



LinkedIn Continuous Deployment

Building a whole new product

September 2024



Keshav Tyagi



Rahul Gade

Agenda

Why a New Platform?

What is LinkedIn Continuous
Deployment?

Architecture

Use Cases

LinkedIn by Numbers

Caters to >1 billion members as of Jan 2024.

Every Minute:

- More than 13K connections are formed
- 7 people are hired through LinkedIn
- More than 9K members apply for jobs
- Globally, more than 1.5M feed updates are viewed
- Globally, ~138 hours of learning content is consumed

What are deployments

- Deployments are last mile in SDLC to delivery feature to customers
- Application(s) can be deployed through
 - Manual deployments
 - Continuous deployments

Landscape of deployments in LinkedIn

Given the scale of LinkedIn, we have to support

- Stateless services
- Stateful services
- Offline services (Flink / Airflow's latest build & DAGs)
- ML Models

Scale of deployments at LinkedIn

Any guess on how many services we have at LinkedIn ?

Any guess on how many deployments we do at LinkedIn ?

Deployments in LinkedIn

Number of services ~ 7k

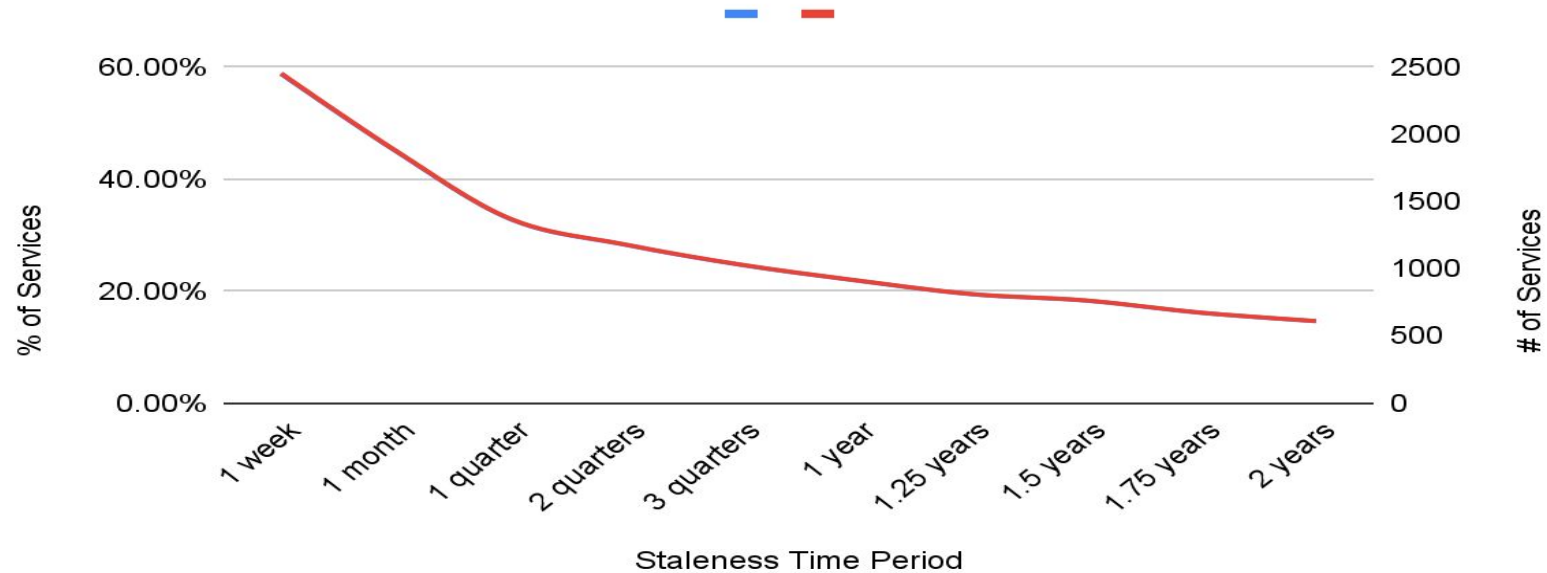
Number of unique pipelines (aka DAGs) ~ 12k

Number of deployments per month ~ 1 million

Deployment Staleness

Service Deployment Staleness

Services that are not deployed in a given time period



Why Rebuild?

