



Managing EKS at Scale with GitOps

Alexis Richardson, CEO Weaveworks

@monadic

alexis@weave.works

<http://weave.works>

AWS Container Days - November 2019

About Weaveworks

- Founded in 2014, backed by Google Ventures & Accel Partners
- Team of industry leaders from multiple projects
- Mission: accelerate cloud native platform adoption in the enterprise, through our **Weave Kubernetes Platform** and open source projects.



kubernetes



CLOUD NATIVE
COMPUTING
FOUNDATION

ubuntu 



spring

by Pivotal

 RabbitMQ

Accel



Genesis of GitOps

“GitOps” is our name for how we use developer tooling to drive operations....
Git is a part of every developer’s toolkit. Using the practices outlined in this post, our developers operate Kubernetes via Git. We manage and monitor all of our applications and the whole ‘cloud native stack’ using GitOps. ”

Excerpt from original blog post – GitOps Operations by Pull Request, August 2017

GitOps is born

The scene: spring 2016, a peaceful morning in London.

The sun is shining. Birds tweet.

“I’m about to make a change that will probably wipe out all our systems”

“Tom, are you sure we want to do that?”

<click>

“Oops - I’ve just deleted all our Kubernetes clusters on AWS”

CHAOS ENSUES



How Weaveworks rebuilt systems in 45 mins?

Now 5
mins or
less

- We use declarative infrastructure ie. Kubernetes, Docker, Terraform, & more
- Our **entire system** including code, config, monitoring rules, dashboards, is described in GitHub with full audit trail
- *We can roll out major or minor changes as **atomic** pull requests, and **automatically converge** then check for diffs if system diverges from the “source of truth” in Git*



Flux joins CNCF

“[you] made the single source of truth possible in git, and it was so much more clear”

[Kyle Rockman, Under Armour](#)

Argo & Flux merging in CNCF

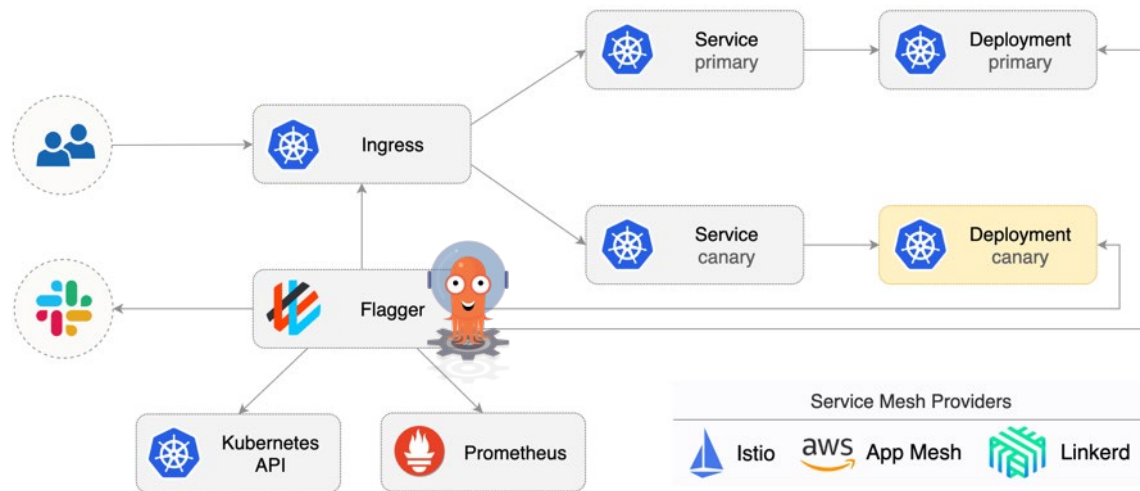
- What does this merger mean?
- Ability to standardize on GitOps
- First joint project - Argo Flux
GitOps Engine



Reproducible Automated Deployment — Policy & Compliance — Progressive Delivery — Alert & Fix Failures Automatically

