

# Running Apache Airflow reliably with Kubernetes

...

and other open source software





# Greg Neiheisel, CTO



# On Deck

- Quick Airflow / Kubernetes overview
- Running Airflow at Scale
- Major system design considerations
- Lessons and best practices we've learned along the way



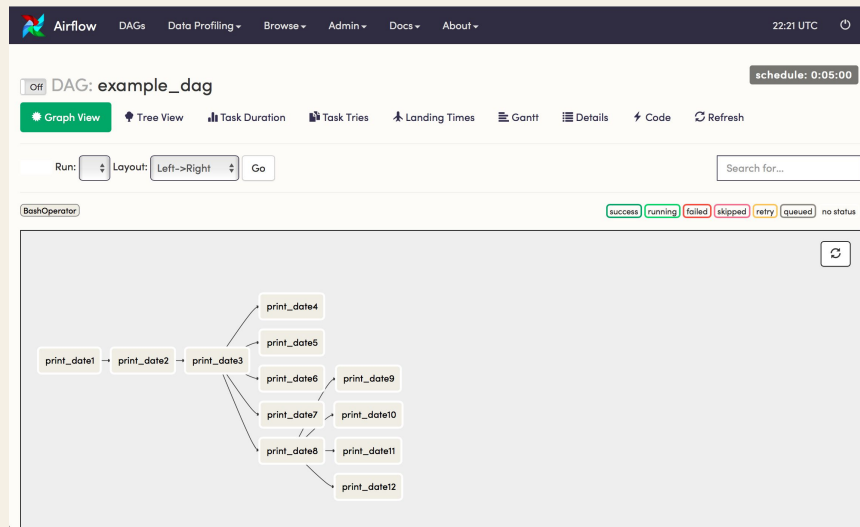
# What is Apache Airflow?

- A task scheduler written in Python to programatically author, schedule, and monitor dependency driven workflows (DAGs)
  - Pluggable architecture, focused on ETL, ML use-cases
  - Lots of existing building blocks
- Top-level Apache Project
  - 11,000+ stars on github
  - 6,000+ commits
  - 700+ contributors



# Airflow core concepts

- **DAGs** - created in code, typically associated with a cron schedule
- **DAG Runs** - typically execution of a dag for a given execution date
- **Task Instances** - represents an execution of a node in the DAG

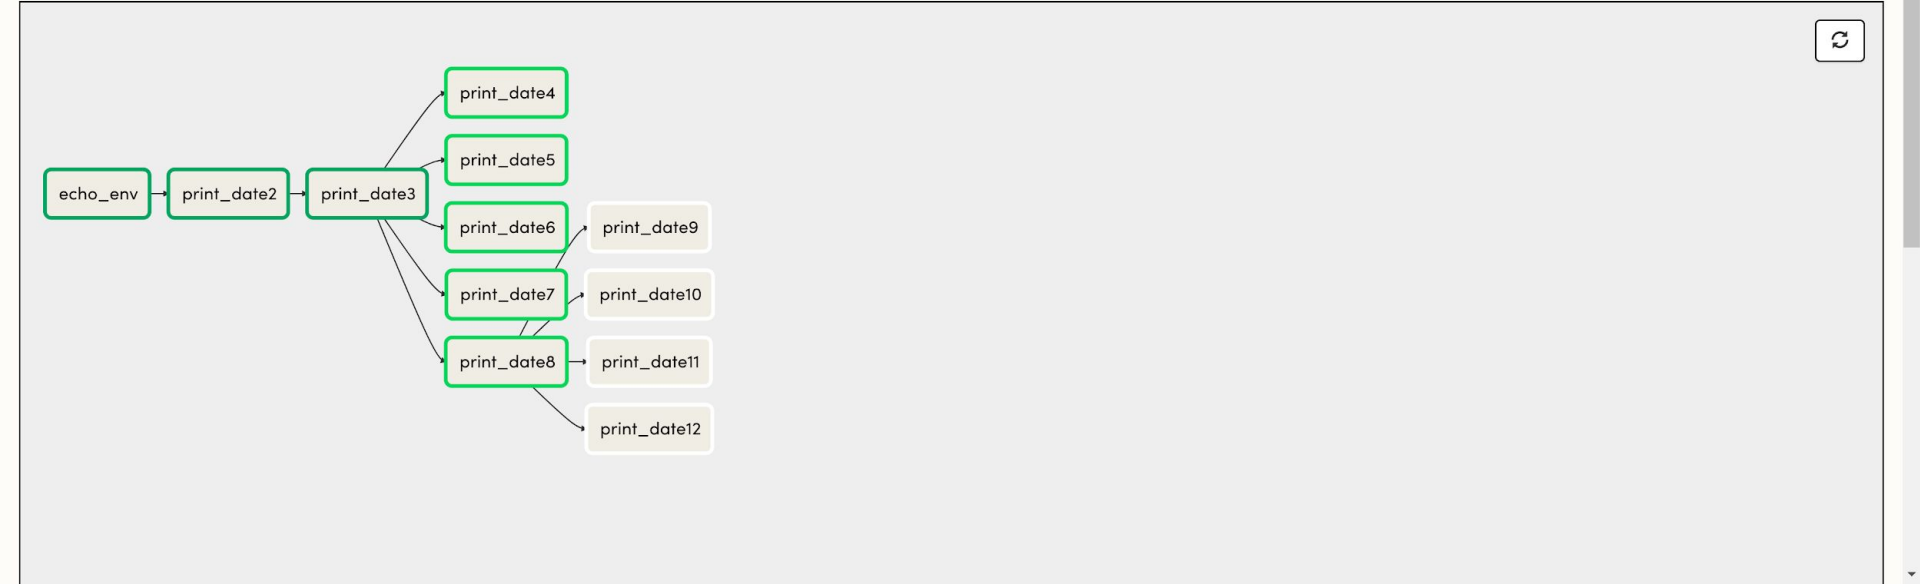


On DAG: example\_dag\_two schedule: 0:05:00

Graph View Tree View Task Duration Task Tries Landing Times Gantt Details Code Refresh Delete

running Base date:  Number of runs:  Run:  Layout:

BashOperator success running failed skipped retry queued no status

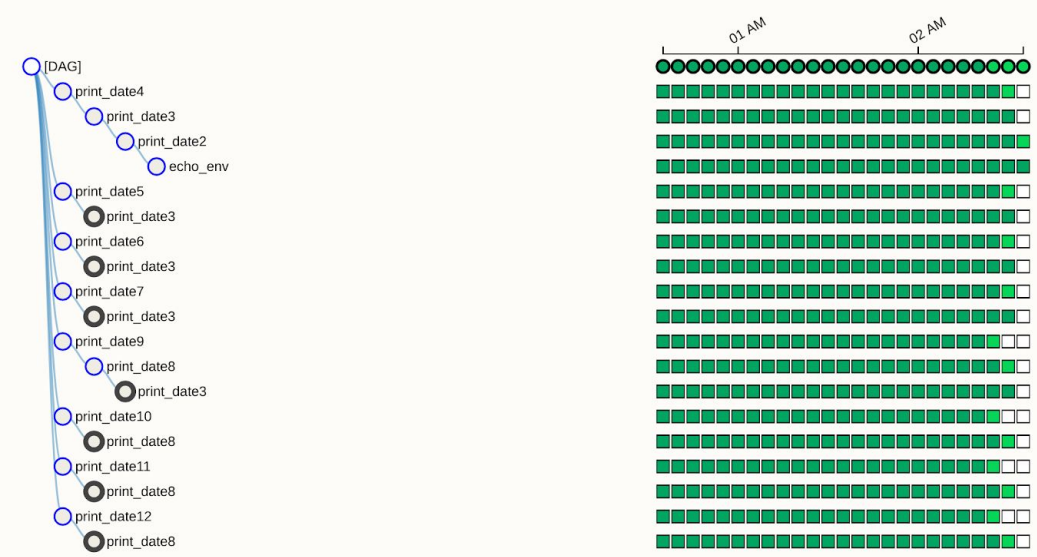


On DAG: example\_dag\_two schedule: 0:05:00

- Graph View
- Tree View
- Task Duration
- Task Tries
- Landing Times
- Gantt
- Details
- Code
- Refresh
- Delete

Base date: 2018-01-03 07:35:00 Number of runs: 25 Go

BashOperator success running failed skipped retry queued no status



☒ DAG: example\_dag\_two

schedule: 0:05:00

 Graph View
  Tree View
  Task Duration
  Task Tries
  Landing Times
  Gantt
  Details
  Code
  Refresh
  Delete

Task Instance: echo\_env 2018-01-03 07:25:00

[Task Instance Details](#)
[Rendered Template](#)
[Log](#)
[XCom](#)

## Log by attempts

```
Dependencies all met for <TaskInstance: example_dag_two.echo_env 2018-01-03T07:25:00+00:00 [queued]>
Dependencies all met for <TaskInstance: example_dag_two.echo_env 2018-01-03T07:25:00+00:00 [queued]>
```

Starting attempt 1 of 2

```
Executing <Task(BashOperator): echo_env> on 2018-01-03T07:25:00+00:00
```

```
Running: ['bash', '-c', 'airflow run example_dag_two echo_env 2018-01-03T07:25:00+00:00 --job_id 7976 --raw -sd DAGS_FOLDER/example-dag.py --cfg_path /tmp/tmptu2qbsac']
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:20,812] {settings.py:174} INFO - settings.configure_orm(): Using pool settings. pool_size=5, pool recycle=1800, pid=55847
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,510] {__init__.py:51} INFO - Using executor CeleryExecutor
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,777] {models.py:273} INFO - Filling up the DagBag from /usr/local/airflow/dags/example-dag.py
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,806] {cli.py:520} INFO - Running <TaskInstance: example_dag.two.echo_env 2018-01-03T07:25:00+00:00 [running]> on host galactic-cosmology-3441-
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,842] {bash_operator.py:77} INFO - Tmp dir root location:
```

```
Job 7976: Subtask echo_env /tmp
```

```
Job 7976: Subtask echo env [2019-04-11 16:48:21.842] {bash operator.py:86} INFO - Exporting the following env vars:
```

```
Job 7976: Subtask echo env AIRFLOW CTX DAG ID=example dag two
```

```
Job 7976: Subtask echo env ATRELOW CTX TASK ID=echo env
```

```
Job 7976: Subtask echo env AT/BELOW CTX EXECUTION DATE=2018-01-03T07:25:00+00:00
```

```
Job 7976: Subtask echo env ATRELOW CTX DAG RUN ID=scheduled 2018-01-03T07:25:00+00:00
```

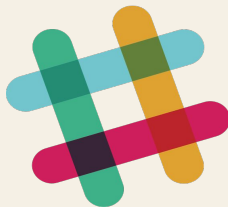


example\_dag\_two

```
1 from airflow import DAG
2 from airflow.operators.bash_operator import BashOperator
3 from datetime import datetime, timedelta
4
5 default_args = {
6     'owner': 'schnie',
7     'depends_on_past': False,
8     'start_date': datetime(2018, 1, 1),
9     'email_on_failure': False,
10    'email_on_retry': False,
11    'retries': 1,
12    'retry_delay': timedelta(minutes=5),
13 }
14
15 dag = DAG('example_dag_two',
16          max_active_runs=3,
17          schedule_interval=timedelta(minutes=5),
18          default_args=default_args)
19
20 t1 = BashOperator(
21     task_id='echo_env',
22     bash_command='sleep 120s',
23     dag=dag)
24
25 t2 = BashOperator(
26     task_id='print_date2',
27     bash_command='sleep 120s',
28     dag=dag)
29
```

# Times are changing

- Wider Use Cases
  - ETL, ML, Reporting, Data Integrity
- Higher Usage
  - More teams with different skill sets and goals for Airflow usage
  - More DAGs running more frequently
- Stricter SLAs
- More complex core components (executors, operators, etc)
  - Kubernetes, Mesos, Spark, etc.
- Immutable infrastructure



10 data engineers  
240+ active DAGs  
5400+ tasks per day

...as of April '18...