- Existing solution was built around single service paradigm
- Does not support complex workflows
- Lack of trust on validation tools
- Onboarding new deployment schedulers / validation tools was expensive and painful
- Workflow schema was too verbose and leaked implementation details

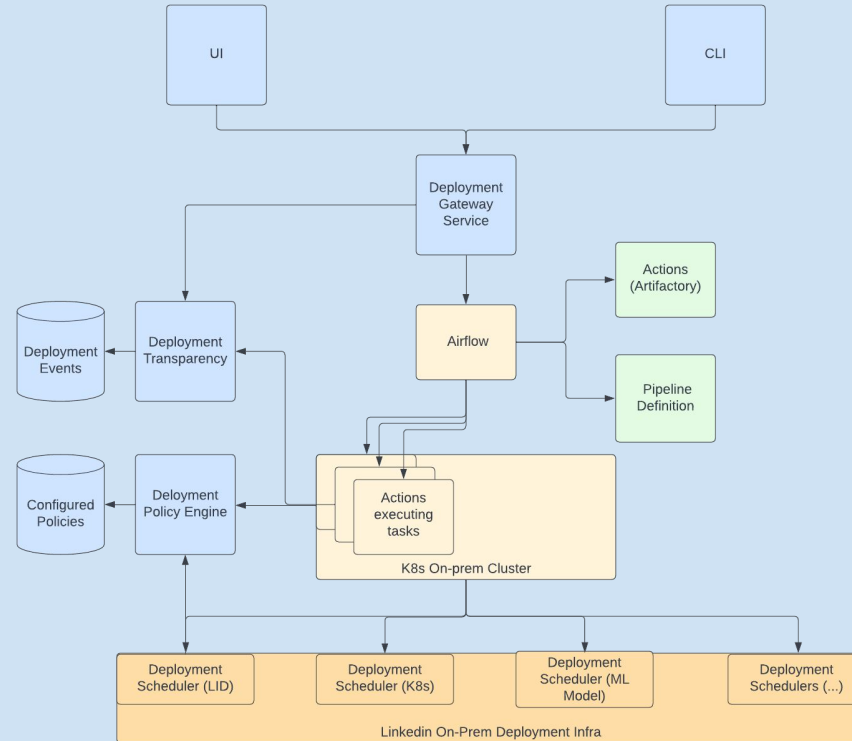
Personas

- Platform Admin
- Our Partners
- Our Users

LinkedIn's continuous deployment vision

- Modular Design
- Democratise the Platform for Partners
- No Code Low Code Pipelines
- Single Deployment Solution across LinkedIn
- Increase continuous deployment adoption
- Engineering efficiency metrics
- Enforce best deployment practices
- Re-enforce trust in validation tools across LinkedIn

10000 Ft view | Architecture



- Chat
- Multiproducts
- Campaigns
- Crews
- Users
- Manage RDevs
- Deployments
- Code Health Hub
- Company Tech Map
- Products
- Workflows
- Mae
- Support
- Mariia Paniutina

NEW PIPELINE EXECUTION

PIPELINE **RUN HISTORY** ACTIONS PREFERENCES

Auto-refresh

Targets	Started By	Start Time	End Time	Duration	Status	Actions
my-mp-1 59.0.885 > All targets 0.2254.0	user1	09/04/2024 09:03 PDT	N/A	N/A	In Progress	VIEW INTENT DETAILS
my-mp-1 59.0.866 > All targets 0.2248.0	github-app-crt2	09/02/2024 23:08 PDT	09/03/2024 12:32 PDT	13h 23min 58s	Failed	VIEW INTENT DETAILS
my-mp-1 59.0.860 > All targets 0.2245.0	github-app-crt2	08/30/2024 13:52 PDT	09/02/2024 09:19 PDT	2d 19h 27min 7s	Succeeded	VIEW INTENT DETAILS
my-mp-1 59.0.859 > All targets 0.2244.0	github-app-crt2	08/30/2024 11:49 PDT	08/30/2024 11:52 PDT	2min 39s	Failed	VIEW INTENT DETAILS
my-mp-1 59.0.858 > All targets 0.2242.0	github-app-crt2	08/29/2024 15:32 PDT	08/30/2024 11:07 PDT	19h 35min 4s	Succeeded	VIEW INTENT DETAILS

Rows per page: 10 1-10 < >

my-custom-pipeline-a

Some description here: aka getting whole MP to be deployed from staging to prod... etc

[Similar pipelines](#)

Entities

- entity-1 (MultiProduct)
- entity-2 (MultiProduct)

LIVE VIEW RUNS HISTORY **ACTIONS**

Recommended best practices

Actions to ensuring stability and reducing errors.