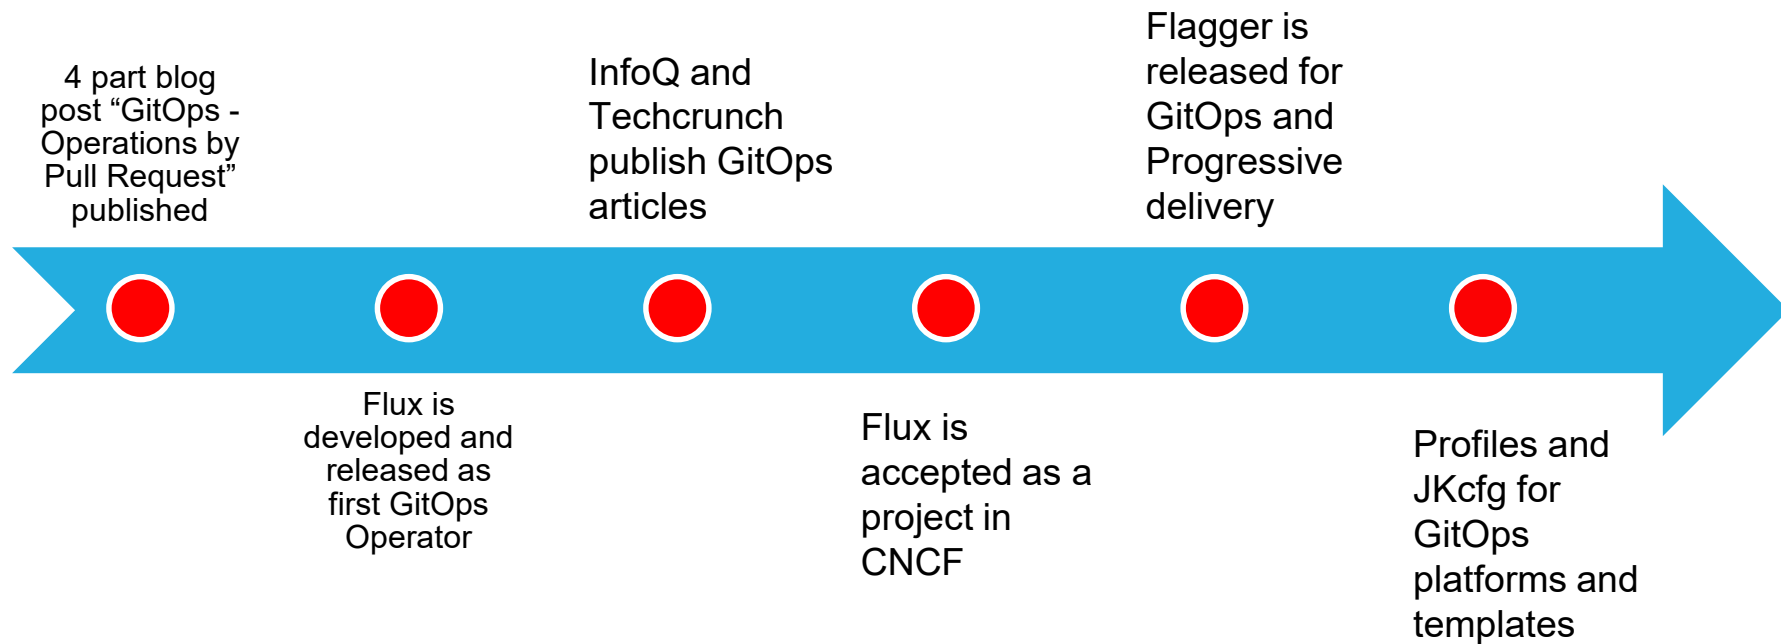
Future

- Extending GitOps to progressive delivery
- Canary, Blue/Green and other controlled deployments
- Experiments under policy



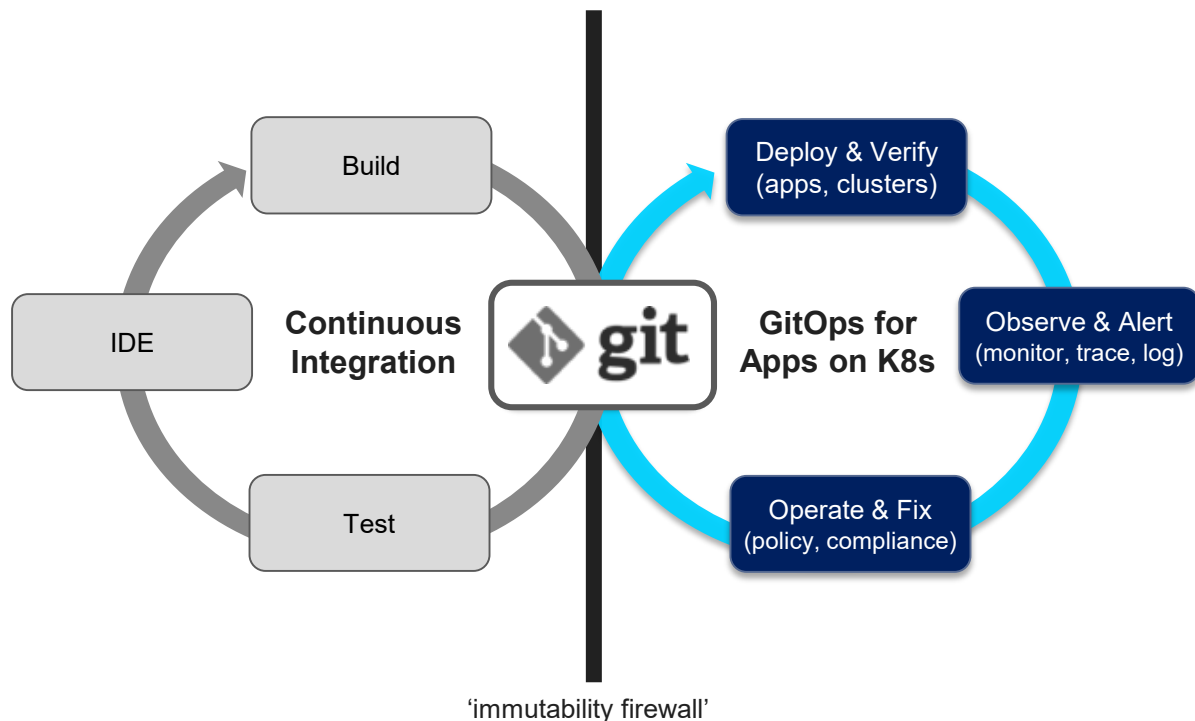
Reproducible Automated Deployment — Policy & Compliance — Progressive Delivery — Alert & Fix Failures Automatically

GitOps @ Weaveworks Timeline



Reproducible Automated Deployment — Policy & Compliance — Progressive Delivery — Alert & Fix Failures Automatically

Can GitOps be a “complete solution”?

