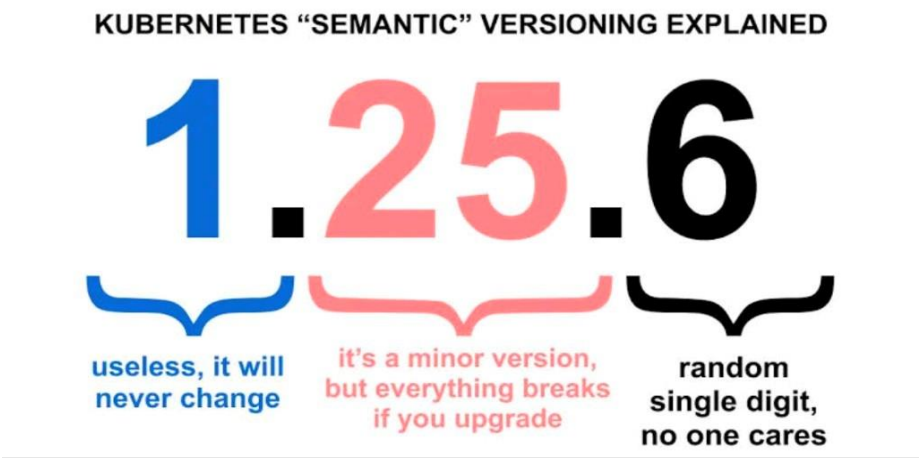

Upgrades are still a major theme for managed Kubernetes providers

- Upgrade environment promotion. Rollout sequencing
- Each cluster is a snowflake -
OS x Components x Versions x Integrations x ...
- Deprecated and removed APIs.
- Deprecated features
- Feature gates changes
- No rollback, yet.

Why rebuild clusters

Disaster Recovery

Keep process aligned with evolving infrastructure and engineers regularly practice the process

Unsupported in-place Changes

CNI upgrade, such as Dataplane v2 (GKE), Service CIDR range update (prior to v1.31), Storage solution changes

Architecture Changes

Cluster topology changes. Network architecture change such as rebuilding to a different VPC or different IP ranges

More Reasons

No support for downgrade, reducing blast radius, k8s version is too outdated, Full end-to-end test of new version