<https://speakerdeck.com/vananth22/operating-data-pipeline-with-airflow-at-slack?slide=6>



# Airflow is a highly-available, mission-critical service

- Automated Airflow deployments
- Continuous delivery
- Support 100s of users and 1,000s of tasks per day
- Security
- Access controls
- Observability (Metrics / Logs)
- Autoscaling / Scale to zero-ish



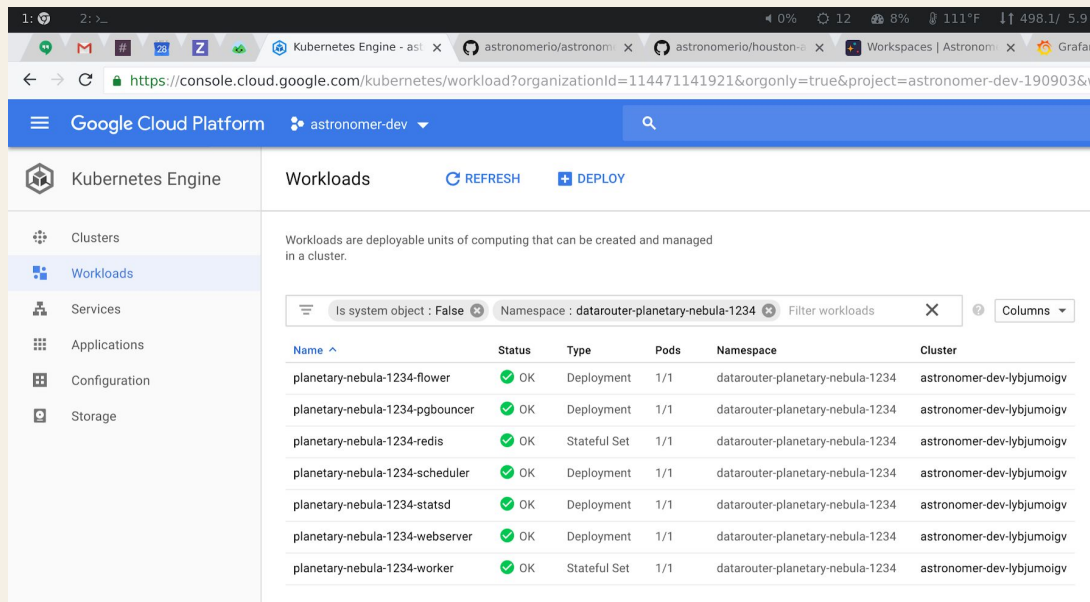
# Kubernetes

Kubernetes is a portable, extensible open-source platform for managing containerized workloads and services, that facilitates both declarative configuration and automation



# Kubernetes

Applications are broken into  
smaller, independent pieces  
and can be deployed and  
managed dynamically



The screenshot shows the Google Cloud Platform console for the 'astronomer-dev' project. The left sidebar lists navigation options: Clusters, Workloads (selected), Services, Applications, Configuration, and Storage. The main panel displays the 'Workloads' page, which includes a 'REFRESH' button and a 'DEPLOY' button. Below this, a text box explains: 'Workloads are deployable units of computing that can be created and managed in a cluster.' A filter bar shows 'Is system object : False', 'Namespace : datarouter-planetary-nebula-1234', and a 'Filter workloads' button. A table lists the workloads with columns for Name, Status, Type, Pods, Namespace, and Cluster.

Name	Status	Type	Pods	Namespace	Cluster
planetary-nebula-1234-flower	OK	Deployment	1/1	datarouter-planetary-nebula-1234	astronomer-dev-lybjumoigv
planetary-nebula-1234-pgbouncer	OK	Deployment	1/1	datarouter-planetary-nebula-1234	astronomer-dev-lybjumoigv
planetary-nebula-1234-redis	OK	Stateful Set	1/1	datarouter-planetary-nebula-1234	astronomer-dev-lybjumoigv
planetary-nebula-1234-scheduler	OK	Deployment	1/1	datarouter-planetary-nebula-1234	astronomer-dev-lybjumoigv
planetary-nebula-1234-statsd	OK	Deployment	1/1	datarouter-planetary-nebula-1234	astronomer-dev-lybjumoigv
planetary-nebula-1234-webserver	OK	Deployment	1/1	datarouter-planetary-nebula-1234	astronomer-dev-lybjumoigv
planetary-nebula-1234-worker	OK	Stateful Set	1/1	datarouter-planetary-nebula-1234	astronomer-dev-lybjumoigv



# Kubernetes

- **Pod** - One or more colocated containers, share volumes, ports
- **Deployment** - Higher level abstraction, manages pods, replica sets
- **Stateful Set** - Similar to Deployment, except each replica gets a stable hostname and can mount persistent volumes
- **Daemon Set** - Replica pods deployed to each node
- **Namespace** - Virtual cluster backed by the same physical cluster



# Declarative Service Definition with Kubernetes / Helm

Helm helps you manage Kubernetes applications — Helm Charts helps you define, install, and upgrade even the most complex Kubernetes application.

```
1: 2: >_
helm.astronomer.io worker-statefulset.yaml
1 #####
2 ## Airflow Worker StatefulSet
3 #####
4 {{- if eq .Values.executor "CeleryExecutor" }}
5 kind: StatefulSet
6 apiVersion: apps/v1
7 metadata:
8   name: {{ .Release.Name }}-worker
9   labels:
10     tier: airflow
11     component: worker
12     release: {{ .Release.Name }}
13     workspace: {{ .Values.platform.workspace | quote }}
14     chart: "{{ .Chart.Name }}"
15     heritage: {{ .Release.Service }}
16 spec:
17   serviceName: {{ .Release.Name }}-worker
18   replicas: {{ .Values.workers.replicas }}
19   podManagementPolicy: Parallel
20   selector:
21     matchLabels:
22       tier: airflow
23       component: worker
24       release: {{ .Release.Name }}
25       workspace: {{ .Values.platform.workspace | quote }}
```

<https://github.com/astronomer/helm.astronomer.io/tree/master/charts/airflow>



```
helm install -n airflow-prod charts/airflow
```





# Airflow Executors

A pluggable way to scale out Airflow workloads. Responsible for running  
`airflow run ${dag_id} ${task_id} ${execution_date}`  
somewhere.

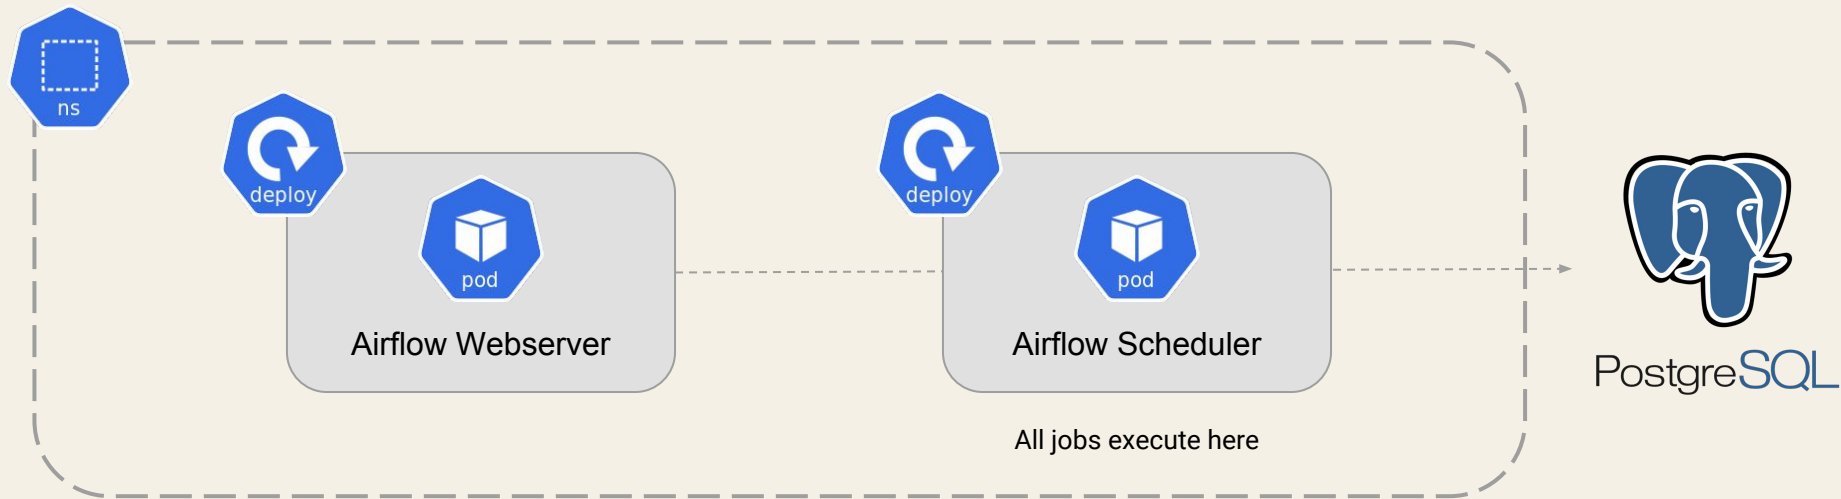


# Executors - Sequential/Local

- Fork off and run tasks in subprocess
- Good for simple workloads
- Eventually things need to scale out