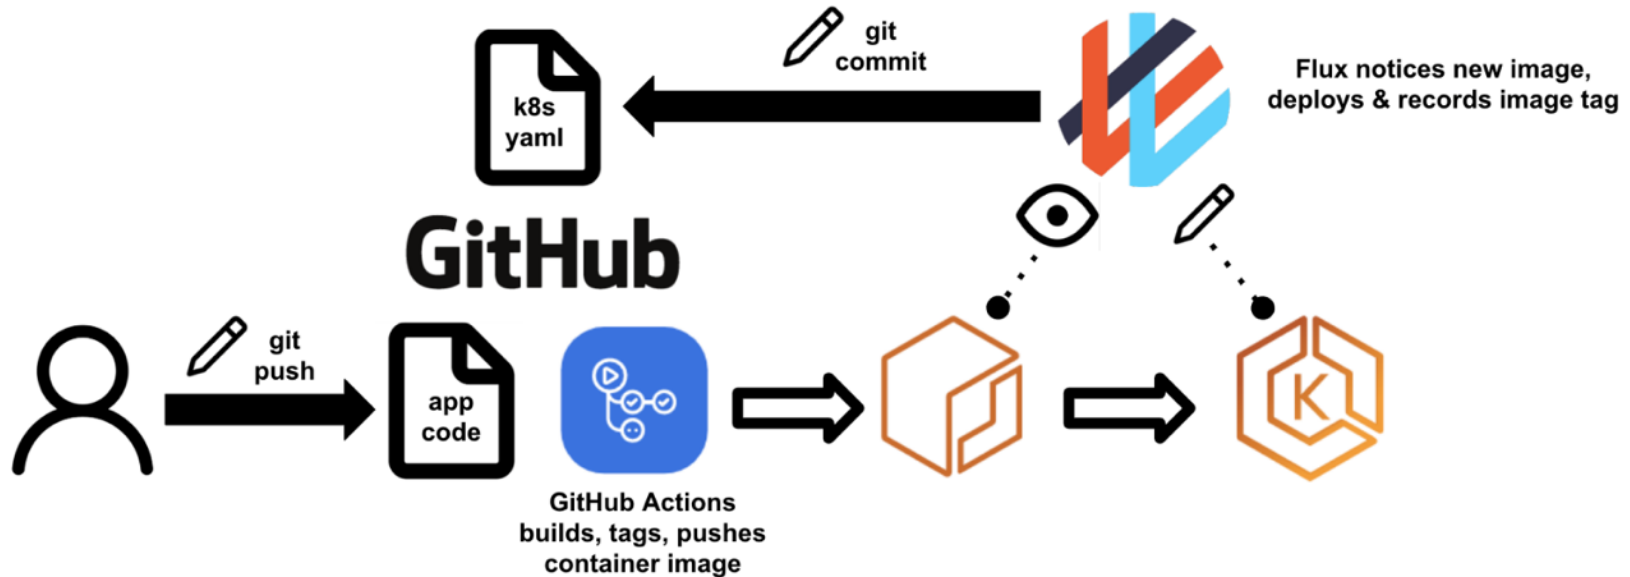


Git as the single source of truth of a system's desired state, enabling reproducible automated deployment, management and drift alerts

GitOps Diffs compare desired state with observed state continuously

... What about **CLUSTERS**

GitOps gives us a developer workflow – what can we do with it



Give me the stack I want

Booting Kubernetes clusters is easy.
Managing a complete application platform is hard.

“Give me an ML stack”

- Make it safe, consistent, auditable, upgradeable
- Can you answer these three questions about your Stack?
 - Do you know if it is in the correct state?
 - Can you reproduce your entire app + cluster stack at will?
 - Can you cleanly upgrade the whole stack?



The Cloud-Native App Delivery Problem

Customers need to deliver **MANY APPs** for **MANY USE CASES**.

Machine Learning is just one of them.

- Figure out what an **'app'** means
- Figure out what 'platform add-ons' you need for your use case, e.g. machine learning
- **Deploy apps and add-ons correctly** and properly wired up to other services, e.g RDS, Lambda
- **Verify** cluster, test add-ons and apps are configured correctly for dev, staging to prod...
- **Operate**, shut down, reproduce stack (think "cattle" not "pets")
- **Manage**, maintain, upgrade, patch any part of this according to POLICY
- **Scale** to fleets of clusters, across many use cases and cloud options

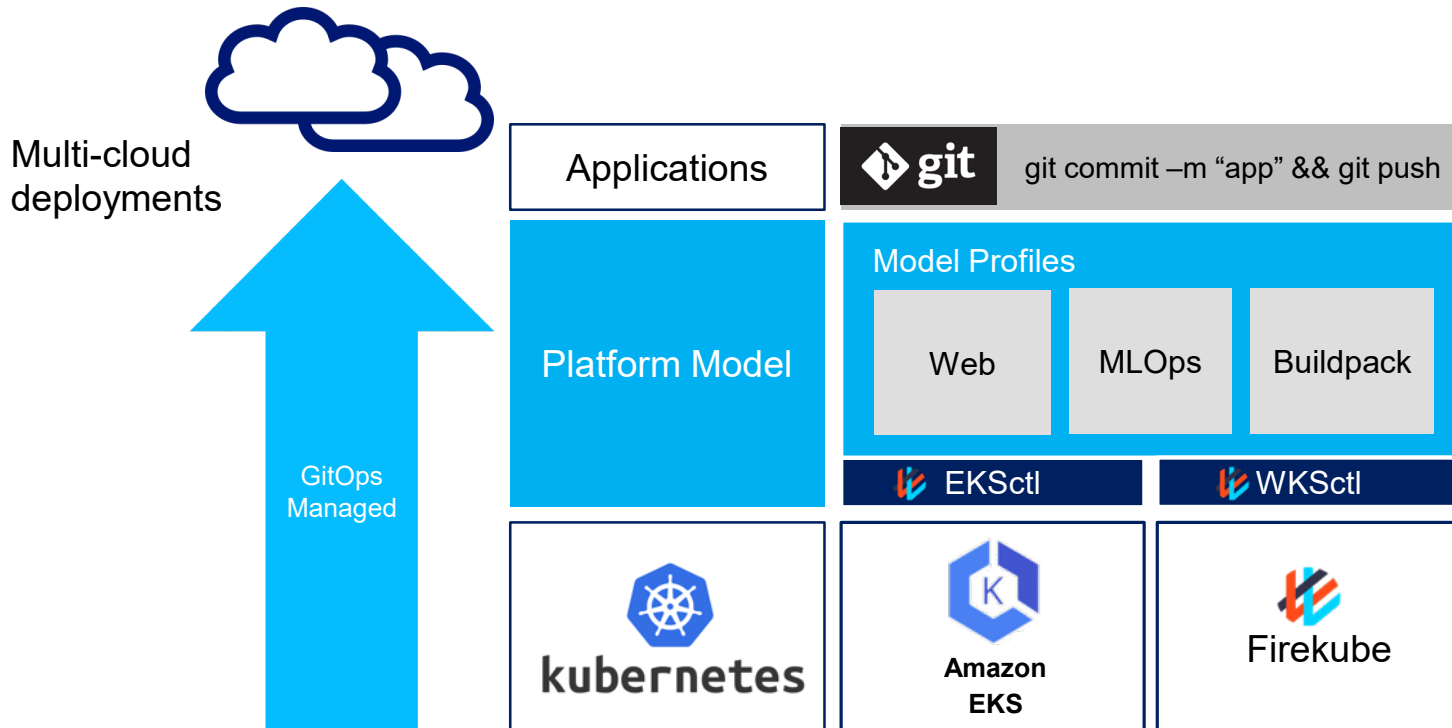
We want easily reproducible cattle and not pets or (worse) snowflakes