Full deployment Deploy a new version through thorough testing and automation to ensure stability and functionality

The deployment will go through stages specified in the pipeline, including building, deploying, verifying, and promoting releases

MultiProduct entity-1	MultiProduct entity-2
Application Version 0.1.23	0.4.1
Config Version v.1.212	Config Version
<small>If not selected - latest config version will be applied</small>	<small>If not selected - latest config version will be applied</small>
0.1.23_0.0.6	
<small>If not selected - latest manifest version will be applied</small>	

Policy Violation: Global moratorium
 This deployment falls within a moratorium period, during which no deployments are allowed due to critical system maintenance, high-traffic periods, or other significant reasons.

I acknowledge that I am bypassing policy deployment restrictions and proceeding with the deployment, understanding the potential risks and implications.

CANCEL EXECUTE

Emergency response strategy

Actions to use during emergencies to quickly fix urgent issues to restore site stability.

- Rollback** - Quickly revert to a previous stable version, sacrificing some features to restore normal operations
- Roll-forward** - Accelerate a tested, safe version to production to resolve issues quickly
- Hot-fix** - Deploy an urgent fix rapidly to address critical issues in production

my-custom-pipeline-a

Some description here: aka getting whole MP to be deployed from staging to prod... etc

[Similar pipelines](#)

Entities

- entity-1 (MultiProduct)
- entity-2 (MultiProduct)

LIVE VIEW RUNS HISTORY ACTIONS **POLICIES**

Policy violations

Runs with overridden policies

Policy name	Description	Severity	Runs
Deprecated dependency	Deployment includes deprecated direct or transitive library dependencies	Medium	#1234 , #12345
Off-scheduled deployment	Deployment outside of scheduled window	Medium	#1234 , #1232 , #12352 , #12352 , #1123