# Executors - Local Executor

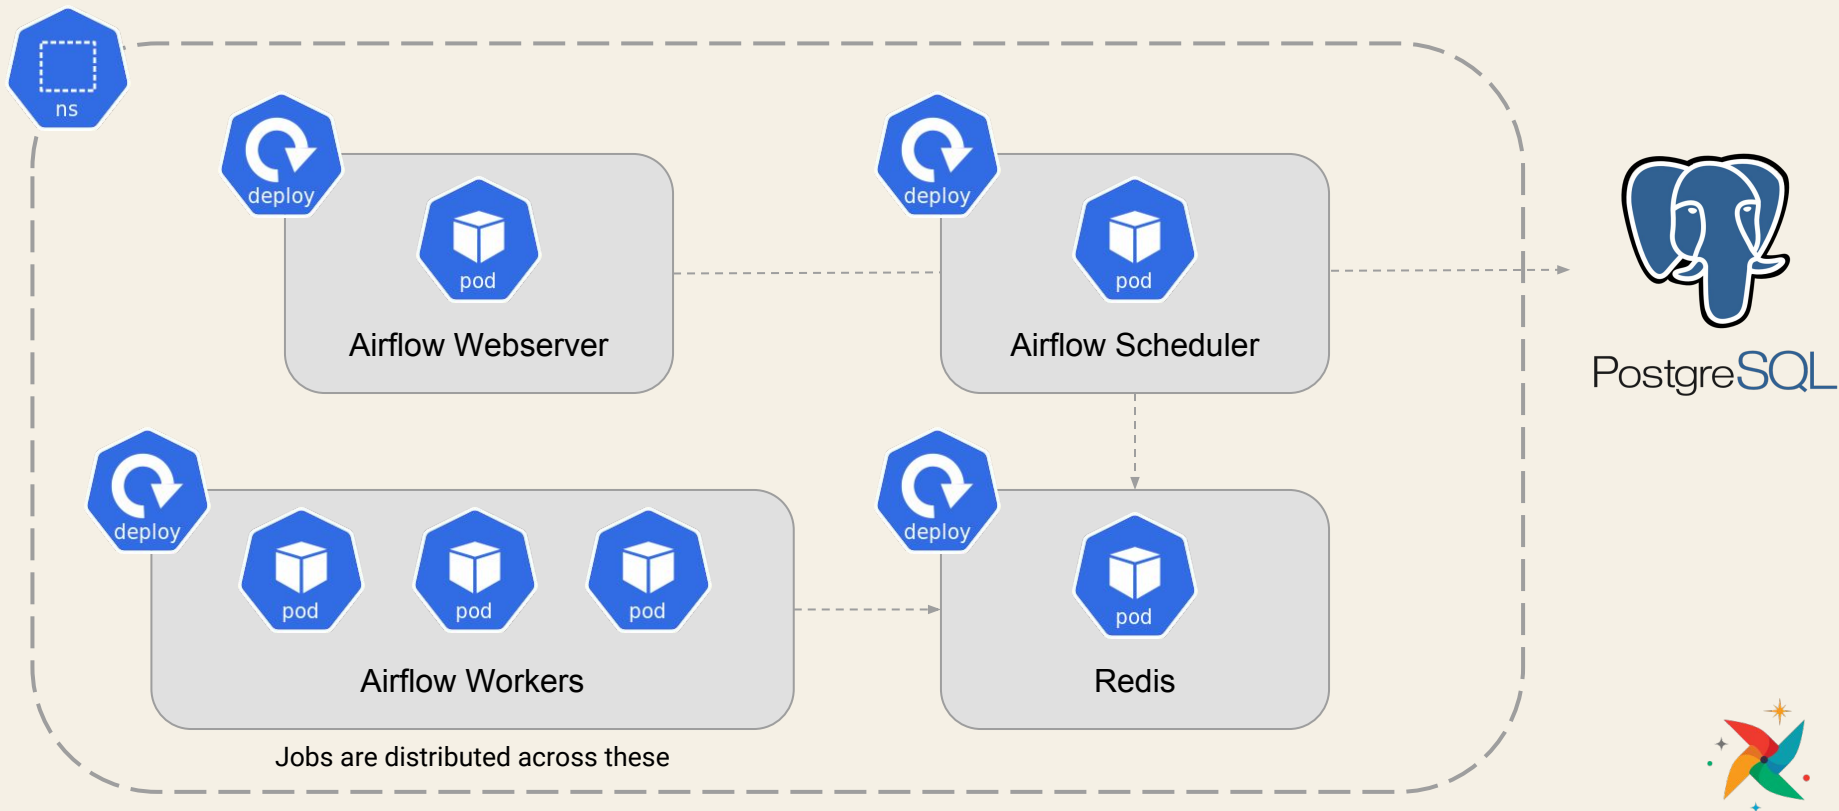


# Executors - Celery Executor

- Distributed Task Queue
- Redis, RabbitMQ, etc dependency
- Configure number of workers
  - Kubernetes HorizontalPodAutoscaler
- Configure worker size
  - Kubernetes resource requests / limits