Reproducible clusters

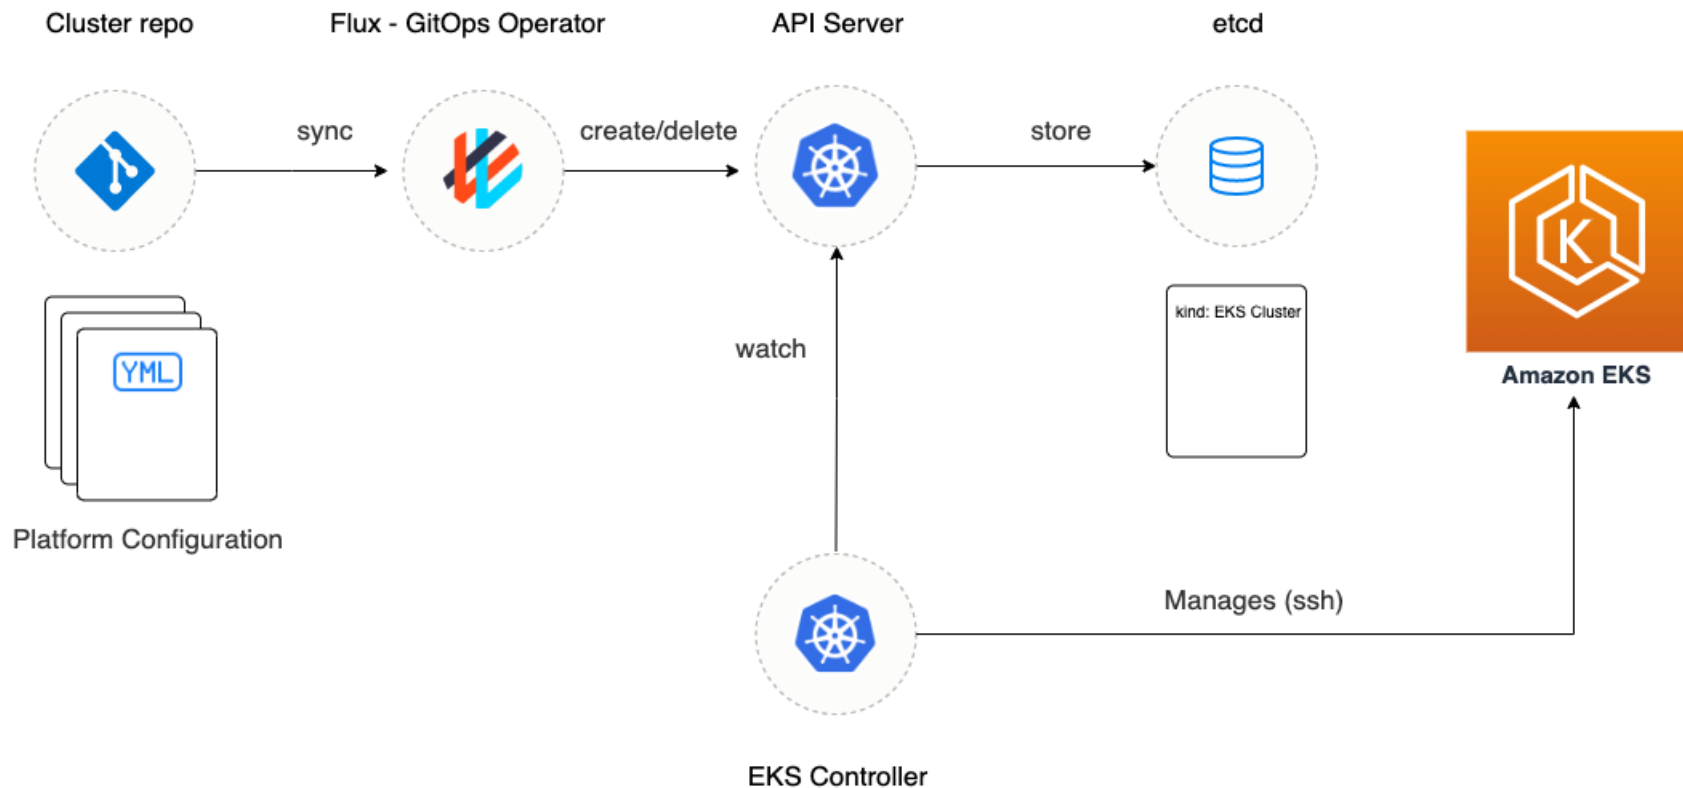
- Consistent Kubernetes App platforms with Argo Flux CD
 - A single git repository can contain:
 - Cluster definition (eg. K8s version, IP range for pods, CRI, CNI, CSI)
 - List of nodes
 - Core apps used by your team eg. Machine Learning needs MLOps
 - Additional custom apps specific to your team (eg. Authentication, CI/CD, Monitoring, etc.)
 - Cluster creation can be driven completely via GitOps with pull requests