Challenges

Develop and maintain in-house tooling

Infrastructure touchpoints, e.g. IPs changes

Stateful applications

Singleton jobs and services

ANZ Plus

Demo

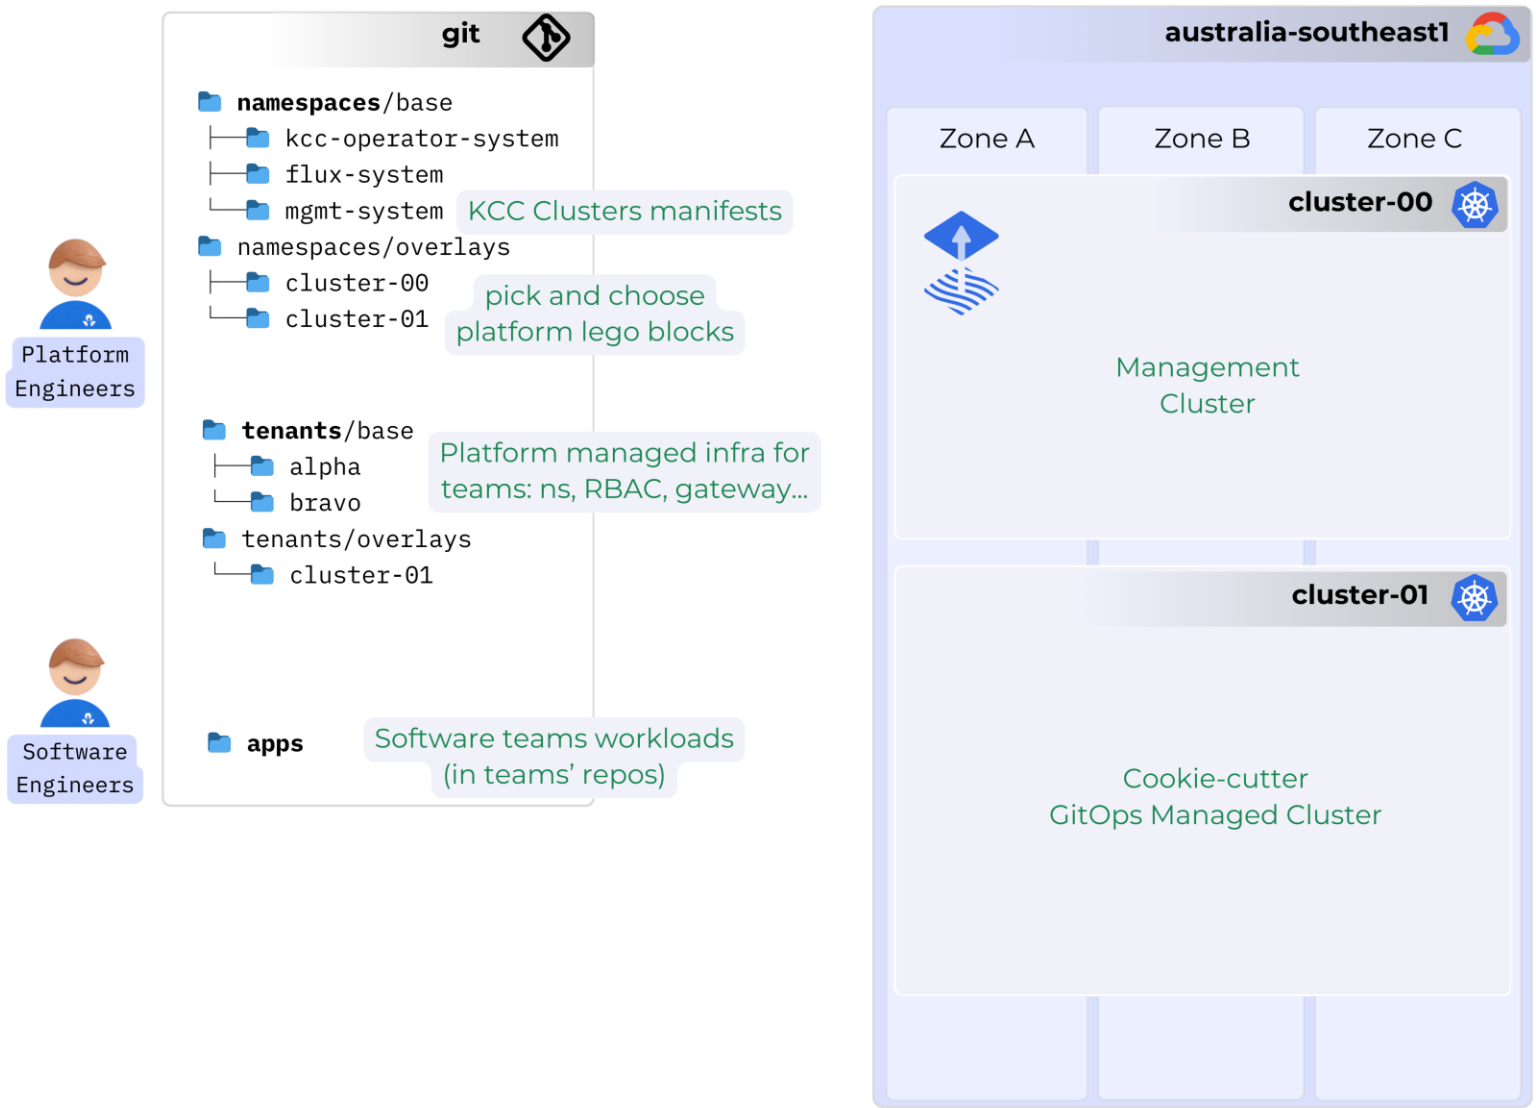
git
namespaces/base
kcc-operator-system
flux-system
mgmt-system
KCC Clusters manifests
namespaces/overlays
cluster-00
cluster-01
pick and choose
platform lego blocks
Platform
Engineers
tenants/base
alpha
bravo
Platform managed infra for
teams: ns, RBAC, gateway...
tenants/overlays
cluster-01
apps
Software teams workloads
(in teams' repos)
Software
Engineers
australia-southeast1
Zone A
Zone B
Zone C
cluster-00
Management
Cluster
cluster-01
Cookie-cutter
GitOps Managed Cluster