# Executors - Celery Executor

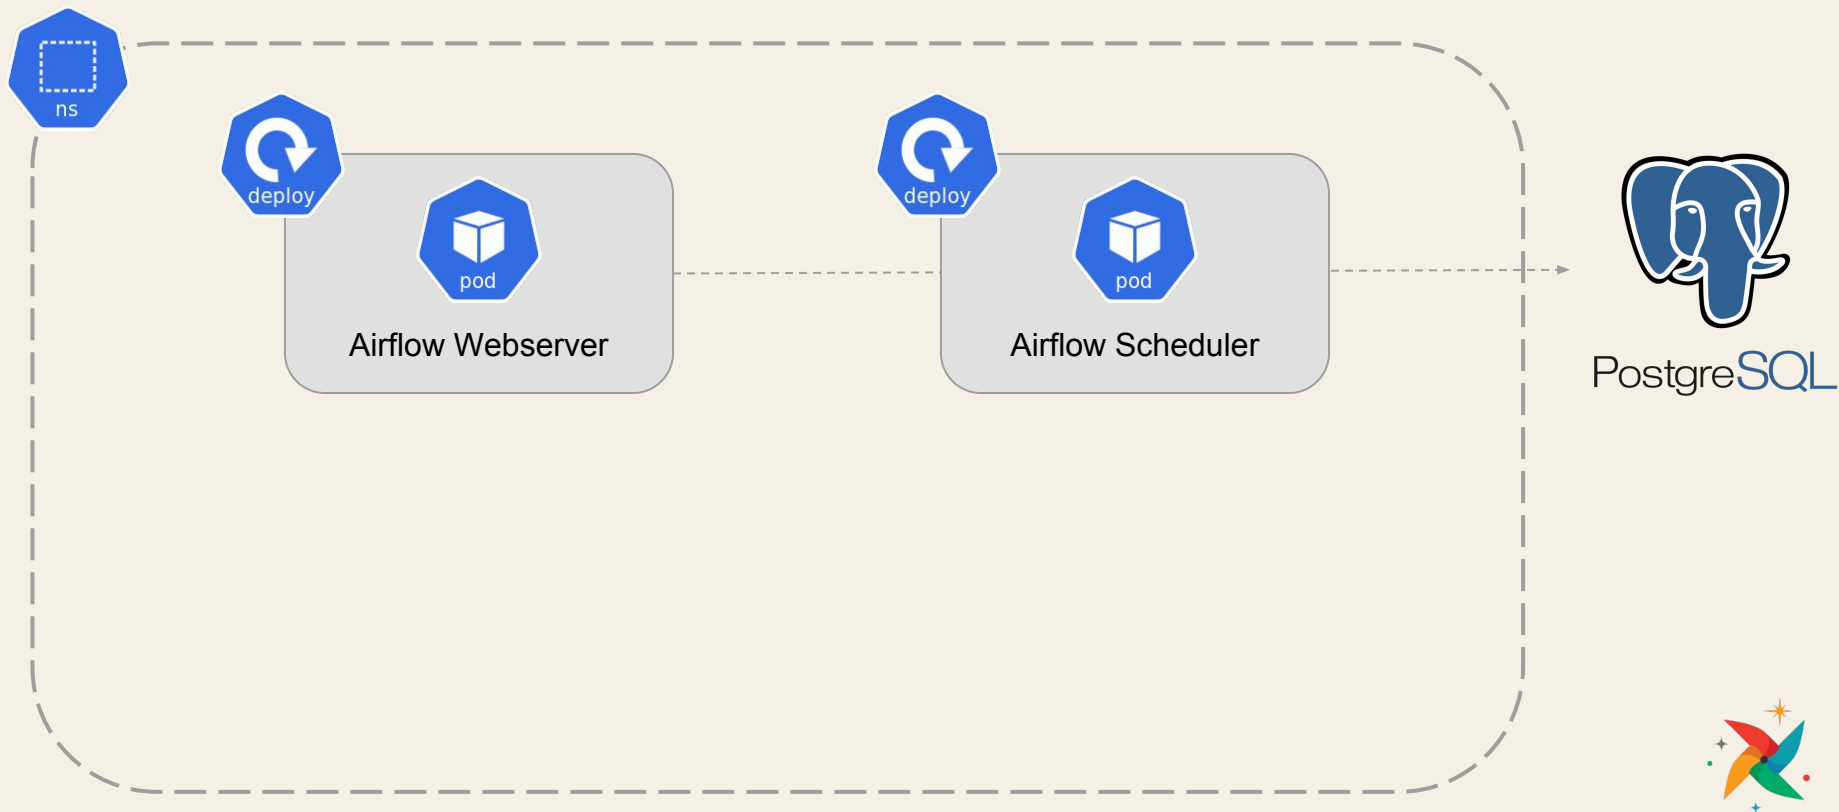


# Executors - Kubernetes Executor

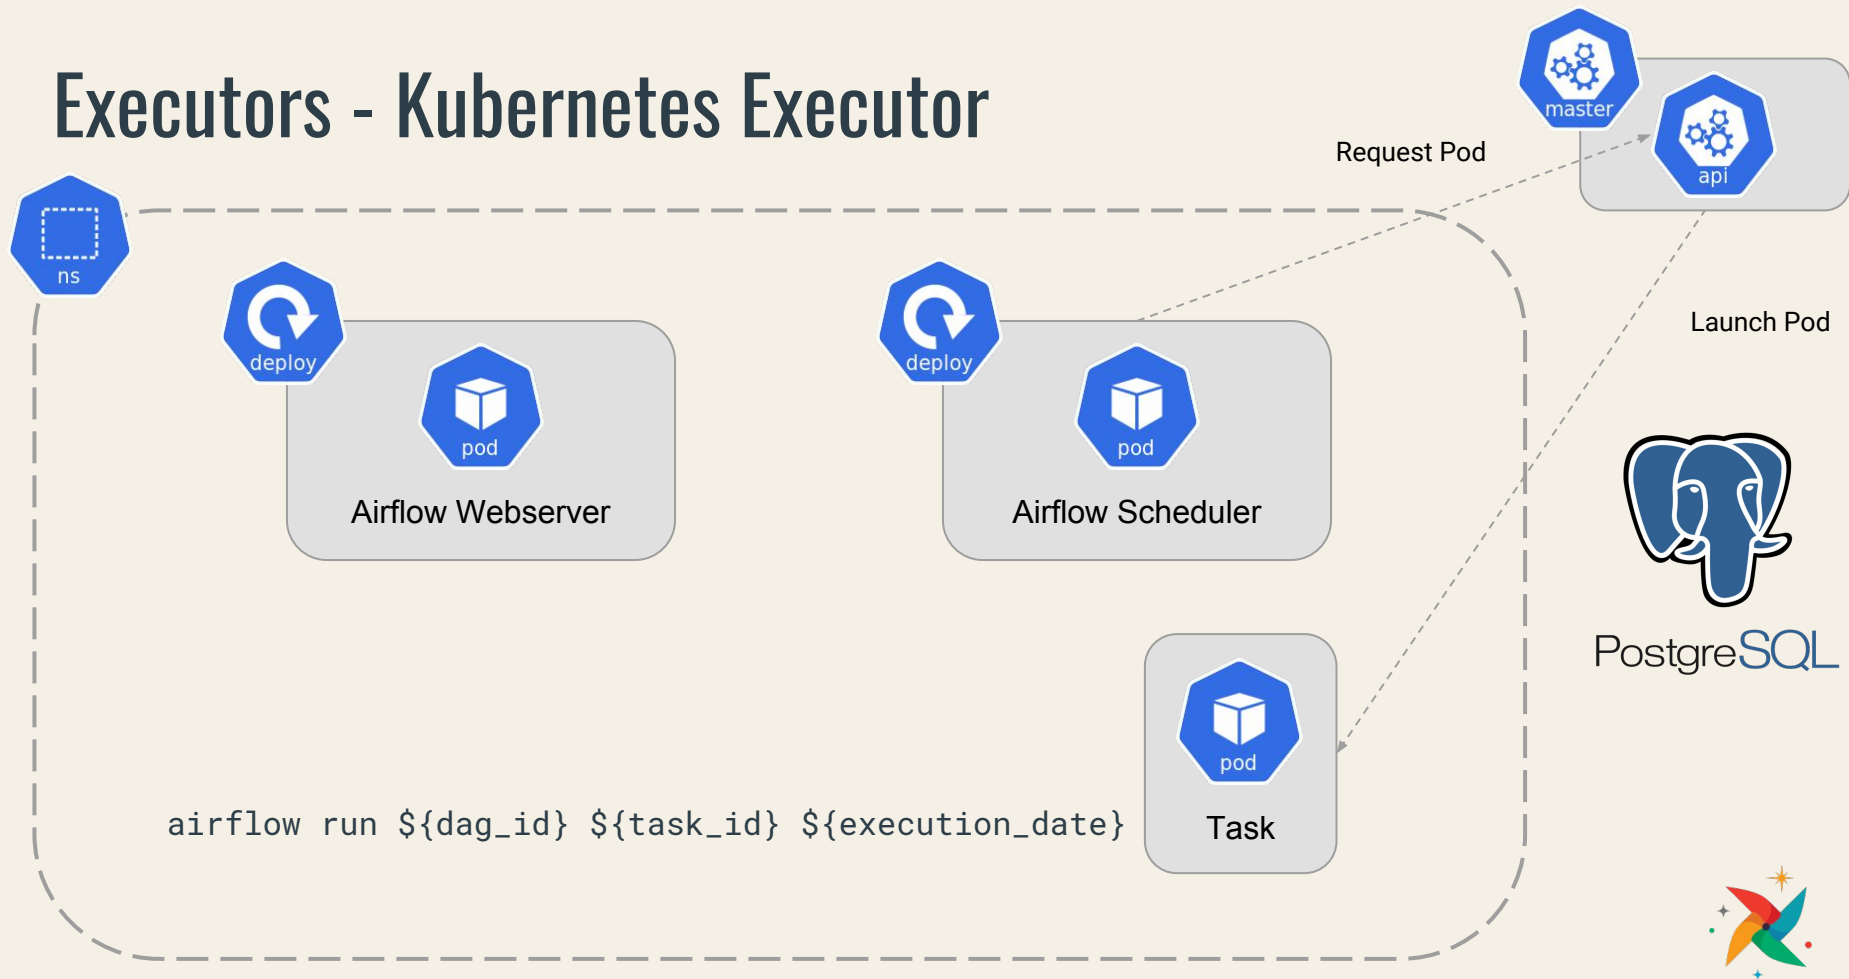
- Scale to zero / near-zero
- Each task runs in a new pod
  - Configurable resource requests (cpu/mem)
- Scheduler subscribes to Kubernetes event stream
- Pods run to completion
- Straightforward and natural
- DAG distribution
  - Git clone with init container for each pod
  - Mount volume with DAGs
  - Ensure the image already contains the DAG code



# Executors - Kubernetes Executor

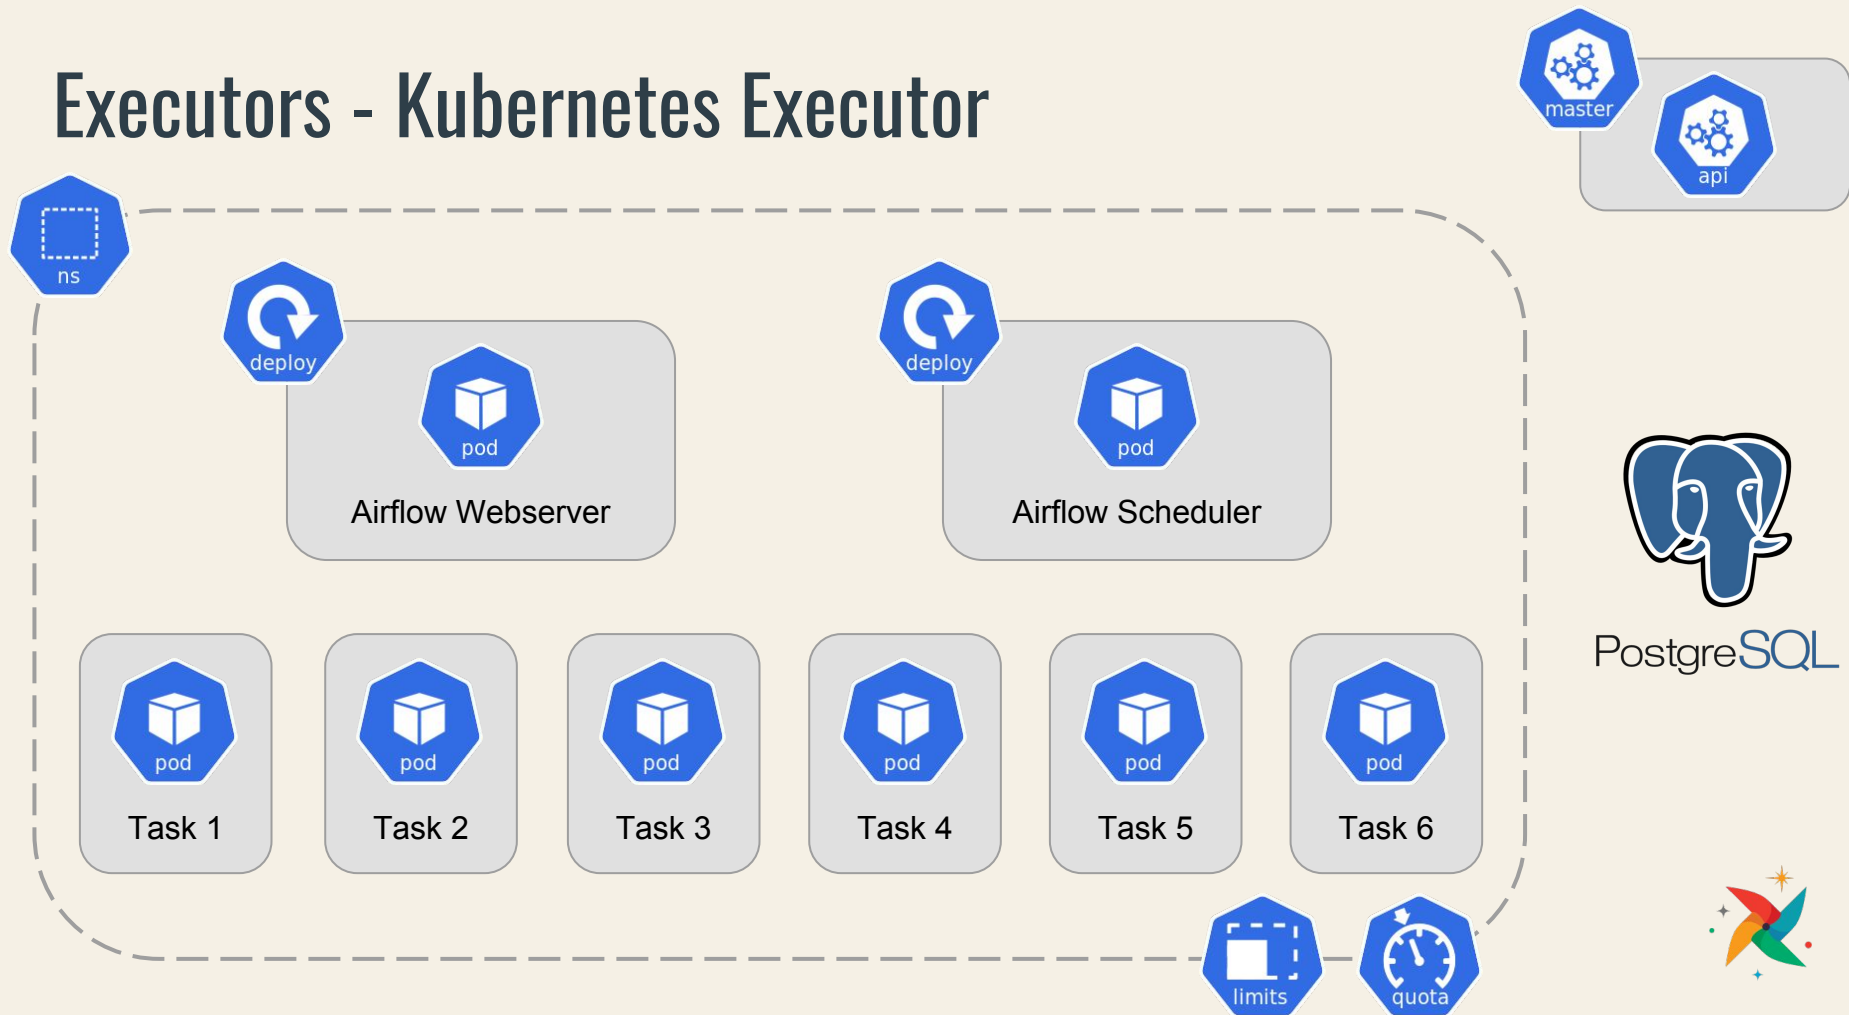


# Executors - Kubernetes Executor





# Executors - Kubernetes Executor



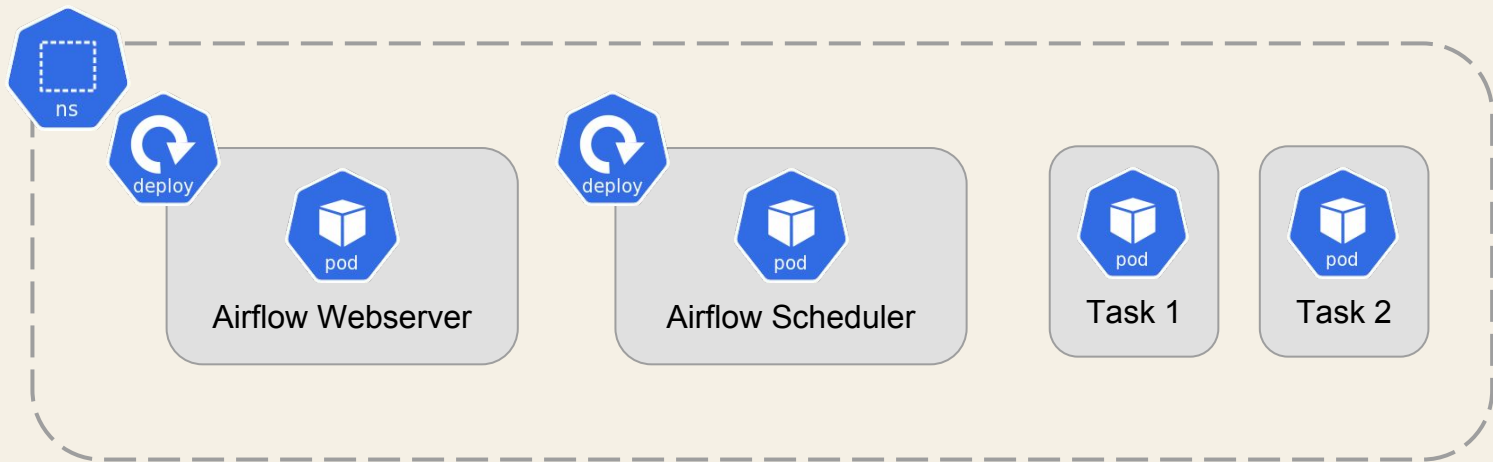
How do we deploy DAG updates to a running environment?



```
helm upgrade airflow-prod charts/airflow --set tag=v0.0.2
```



# DAG Updates



- `helm upgrade` updates the Deployments state in Kubernetes
- Kubernetes gracefully terminates the webserver and scheduler and reboots pods with updated image tag
- Task pods continue running to completion
- You experience negligible amount of downtime
- Can be automated via CI/CD tooling



How do we monitor and alert across a number of Airflow deployments?