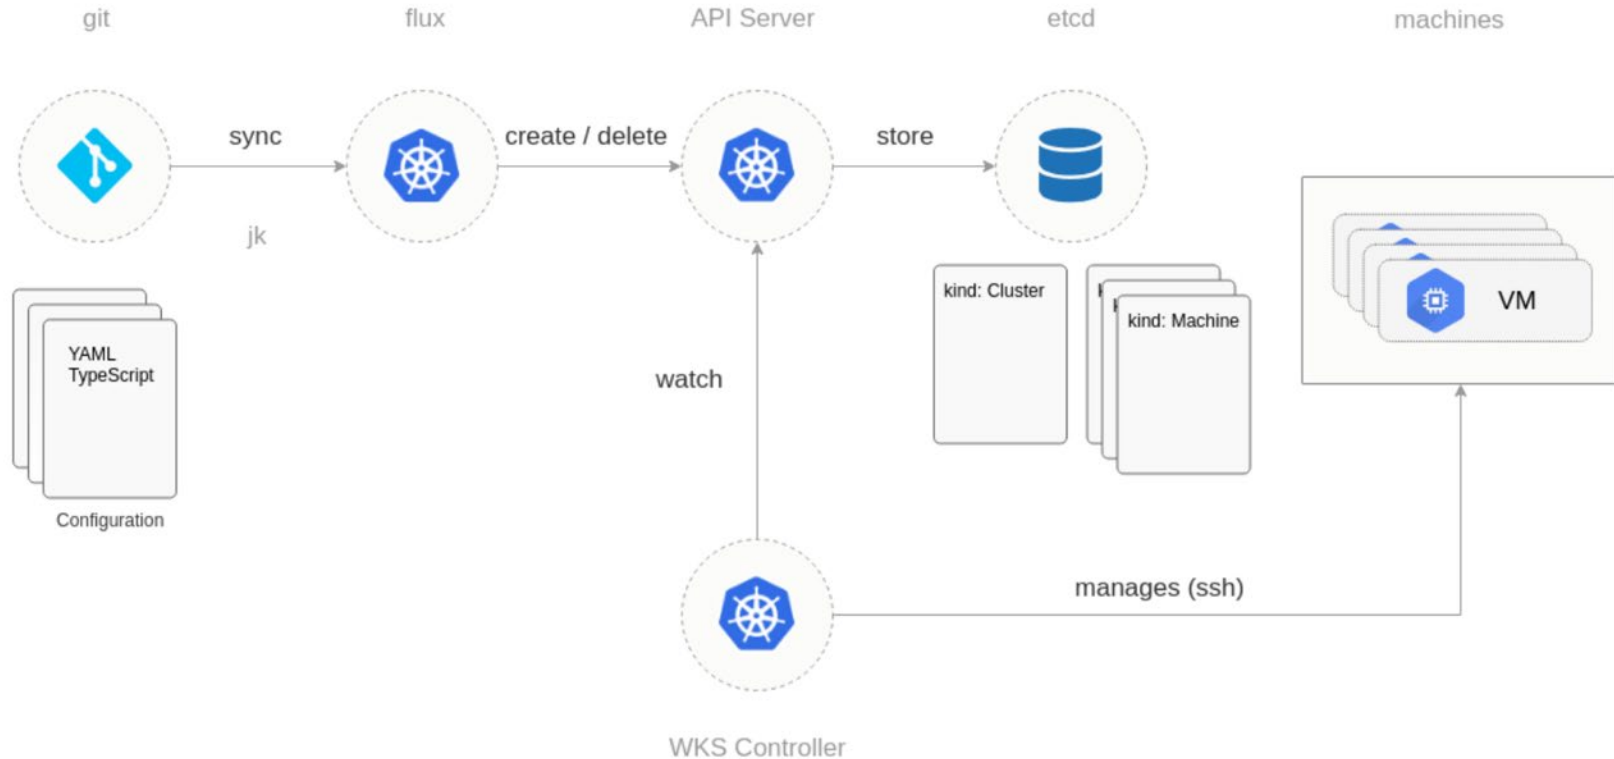
Build any stack based on a MODEL



Reproducible platforms on EKS



Reproducible platforms and OSS K8s



Declarative Clusters

- Example: Advanced cluster architecture with two AWS auto-scaling groups (nodegroups)

```
apiVersion: eksctl.io/v1alpha5
kind: ClusterConfig
metadata:
  name: jpmc-demo-cluster-1
  region: eu-west-1
  version: "1.13"
nodeGroups:
  - name: applications
    labels: {applications-only: "true"}
    minSize: 5
    maxSize: 15
    instanceType: m5.xlarge
```

```
- name: management
  taint: {management-only: "true:NoSchedule"}
  labels: {management-only: "true"}
  iam:
    attachPolicyARNs:
      - arn:aws:iam::123:policy/ng-management-1
    withAddonPolicies:
      albIngress: true
      externalDNS: true
  desiredCapacity: 2
  instanceType: m5.large
  securityGroups:
    attachIDs: [sg-0b44c48bcba5b7362]
```