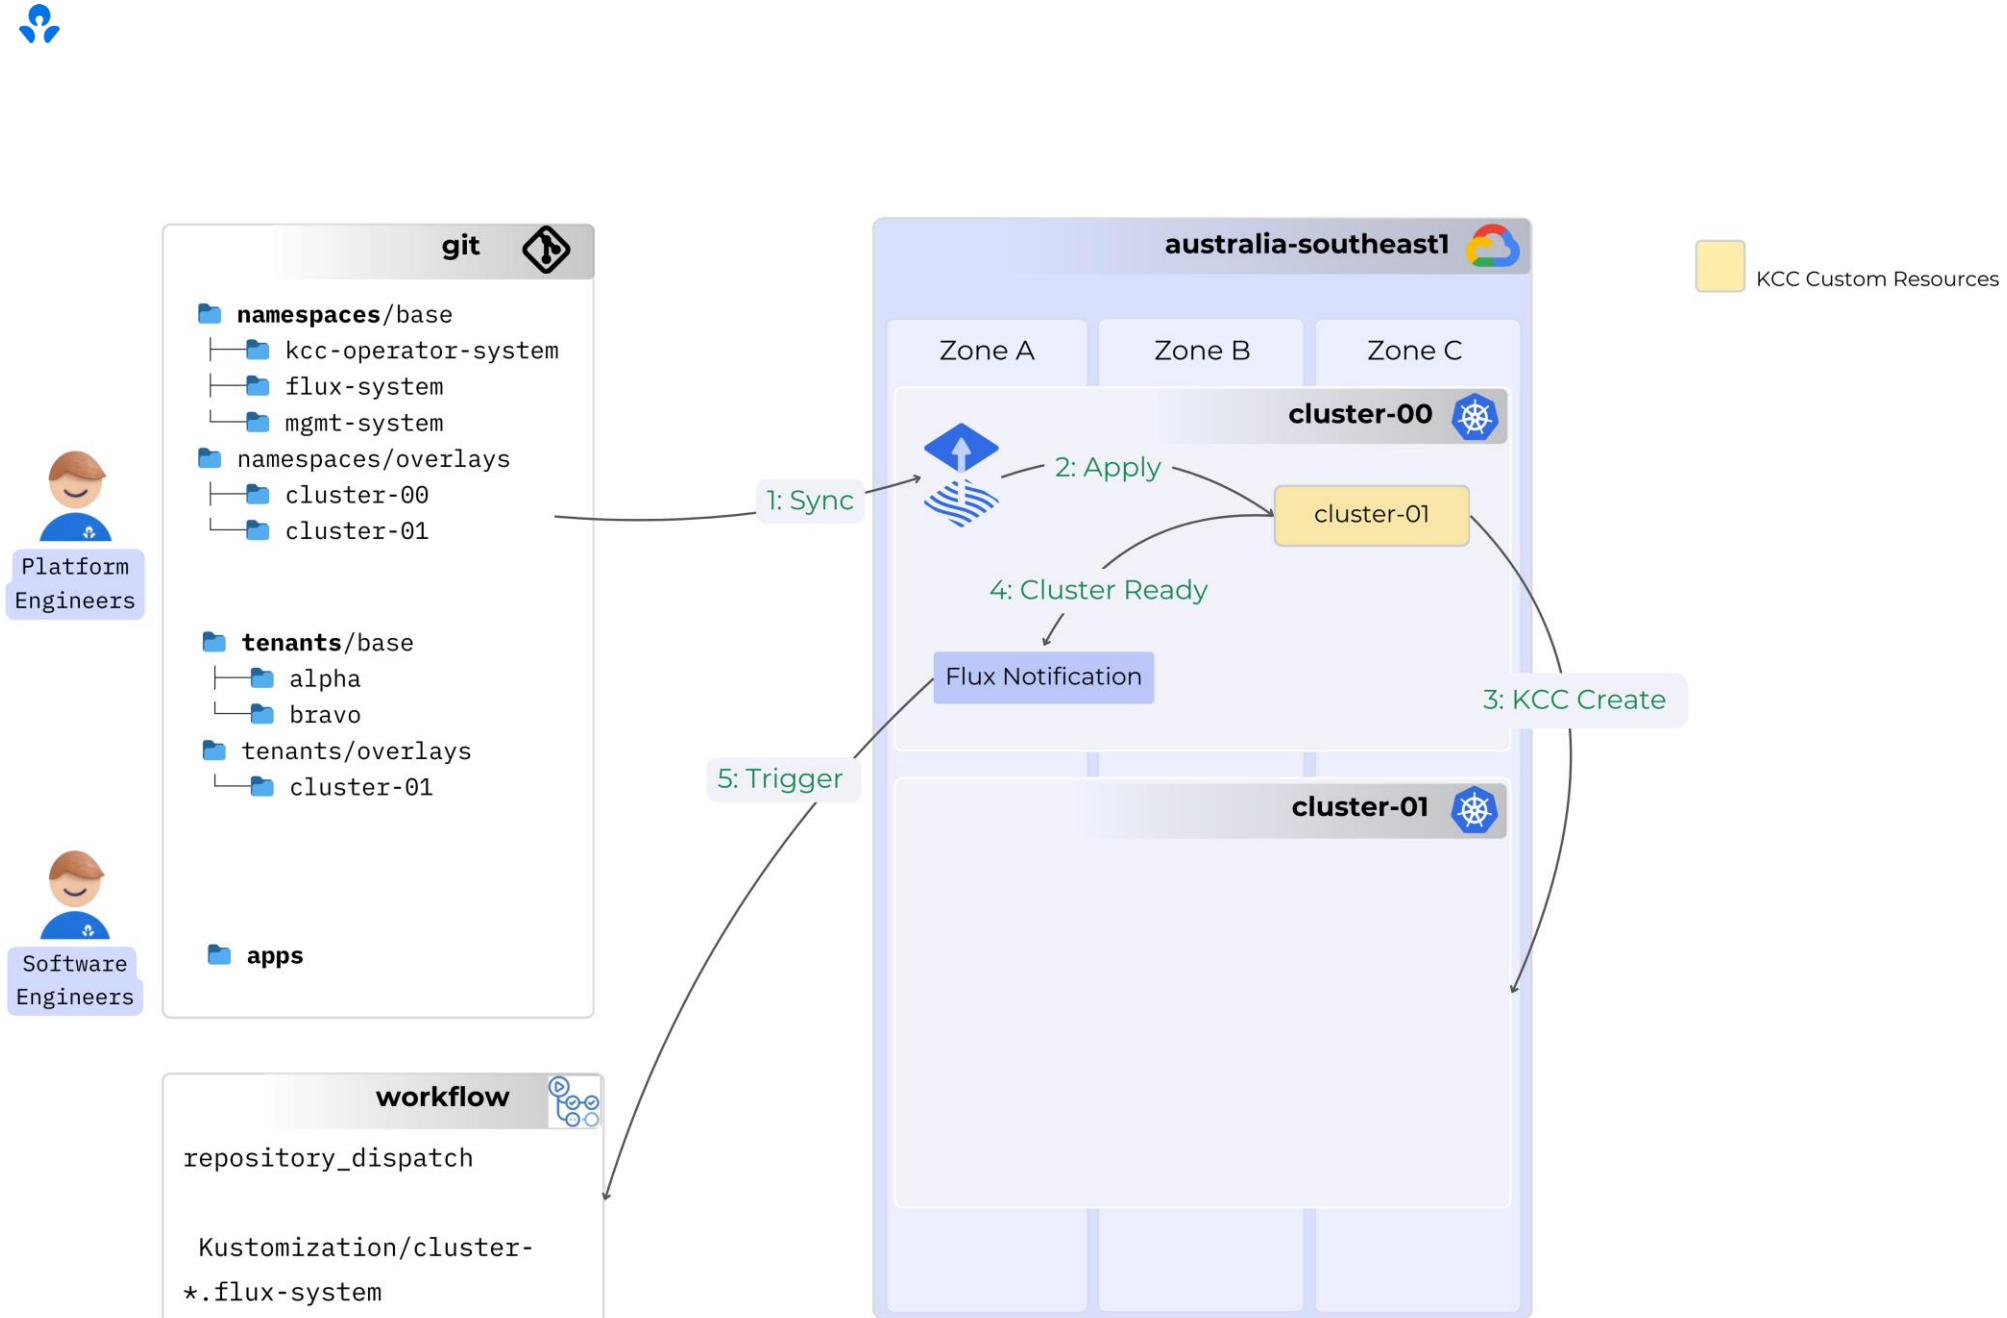
git
namespaces/base
kcc-operator-system
flux-system
mgmt-system
namespaces/overlays
cluster-00
cluster-01
tenants/base
alpha
bravo
tenants/overlays
cluster-01
apps
Platform Engineers
Software Engineers
workflow
repository_dispatch
Kustomization/cluster-*.flux-system
australia-southeast1
Zone A
Zone B
Zone C
cluster-00
cluster-01
1: Sync
2: Apply
3: KCC Create
4: Cluster Ready
5: Trigger
Flux Notification
KCC Custom Resources