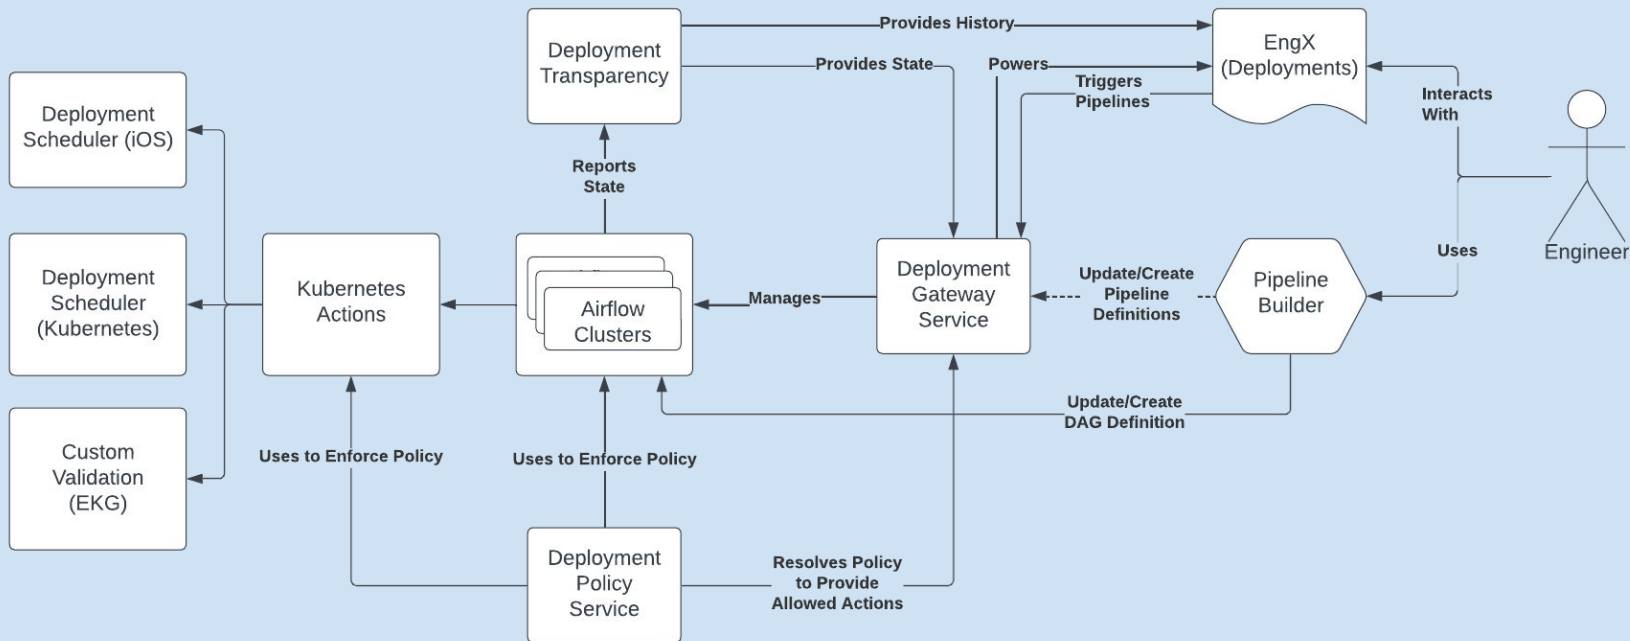
Rows per page: 10 1-1 of 1

Applied Policy

Policy name	Description	Can be overridden
Deprecated dependency	Deployment includes deprecated direct or transitive library dependencies.	YES
Off-scheduled deployment	Deployment outside of scheduled window	YES
ACL ownership check	Only ACLs/crews which are part of DEVELOPMENT_TEAM or RELIABILITY_TEAM roles for given entity can authorise deployment	NO
Deprecated version deployment	Deploying deprecated entity version/versions	YES