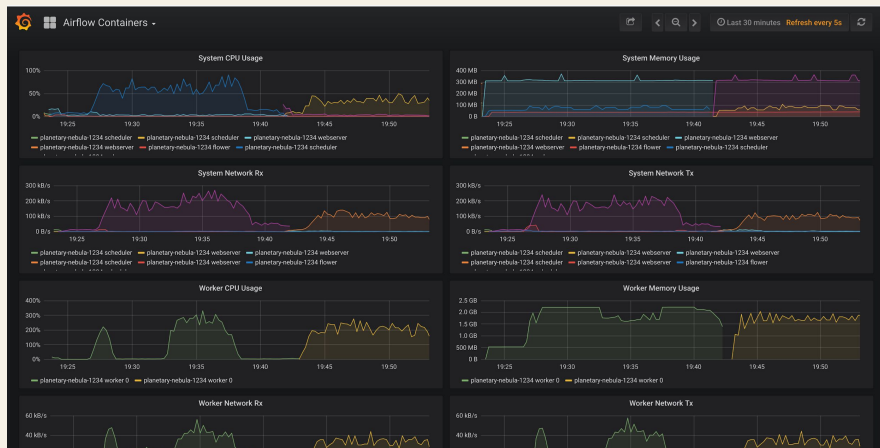


```
helm install stable/prometheus
```

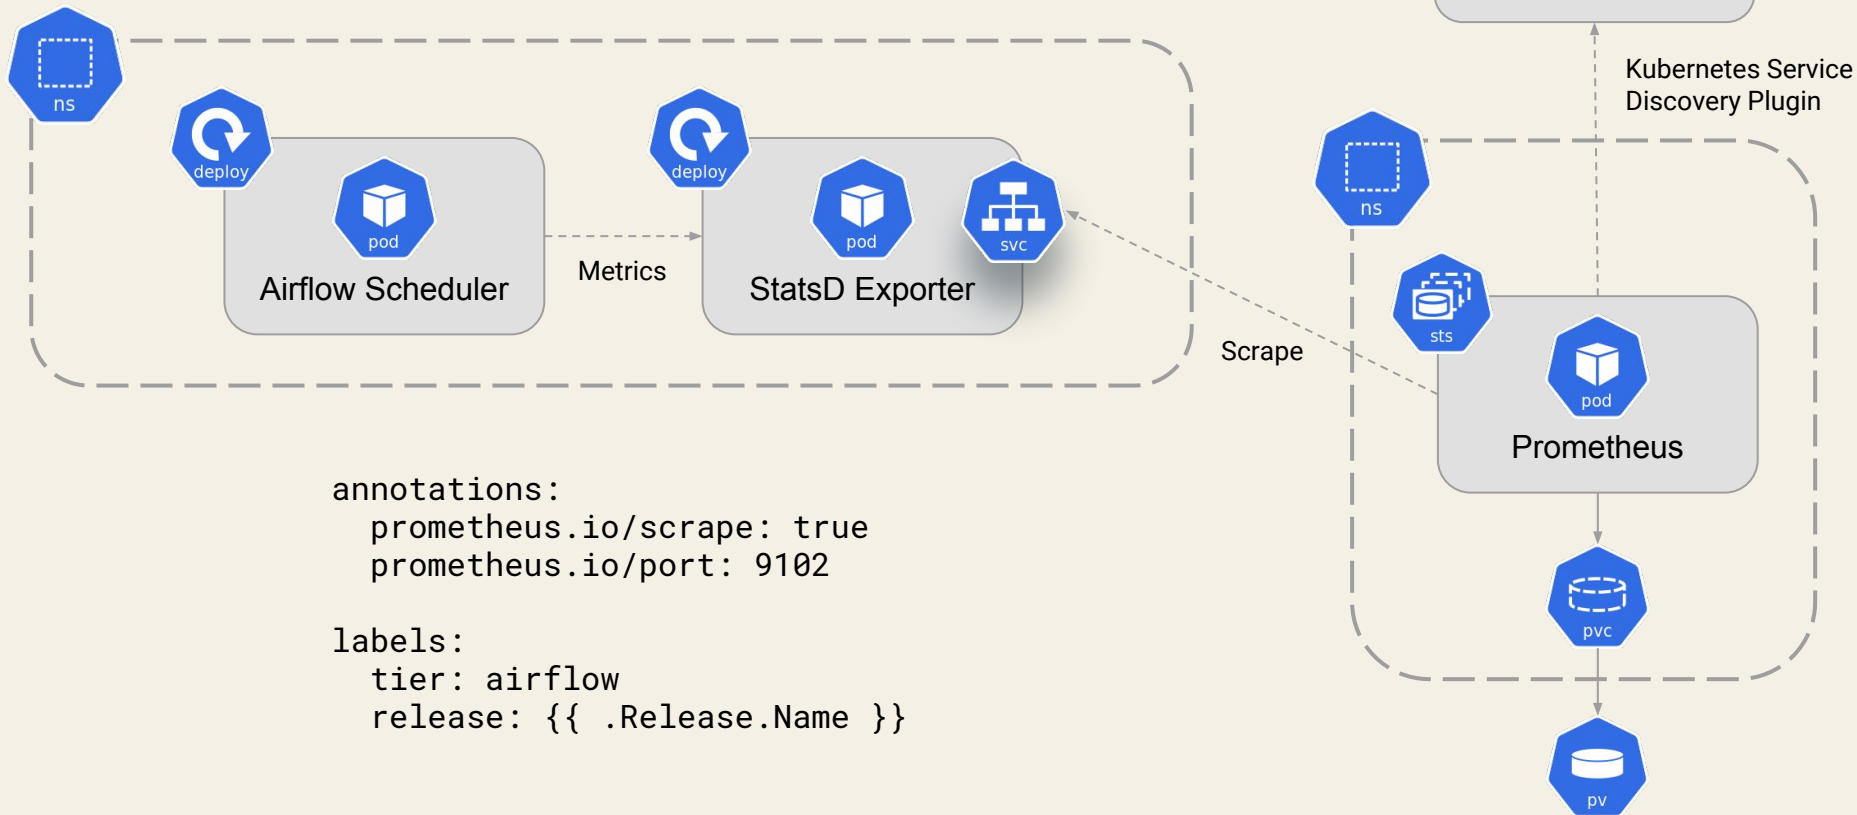


# Monitoring Airflow(s) with Prometheus

- Prometheus
  - Also CNCF project
  - Time series database
  - Pull-based
  - Auto-scrape with kubernetes annotations and SD plugin
  - Works great with Grafana →
- Airflow natively exports statsd metrics
- Statsd Exporter as a bridge to Prometheus