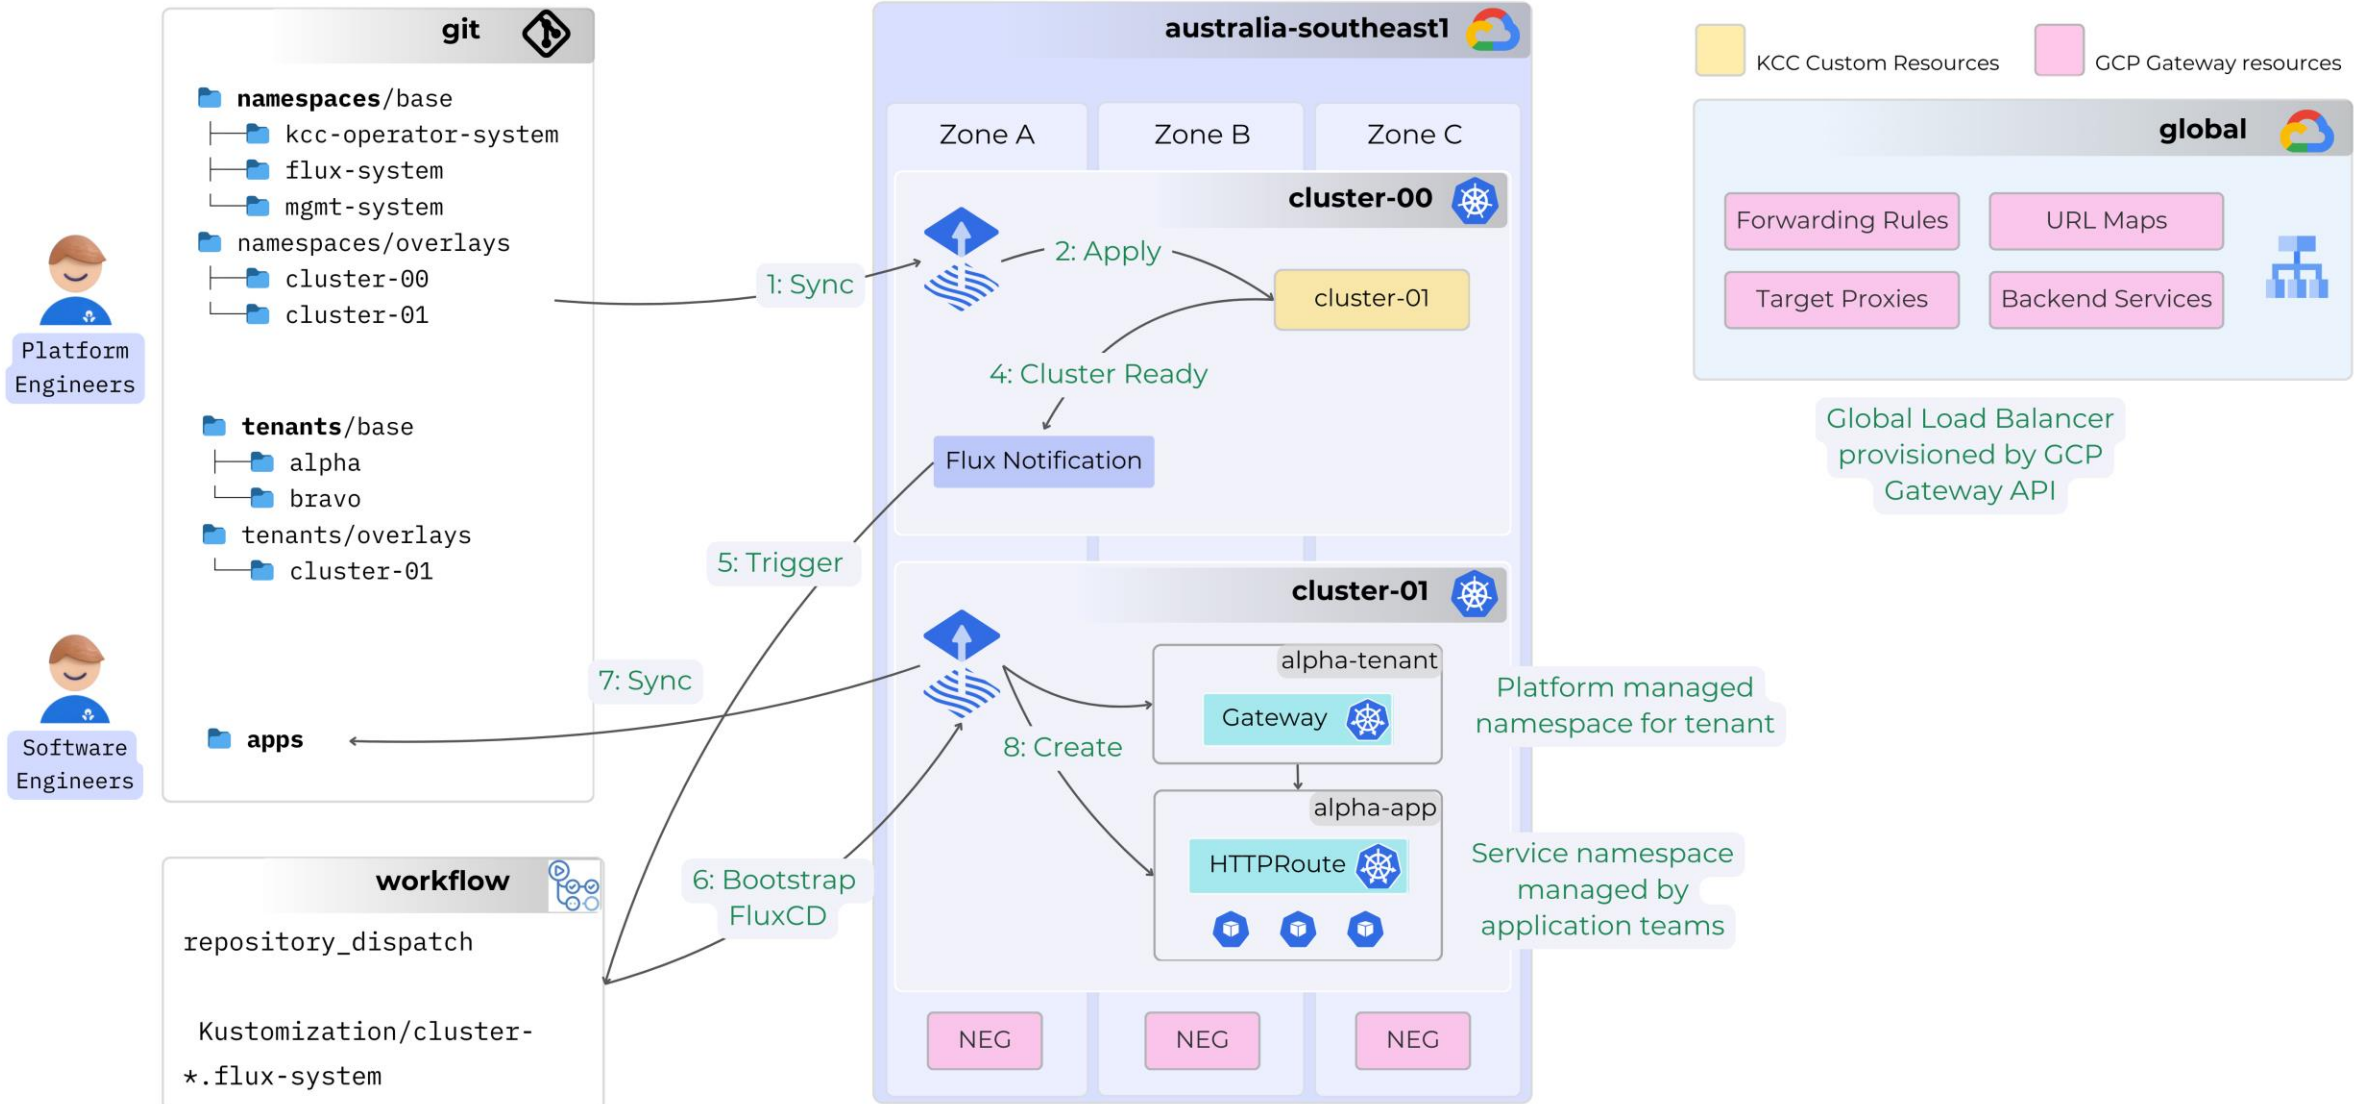
git
namespaces/base
kcc-operator-system
flux-system
mgmt-system
namespaces/overlays
cluster-00
cluster-01
tenants/base
alpha
bravo
tenants/overlays
cluster-01
apps
Platform Engineers
Software Engineers
workflow
repository_dispatch
Kustomization/cluster-*.flux-system
1: Sync
2: Apply
4: Cluster Ready
5: Trigger
6: Bootstrap FluxCD
7: Sync
8: Create
australia-southeast1
Zone A
Zone B
Zone C
cluster-00
cluster-01
Flux Notification
alpha-tenant
Gateway
alpha-app
HTTPRoute
NEG
NEG
NEG
Platform managed namespace for tenant
Service namespace managed by application teams
KCC Custom Resources
GCP Gateway resources
global
Forwarding Rules
URL Maps
Target Proxies
Backend Services
Global Load Balancer provisioned by GCP Gateway API