Rows per page: 10 1-1 of 1

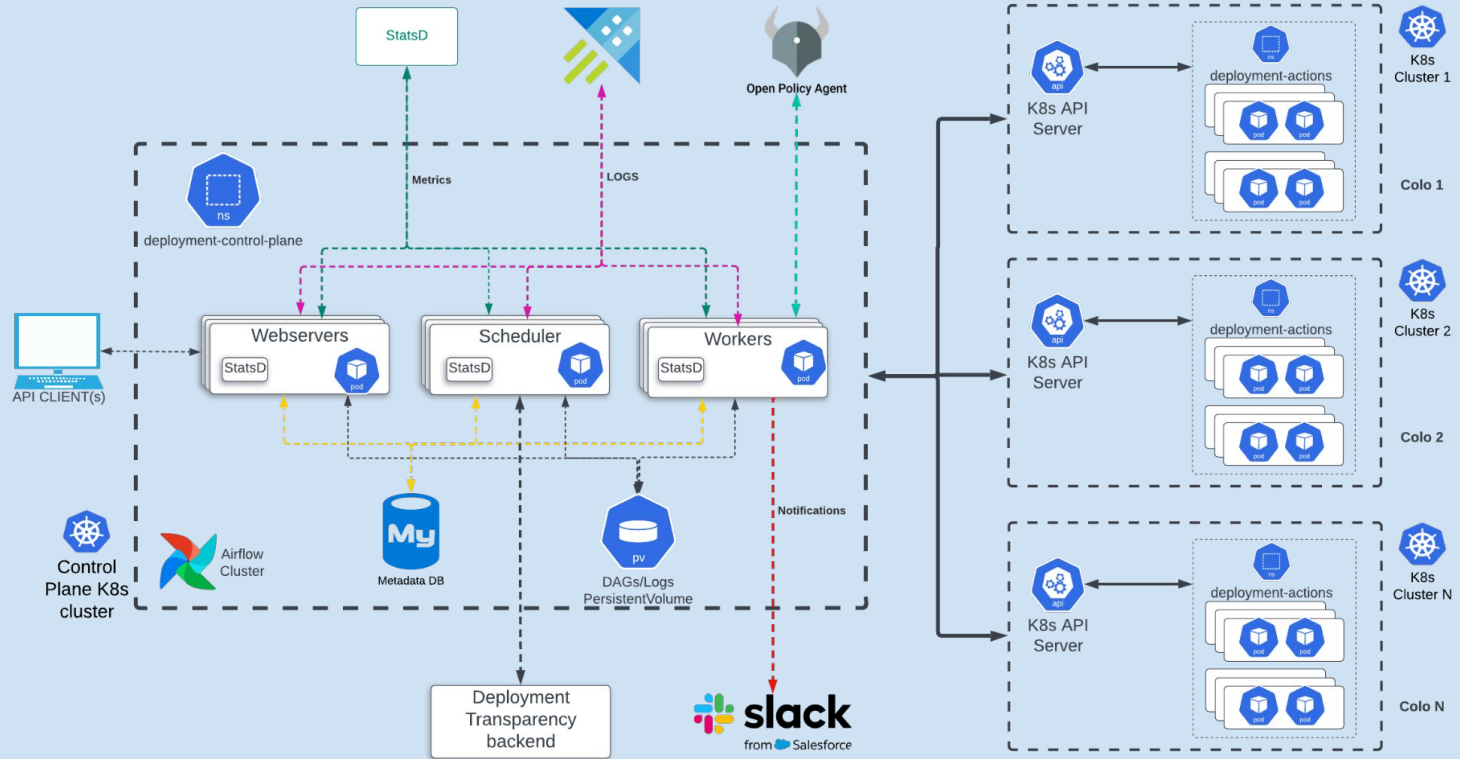
Under the hood | Architecture



Under the hood | Architecture




Orchestration Ecosystem