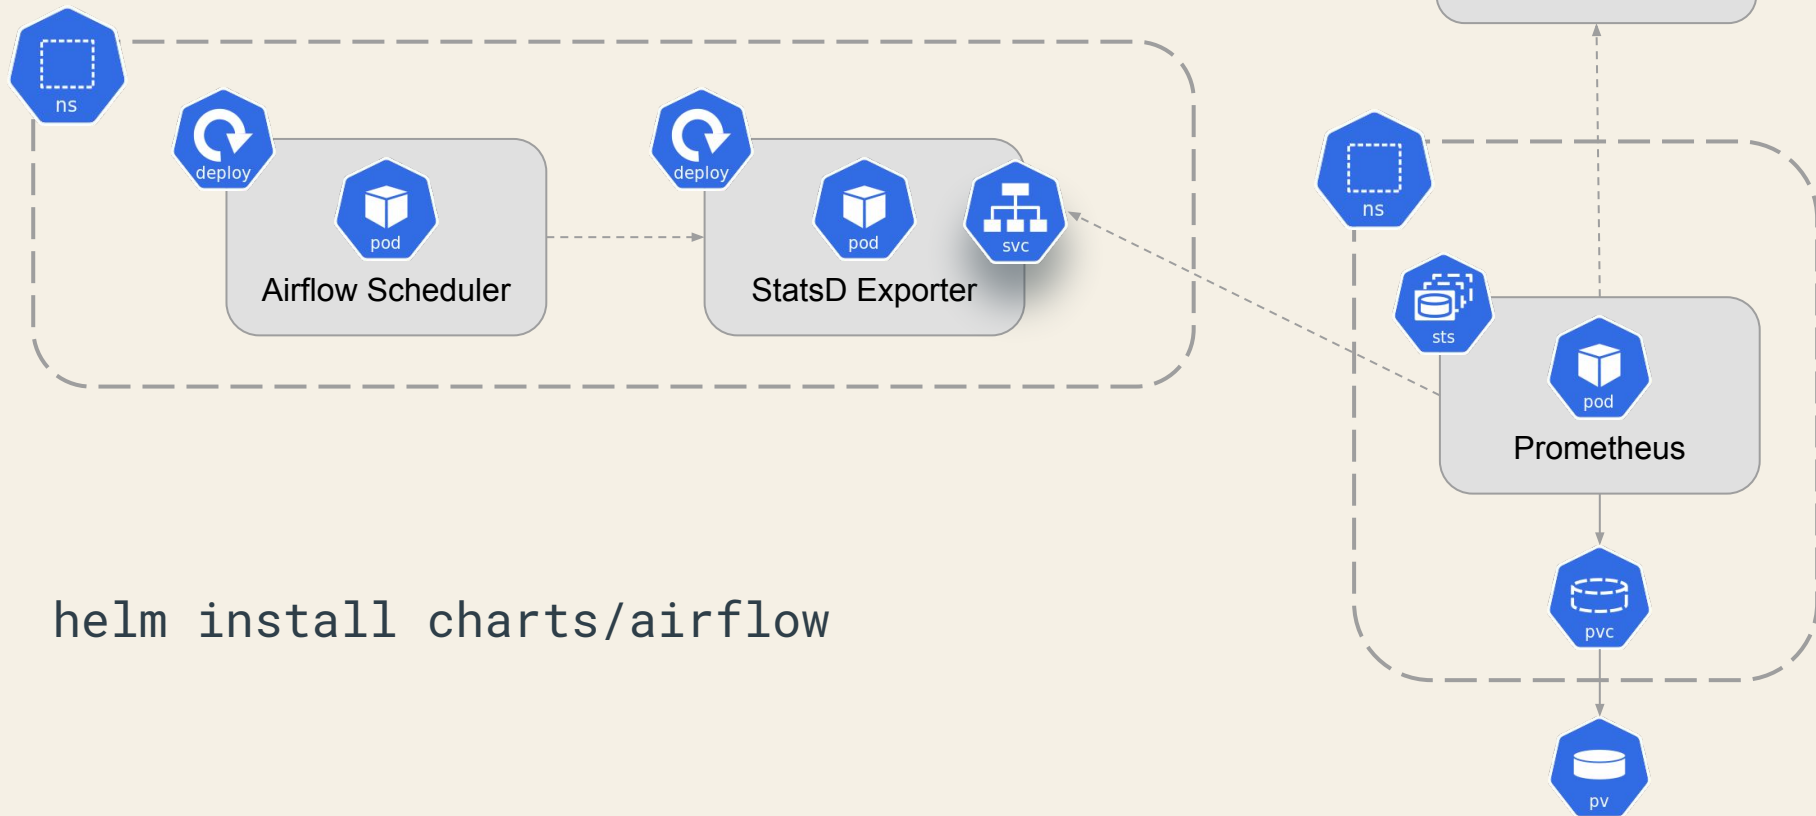


# Monitoring Airflow(s) with Prometheus

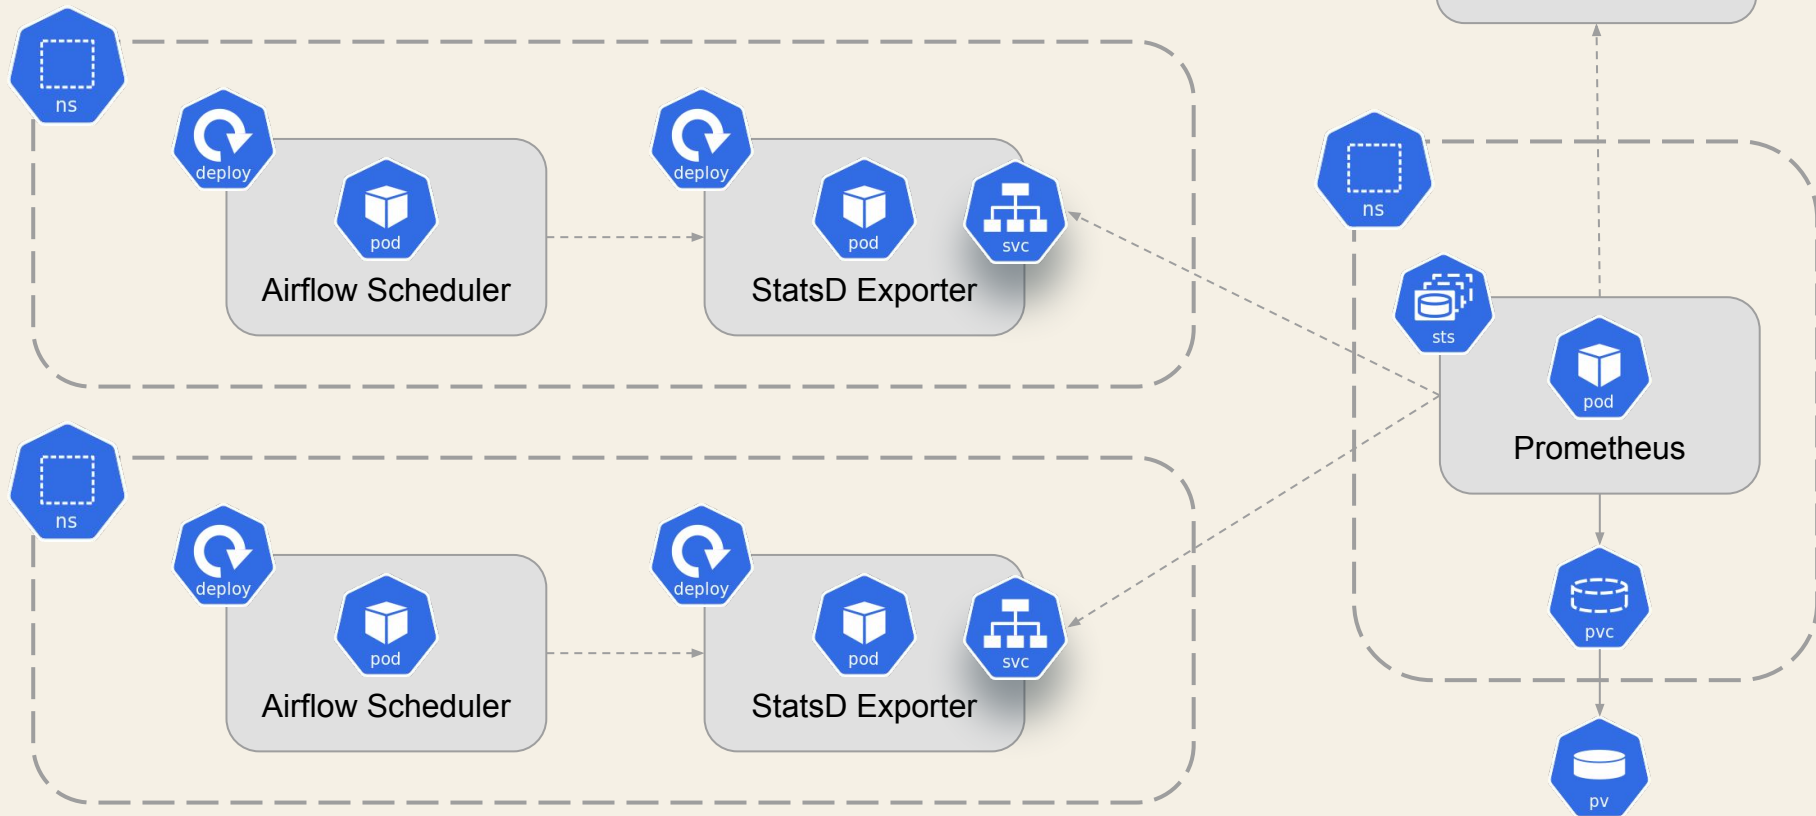




# Monitoring Airflow(s) with Prometheus

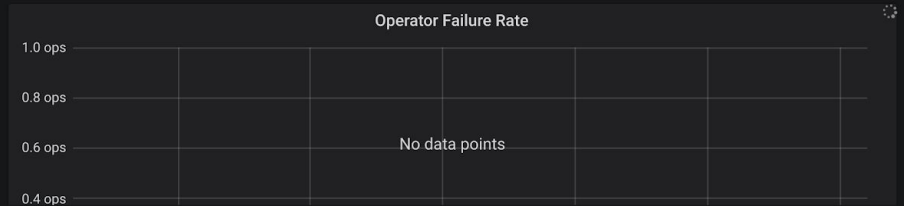
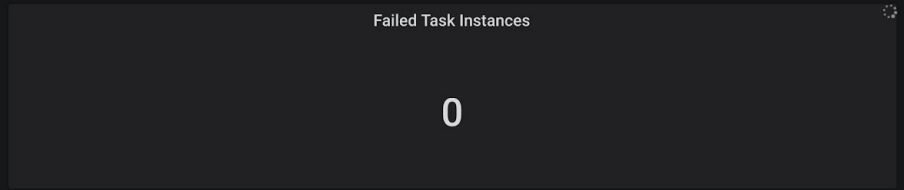
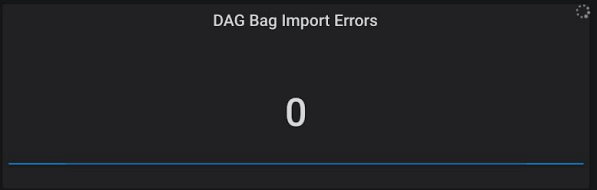
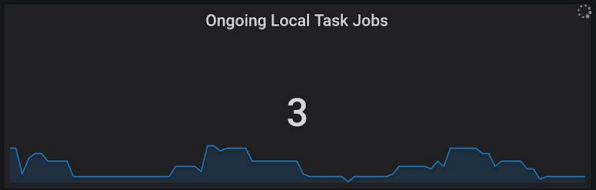
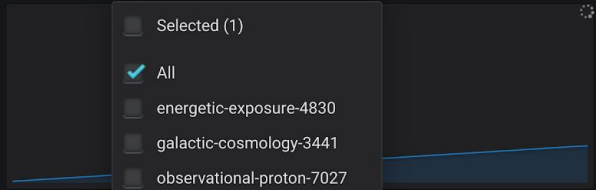


# Monitoring Airflow(s) with Prometheus



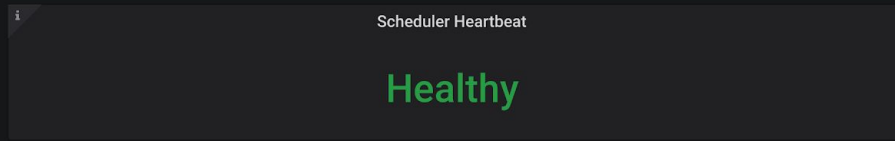
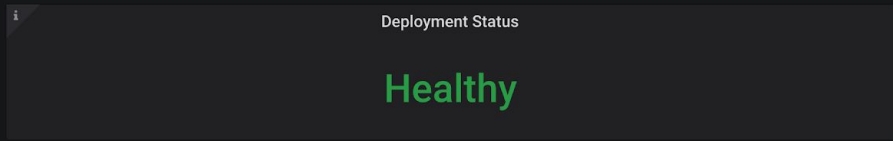
Deployment

☐ Selected (1)  
☒ All  
☐ energetic-exposure-4830  
☐ galactic-cosmology-3441  
☐ observational-proton-7027

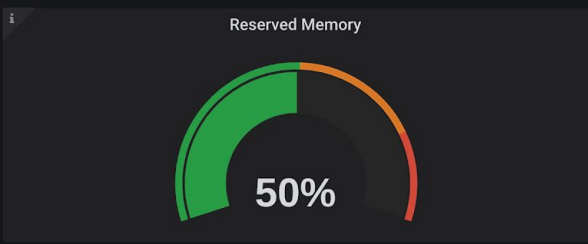
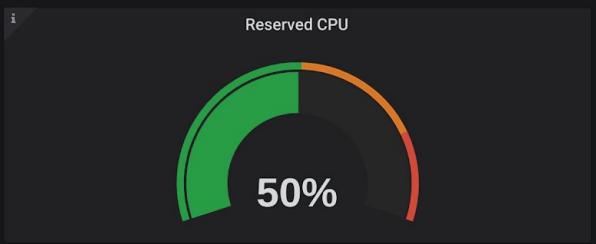
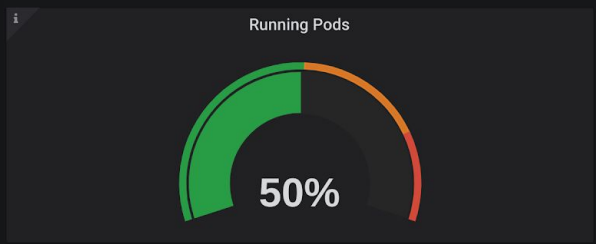
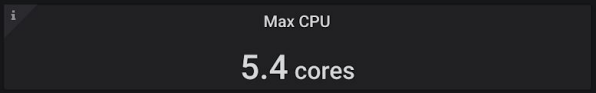
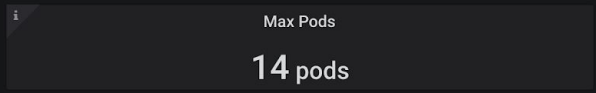