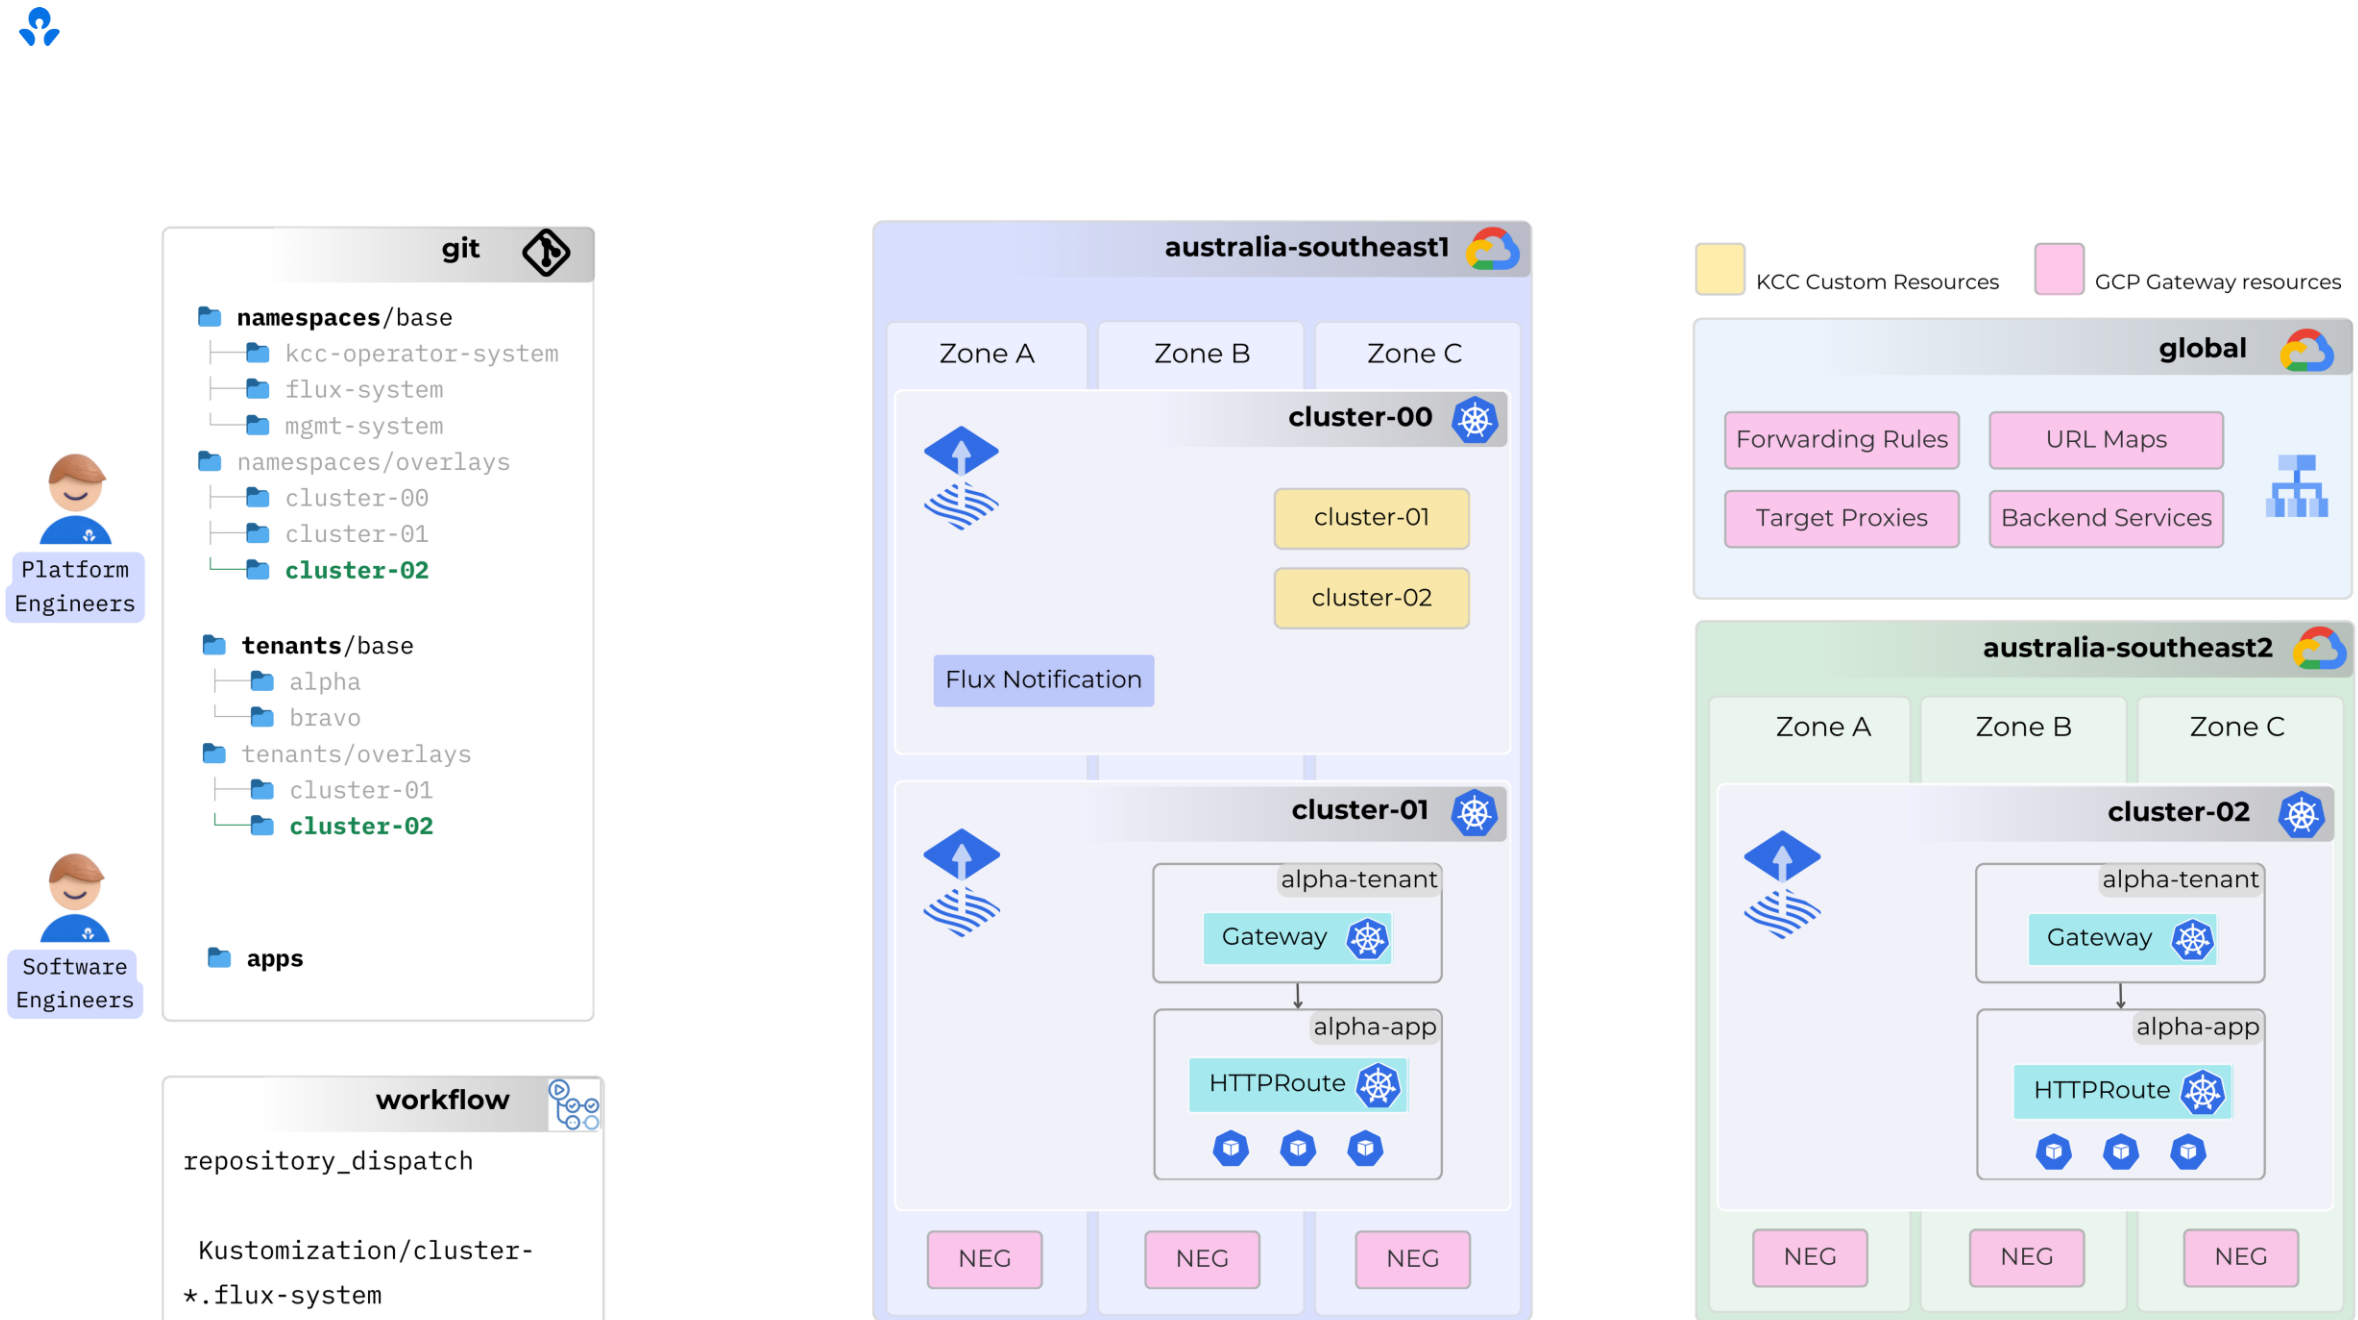
git
namespaces/base
kcc-operator-system
flux-system
mgmt-system
namespaces/overlays
cluster-00
cluster-01
cluster-02
tenants/base
alpha
bravo
tenants/overlays
cluster-01
cluster-02
apps
Platform Engineers
Software Engineers
workflow
repository_dispatch
Kustomization/cluster-*.flux-system
australia-southeast1
Zone A
Zone B
Zone C
cluster-00
cluster-01
cluster-02
Flux Notification
cluster-01
alpha-tenant
Gateway
alpha-app
HTTPRoute
NEG
NEG
NEG
KCC Custom Resources
GCP Gateway resources
global
Forwarding Rules
URL Maps
Target Proxies
Backend Services
australia-southeast2
Zone A
Zone B
Zone C
cluster-02
alpha-tenant
Gateway
alpha-app
HTTPRoute
NEG
NEG
NEG

Source code:

https://github.com/xunholy/k8s-gitops-atomic-clusters

Questions and Feedback