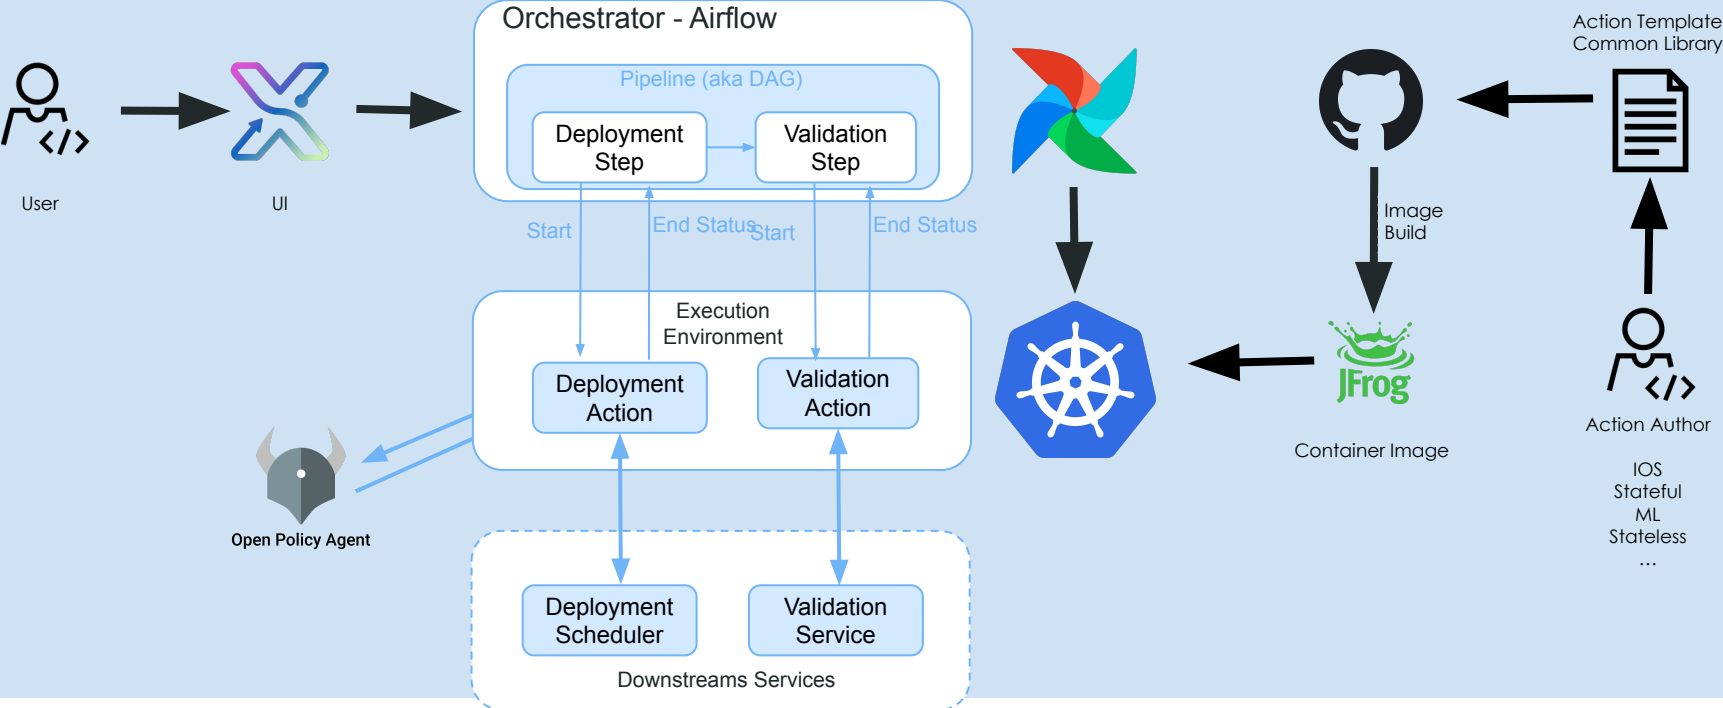


Actions Ecosystem

- Portable pipeline steps as Docker containers
- Able to perform almost ANY pipelines
- Actions can be authored any language