Deployment galactic-cosmology-3441

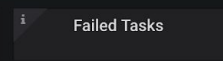
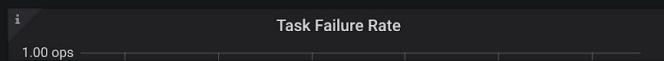
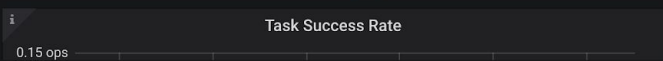
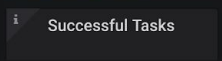
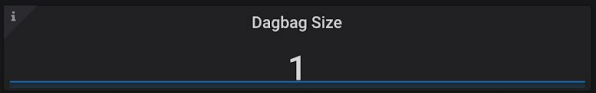
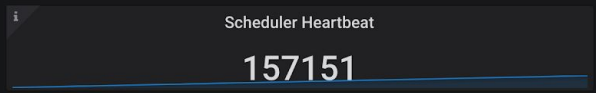
At a Glance



Quotas



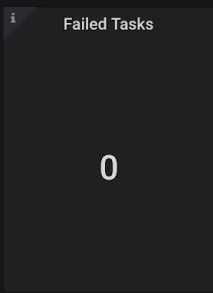
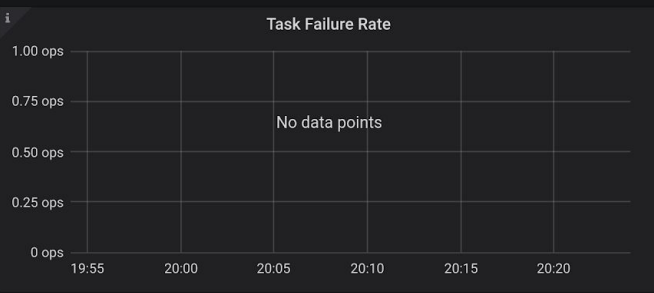
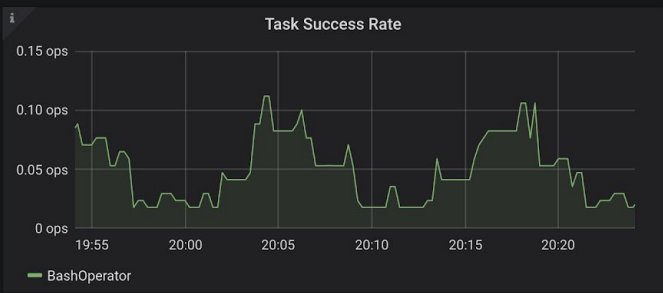
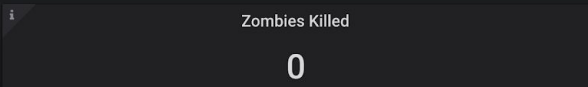
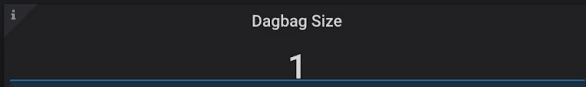
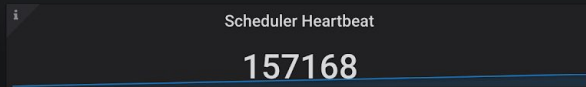
Scheduler



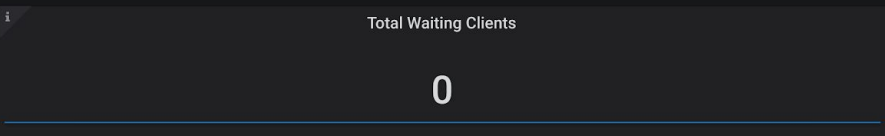
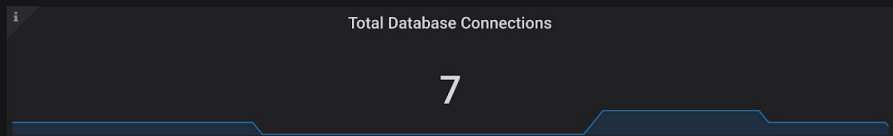


## Airflow Deployment Overview ▾

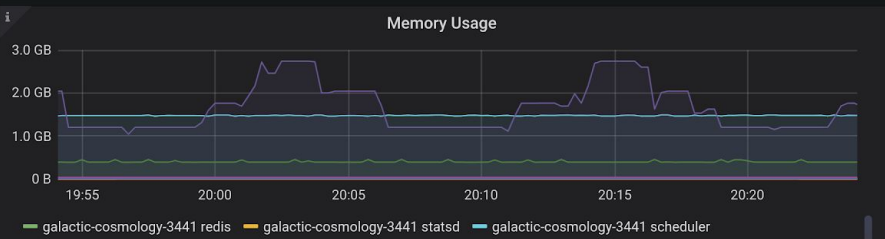
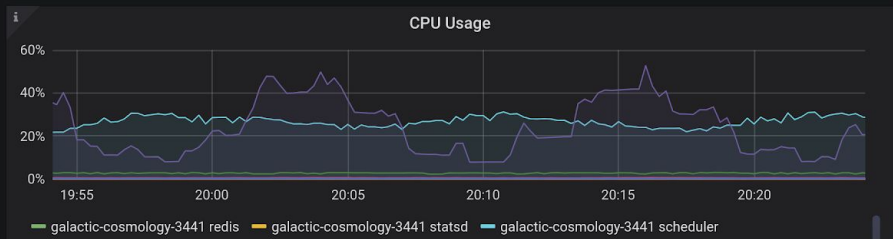
🔄 Last 30 minutes Refresh every 10s 🔍 ↺

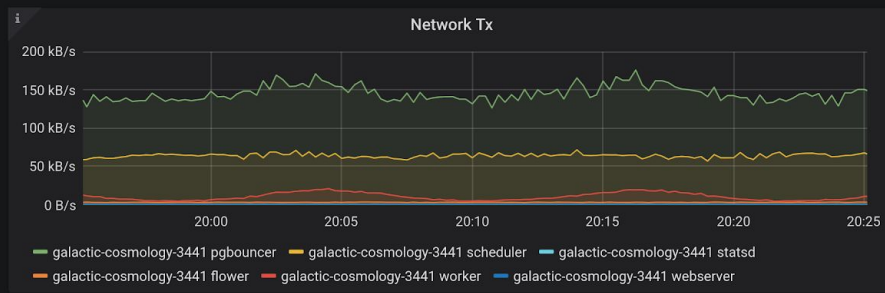
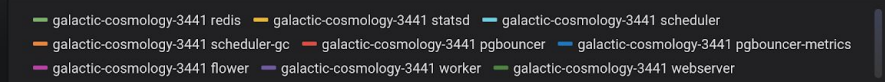
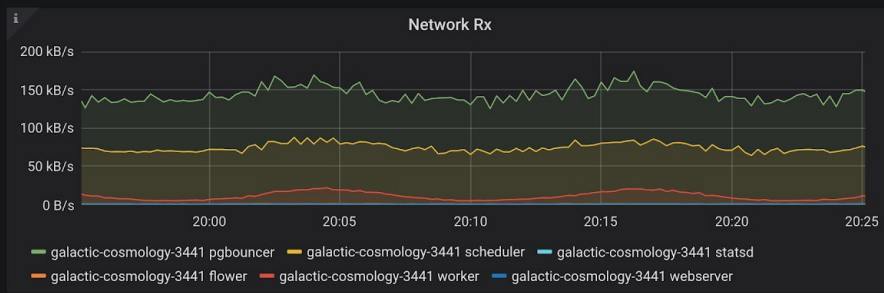
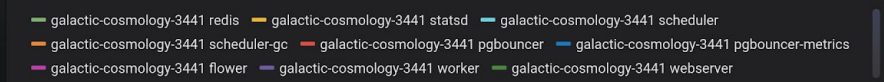


### Database Connections



### Resource Utilization



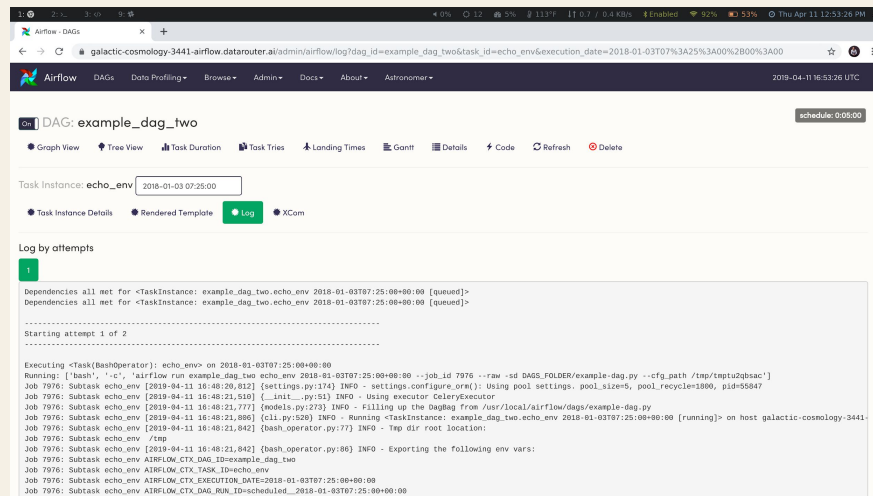


Container Status

Container Status		
Container	Pod	Status
galactic-cosmology-3441-redis	galactic-cosmology-3441-redis-0	Healthy
galactic-cosmology-3441-scheduler	galactic-cosmology-3441-scheduler-55fcf74956-bgggp	Healthy
galactic-cosmology-3441-pgbouncer	galactic-cosmology-3441-pgbouncer-55874674fc-fslcj	Healthy
galactic-cosmology-3441-statsd	galactic-cosmology-3441-statsd-7656dc98f5-9cng2	Healthy
galactic-cosmology-3441-flower	galactic-cosmology-3441-flower-7b7c9f8598-gcv7l	Healthy
galactic-cosmology-3441-webserver	galactic-cosmology-3441-webserver-dd858849d-knw7m	Healthy
galactic-cosmology-3441-pgbouncer-metrics	galactic-cosmology-3441-pgbouncer-55874674fc-fslcj	Healthy
galactic-cosmology-3441-scheduler-gc	galactic-cosmology-3441-scheduler-55fcf74956-bgggp	Healthy
galactic-cosmology-3441-worker	galactic-cosmology-3441-worker-85f6f4d96b-jld7z	Healthy

# Airflow Logging

- Powers the task log view in Airflow UI
- KubernetesExecutor requires remote logging plugin
- Several remote logging backend plugins available
  - Object Storage (S3, GCS, WASB)
  - Elasticsearch



The screenshot shows the Airflow web interface for a DAG named 'example\_dag\_two'. The URL is 'galactic-cosmology-3441-airflow.datarouter.ai/admin/airflow/log?dag\_id=example\_dag\_two&task\_id=echo\_env&execution\_date=2018-01-03T07:25:00%3A25%3A00%2800%3A00'. The interface includes a navigation bar with 'DAGs', 'Data Profiling', 'Browse', 'Admin', 'Docs', 'About', and 'Astronomer'. Below the navigation bar, the DAG name 'example\_dag\_two' is displayed with a 'schedule: 0:05:00' badge. The 'Log' tab is selected, showing the task instance 'echo\_env' for the execution date '2018-01-03 07:25:00'. The log content shows the task dependencies, starting attempt 1 of 2, and the execution of the task 'echo\_env' on host 'galactic-cosmology-3441-'. The log output includes the command 'airflow run example\_dag\_two echo\_env 2018-01-03T07:25:00+00:00 --job\_id 7976 --raw -sd DAGS\_FOLDER/example\_dag.py --cfg\_path /tmp/tmptu2qhsac' and the output 'Using pool settings. pool\_size=5, pool\_recycle=1800, pid=55847'.