Consistent toolchains across environments

- EKSctl profiles for consistent and complete clusters
- Takes a cluster, adds Flux to the cluster and links it to a git repository
- Adds the “app-dev” profile config to your repository
- Flux watches the repository and deploys the workloads



eksctl enable profile



cluster

--git-url



gitops repo
profile manifests

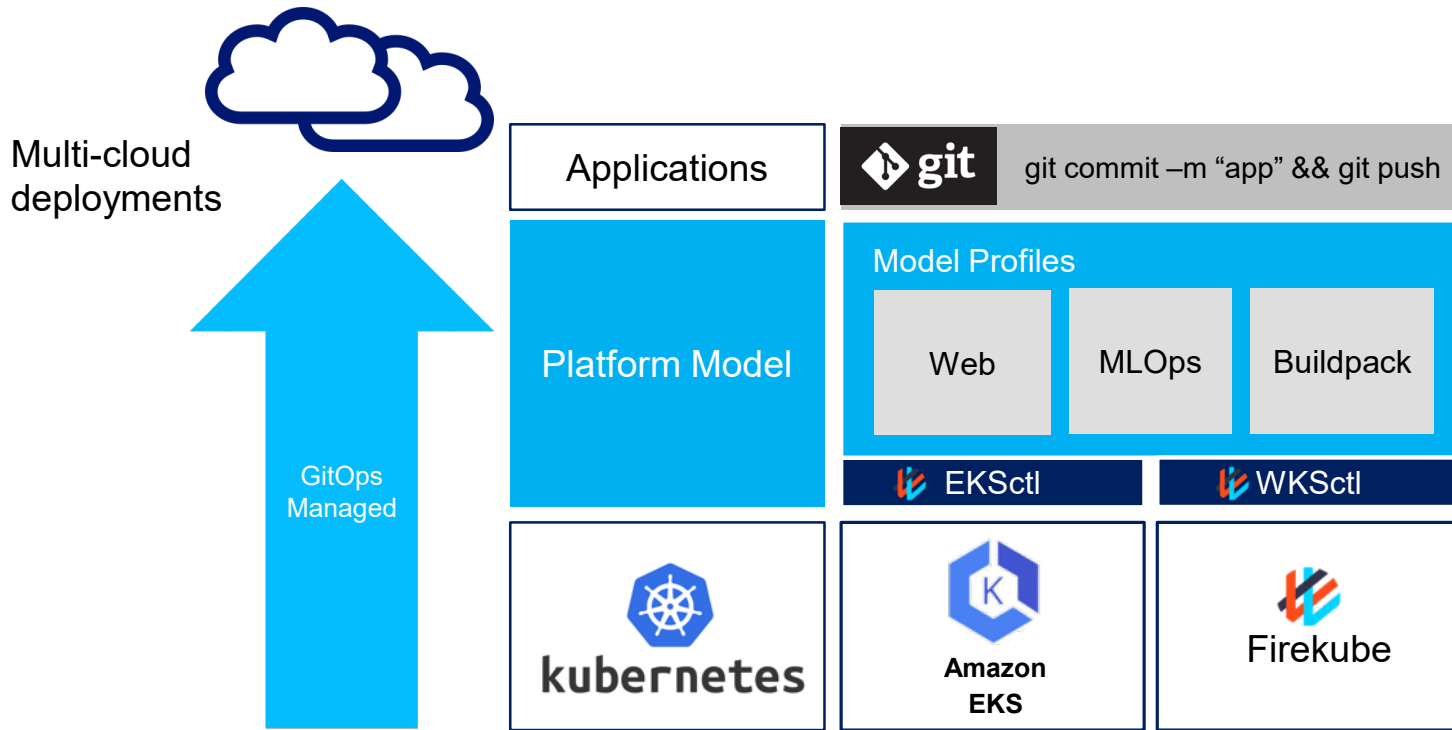
<profile-url or name>



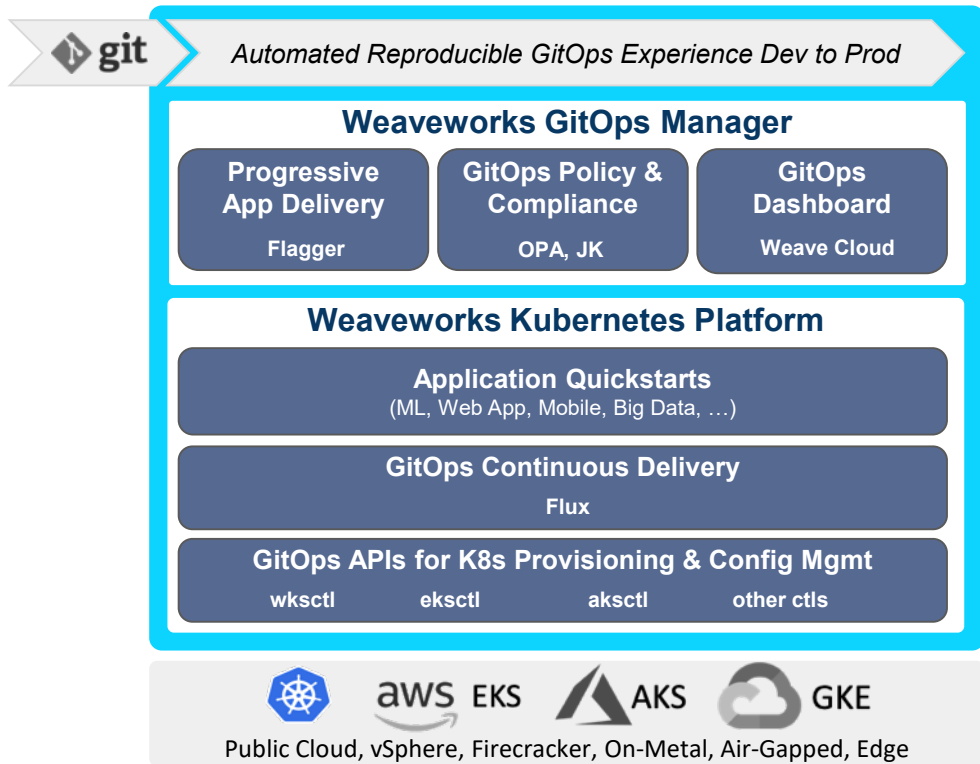
profile



Build any stack based on a MODEL



Weave: Automated Reproducible Clusters



Clusters / Weave Cloud (dev)