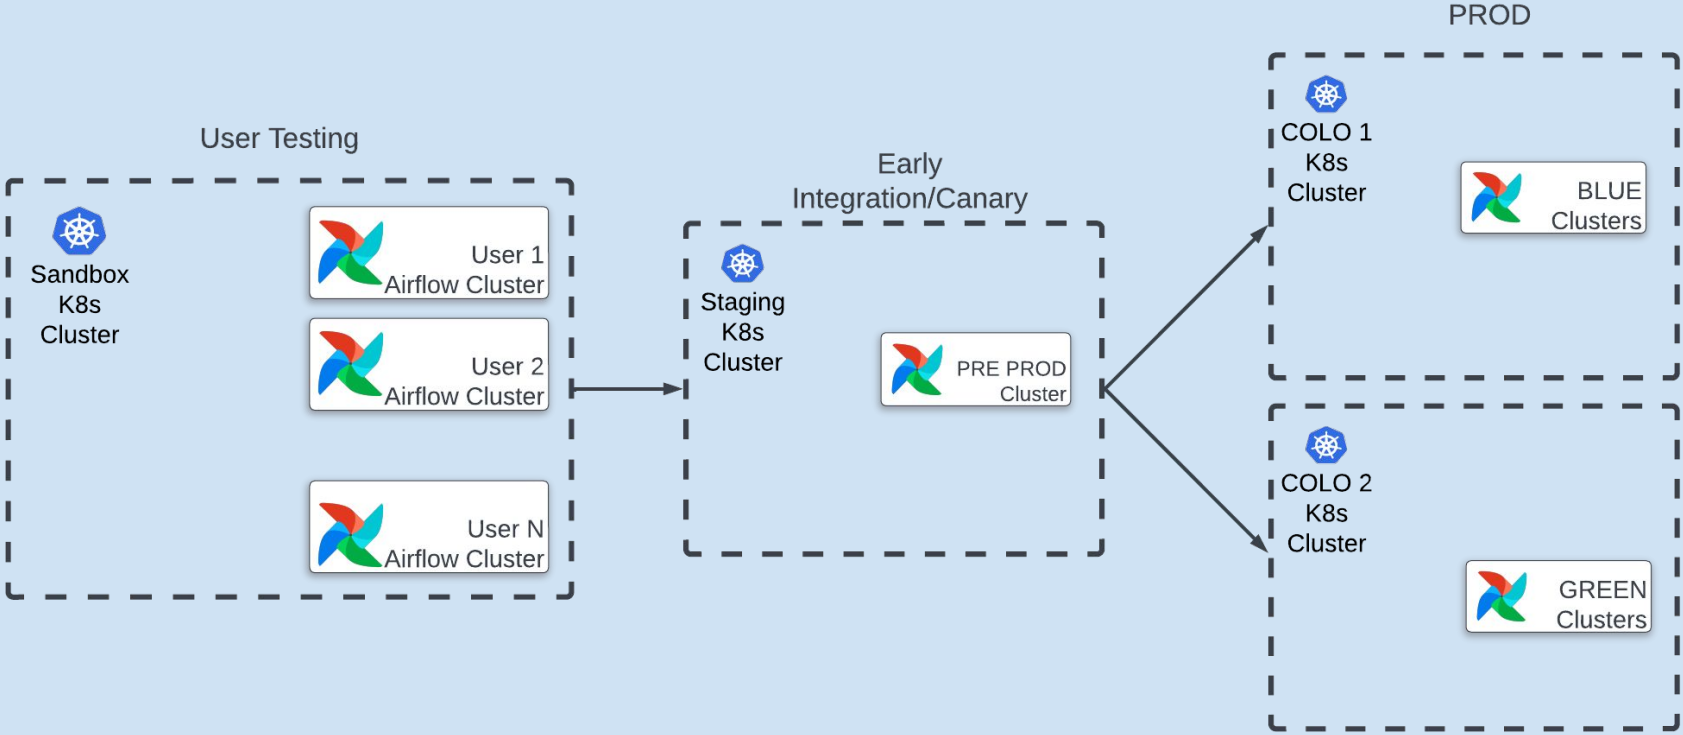
Under the hood | Actions Ecosystem



How we rollout features? | Action Version Upgrades

Variables to the rescue

How we rollout features? | Zero Downtime upgrades



Disclaimer

We are NOT an Airflow as a platform team at LinkedIn!

Pipeline Governance

- Operate on only one production colos at a time
- Pause Between Fabrics
- Decide at runtime which colos to deploy
- Ensure all pipelines are enforcing policies

Observability | Actions Ecosystem


How will different persona will debug things.

- Admins
- Action Owners
- Users

Low code, no code | Onboarding

Build your pipeline by just answering a few questions through CLI/UI.

```
2  "$schema": "https://crt-backend/api/v1/schema/workflow_schema_1.json",
3  "paths": [
4    {
5      "name": "NO_AUTOMATION_PATH",
6      "groups": {
7        "Canary_NO_AUTOMATION_lit": {
8          "execution_type": "CANARY",
9          "fabrics": [
10           {
11             "fabric_name": "lit-1",
12             "display_name": "LIT",
13             "allowed_tag": "DEFAULT"
14           }
15         ],
16         "strategies": {
17           "auto_retry": {
18             "enable": true
19           }
20         }
21       },
22       "Promotion_NO_AUTOMATION_lit": {
23         "execution_type": "CANARY_PROMOTION",
24         "fabrics": [
25           {
26             "fabric_name": "lit-1",
27             "display_name": "LIT",
28             "allowed_tag": "DEFAULT"
29           }
30         ],
31         "strategies": {
32           "auto_retry": {
33             "enable": true
34           }
35         }
36       },
37       "NO_AUTOMATION_ei": {
38         "execution_type": "FULL_DEPLOYMENT",
39         "fabrics": [
40           {
41             "fabric_name": "ei1",
42             "display_name": "ei1",
43             "allowed_tag": "DEFAULT"
44           }
45         ]
46       }
47     }
48 ]
```

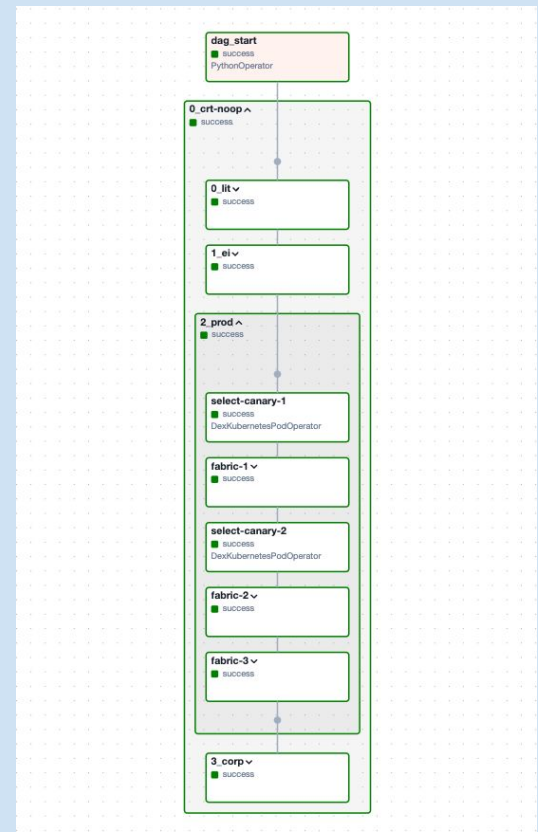


Answer:

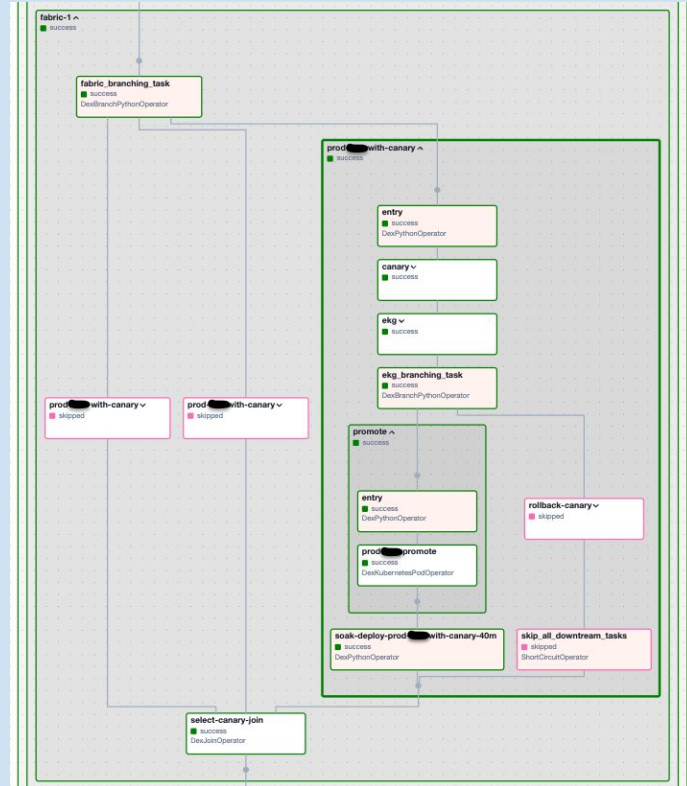

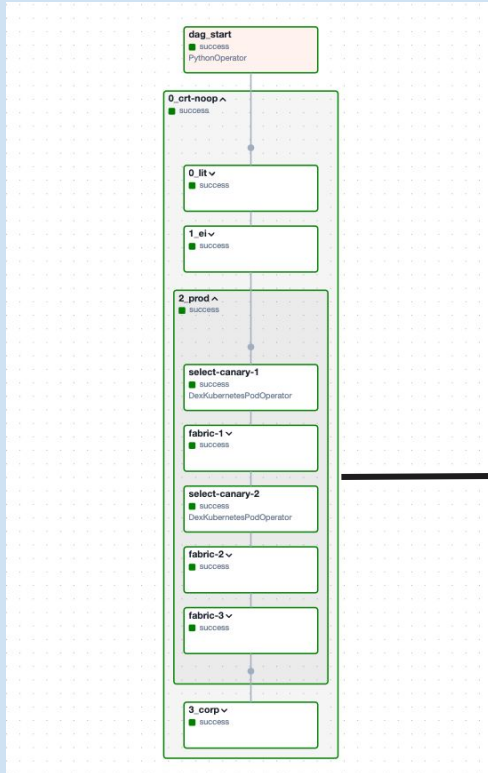
- what : deployable-foo
- how: K8s
- where:
 - Target-colo1 (ei1)
 - Target-colo2 (lit-1)
 - ...

Pipeline intent to Airflow DAG

```
1  version: v1
2  name: deployment-foo
3  targets:
4    - clusters:
5        - ei-1
6          - lit-1
7            - corp-1
8              - prod-1
9                - prod-2
10                 - prod-3
11      types:
12        - lid
13      deployable_name: deployment-foo
14  notification:
15    slack: deployment-foo-alerts
16
```



Pipeline intent to Airflow DAG



Migration Goals and Challenges

- No efforts from users to migrate to our new platform
- Coordinate parallel migration across deployment stack (moving to k8s)

Validation Actions

- Canary Validation
- Test Manager (Run Integration Tests)
- Epsilon (Replays prod traffic on set of hosts)
- Smart Canaries (Real time detection of which colo to deploy based of traffic patterns)
- And many more

Newer use cases supported

- Support for touchless Schema Deployment
- Support for Stateful services (e.g Pinot)
- Support for provisioning clusters
- Scheduled deployments
- (Future use cases) Support for chaos monkey

Q&A