Provider AWS Regions us-east-1



50% CPU



50% Mem



No memory usage metrics available

N/A% CPU

Cluster components 12

⚙️ Add components


Prometheus v6.2.5	Open Prometheus Open Alertmanager
Prometheus Alertmanager v6.2.5	Open Admin
Grafana v6.2.5	Open Admin
Grafana v6.2.5	Open Admin
Grafana v6.2.5	Open Admin

Active Git branches 5

⚙️ Open Git repo

master Last commit 6 days ago by foot@weave.works	↔ 1a2c9ce
fbair-more-robust-prometheus-rules-config Last commit 6 days ago by foot@weave.works	↔ 1a2c9ce
grafana-cluster-component Last commit 6 days ago by foot@weave.works	↔ 1a2c9ce
grafana-cluster-component Last commit 6 days ago by foot@weave.works	↔ 1a2c9ce
grafana-cluster-component Last commit 6 days ago by foot@weave.works	↔ 1a2c9ce

Cluster components



Active Git branches



Cluster Dashboard

Cluster Profile Definitions



Weave Kubernetes Platform

Clusters

Models

Policies

Dashboards

Teams




Models 12

[Add models](#)



Name	K8s v.	Node groups	Components & apps	Last change	
Production cluster Standardised production cluster for base application workloads.	1.15.0	10x m5.large 3x m5.xlarge	Prometheus, metrics server, fluentd, Elasticsearch, Istio, Jaeger, Helm, Tiller, wkp-external-dns, Vault	12:45 9/10/2019 john.smith@acmecorp.com	Create a new cluster
Machine learning cluster Standardised production cluster for machine learning workloads.	1.14.0	5x m5.large 4x p2.8xlarge	Kubeflow, Knative, Prometheus, metrics server, fluentd, Elasticsearch, Istio, Jaeger, Helm, Tiller, wkp-external-dns	12:45 9/10/2019 john.smith@acmecorp.com	Create a new cluster
Search components All components needed to work on the search functionality of the website.	1.15.0	5x m5.large	Elasticsearch-cluster-search, search-ui, search-api, redis	12:45 9/10/2019 john.smith@acmecorp.com	Create a new cluster




Alerting and Events Dashboard across clusters

 Weave Kubernetes Platform

ClustersModelsPoliciesDashboardsAlertsTeams



Firing alerts1

 Monitor Alert

Production Cluster


2019-11-06T11:21:48.834Z

View dashboard

(NodeHighCPUUsage - WARNING FIRING) CPU usage has been over 10% for 5m

Impact: Pods running on that overcommitted node may cause further issues down the line
Node: gke-sock-shop-default-pool-9652982b-tljz
Node: gke-sock-shop-default-pool-9652982b-7pbs
Node: gke-sock-shop-default-pool-9652982b-fvfn
containerName: prom-node-exporter

Events2

 Drift detected

Machine Learning Cluster

2019-11-06T11:00:67.234Z

View in Github

Stack	Stack status	Drift status	Last check timestamp
cluster	UPDATE_COMPLETE	IN_SYNC	2019-11-06T11:00:67.234Z
nodegroup berlin-ai	UPDATE_COMPLETE	DRIFTED	2019-11-06T11:00:67.234Z


The definition is 6 nodes but berlin-ai is running with 5.

Production Cluster

View in Github

Stack	Stack status	Drift status	Last check timestamp
cluster	UPDATE_COMPLETE	IN_SYNC	2019-11-06T11:00:67.234Z
nodegroup group-1	UPDATE_COMPLETE	IN_SYNC	2019-11-06T11:00:67.234Z
nodegroup group-2	UPDATE_ROLLBACK_FAILED	DRIFTED	2019-11-06T11:00:67.234Z

Kubediff has detected a difference in running config.

 Policy Check Failed

Production Cluster

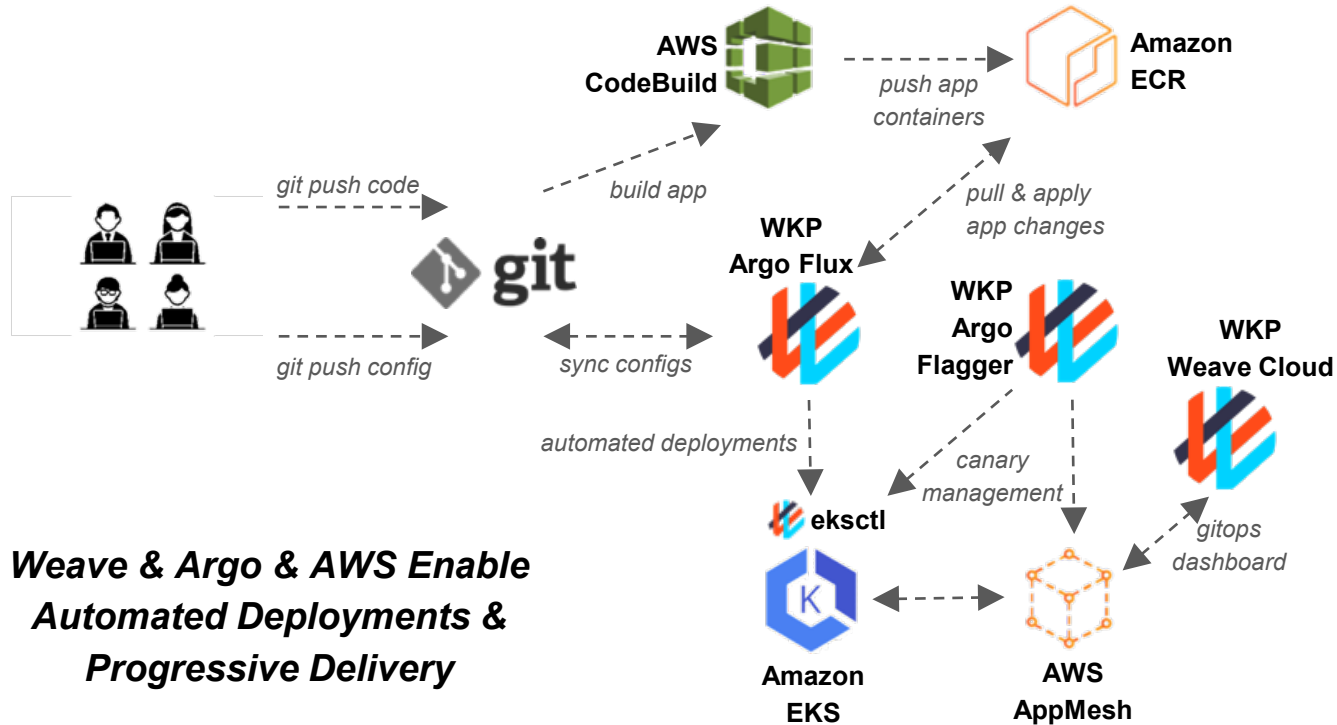
2019-11-06T09:34:23.584Z

View in Github

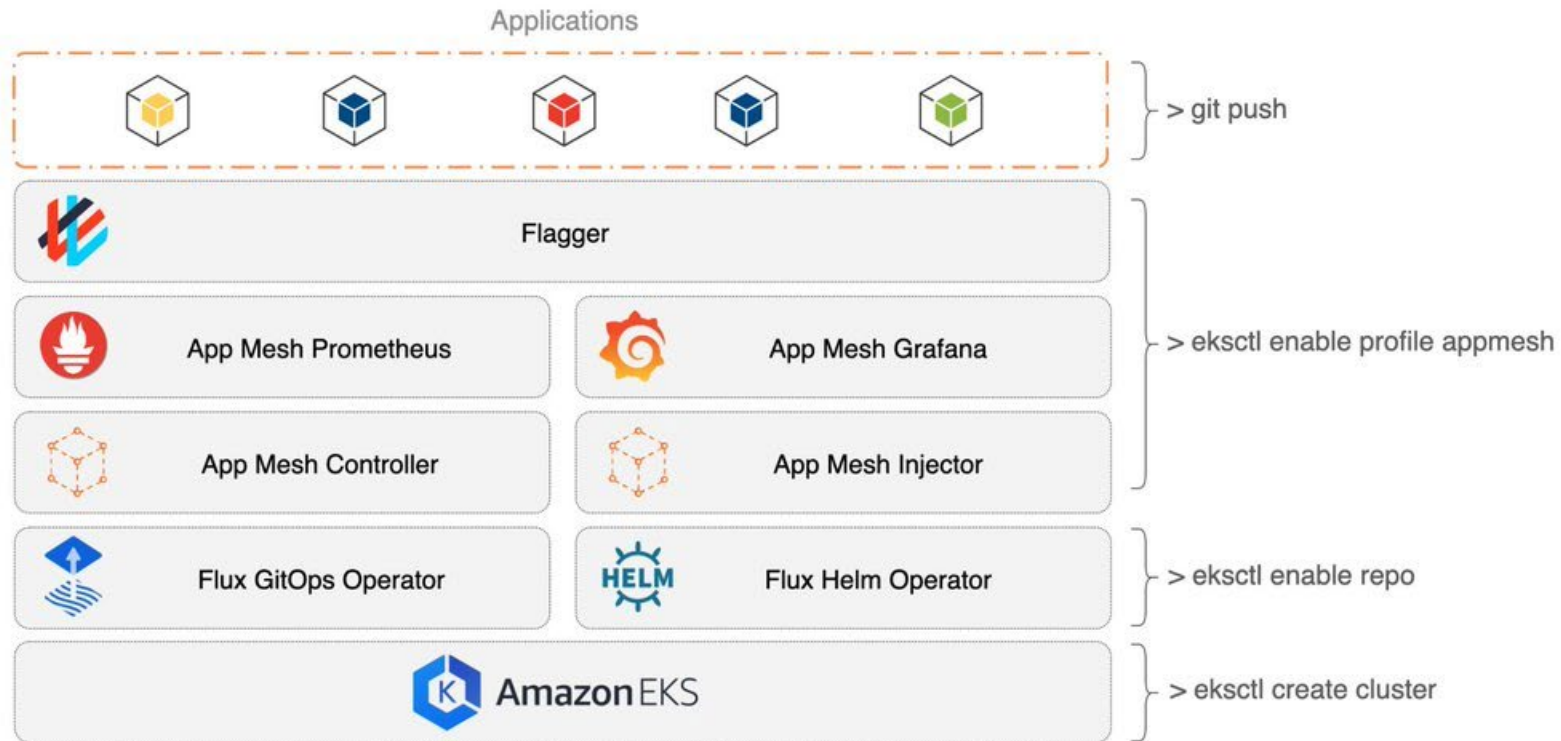
commit cac51df

hello-deployment.yaml -- Critical production deployment: "hello-world my not have fewer than two replicas

We can do Progressive Delivery this way too



Progressive Delivery Toolset



Visit our booth #S51 to meet the team, grab some swag and learn more about **automating Kubernetes with GitOps!**

Try our GitOps Hands-On for a prize.

www.weave.works