☒ DAG: example\_dag\_two

schedule: 0:05:00

 Graph View
  Tree View
  Task Duration
  Task Tries
  Landing Times
  Gantt
  Details
  Code
  Refresh
  Delete

Task Instance: echo\_env 2018-01-03 07:25:00

[Task Instance Details](#)
[Rendered Template](#)
[Log](#)
[XCom](#)

## Log by attempts

```
Dependencies all met for <TaskInstance: example_dag_two.echo_env 2018-01-03T07:25:00+00:00 [queued]>
Dependencies all met for <TaskInstance: example_dag_two.echo_env 2018-01-03T07:25:00+00:00 [queued]>
```

Starting attempt 1 of 2

```
Executing <Task(BashOperator): echo_env> on 2018-01-03T07:25:00+00:00
```

```
Running: ['bash', '-c', 'airflow run example_dag_two echo_env 2018-01-03T07:25:00+00:00 --job_id 7976 --raw -sd DAGS_FOLDER/example-dag.py --cfg_path /tmp/tmptu2qbsac']
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:20,812] {settings.py:174} INFO - settings.configure_orm(): Using pool settings. pool_size=5, pool recycle=1800, pid=55847
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,510] {__init__.py:51} INFO - Using executor CeleryExecutor
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,777] {models.py:273} INFO - Filling up the DagBag from /usr/local/airflow/dags/example-dag.py
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,806] {cli.py:520} INFO - Running <TaskInstance: example_dag_two.echo_env 2018-01-03T07:25:00+00:00 [running]> on host galactic-cosmology-3441-
```

```
Job 7976: Subtask echo_env [2019-04-11 16:48:21,842] {bash_operator.py:77} INFO - Tmp dir root location:
```

```
Job 7976: Subtask echo env /tmp
```

```
Job 7976: Subtask echo env [2019-04-11 16:48:21,842] {bash operator.py:86} INFO - Exporting the following env vars:
```

```
Job 7976: Subtask echo env AIRFLOW CTX DAG ID=example dag two
```

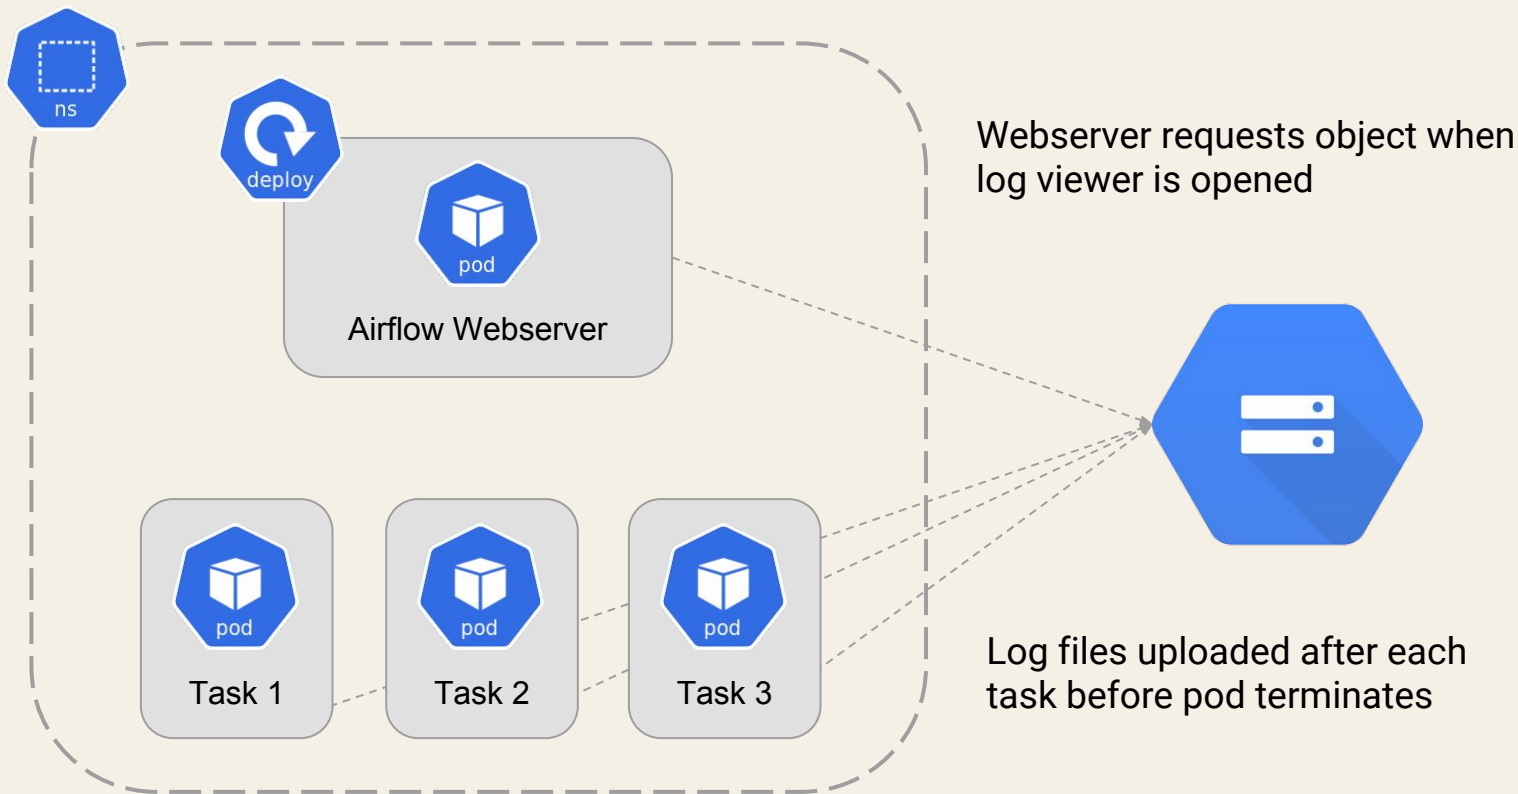
```
Job 7976: Subtask echo env ATRELOW CTX TASK ID=echo env
```

```
Job 7976: Subtask echo env AT/BELOW CTX EXECUTION DATE=2018-01-03T07:25:00+00:00
```

```
Job 7976: Subtask echo env ATRELOW CTX DAG RUN ID=scheduled 2018-01-03T07:25:00+00:00
```



# Airflow Logging - Object Storage



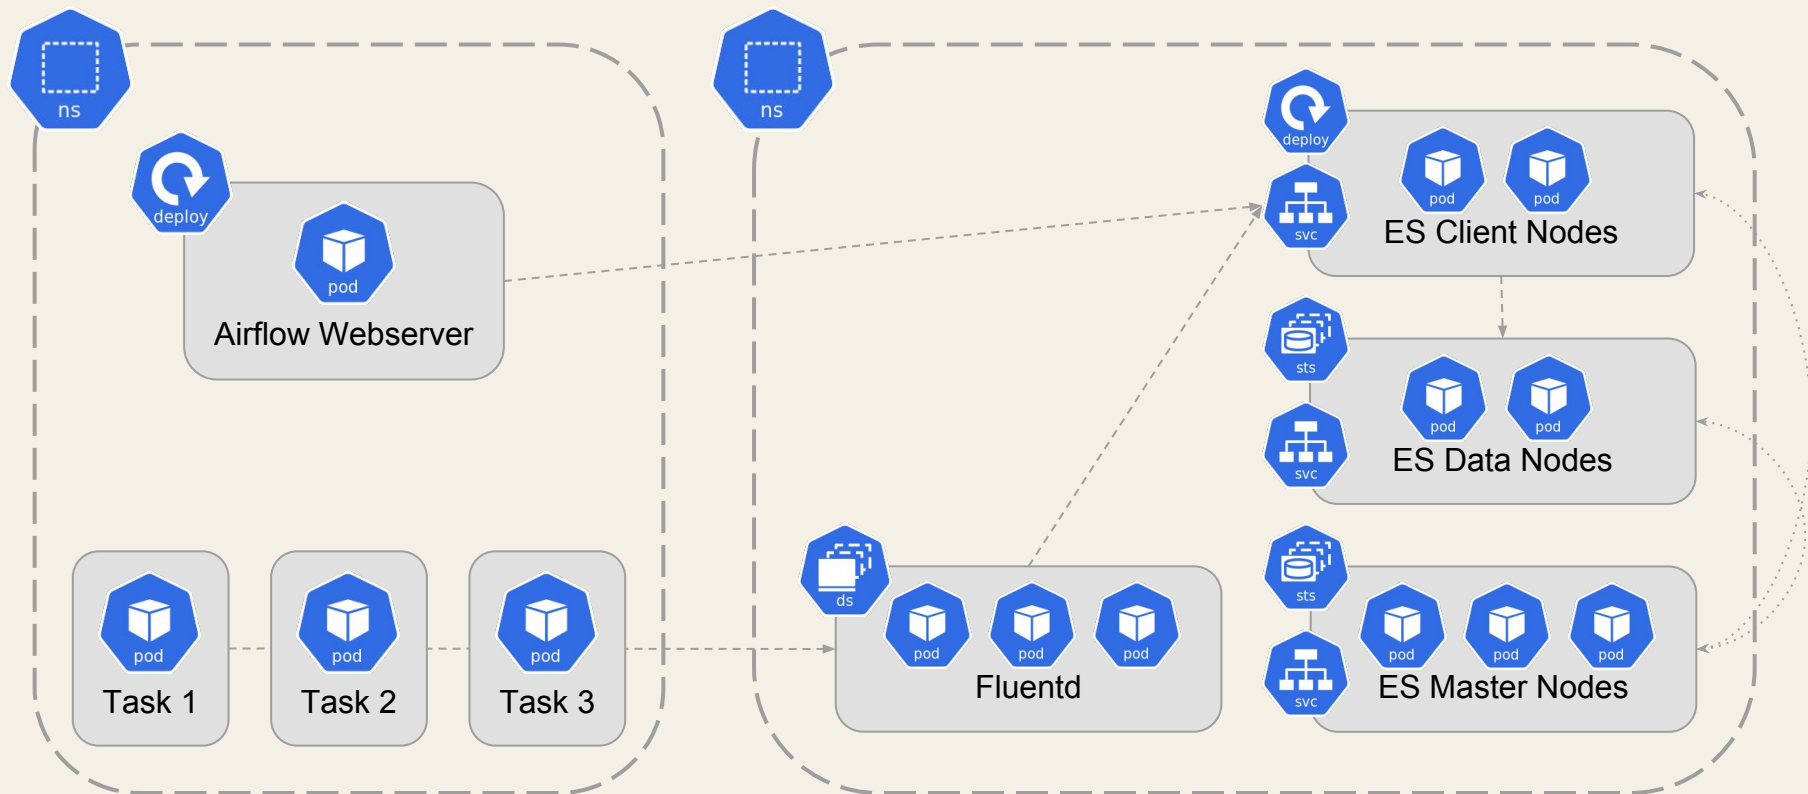
# Airflow Logging - Elasticsearch

```
helm install stable/elasticsearch
```

```
helm install stable/fluentd
```



# Airflow Logging - Elasticsearch



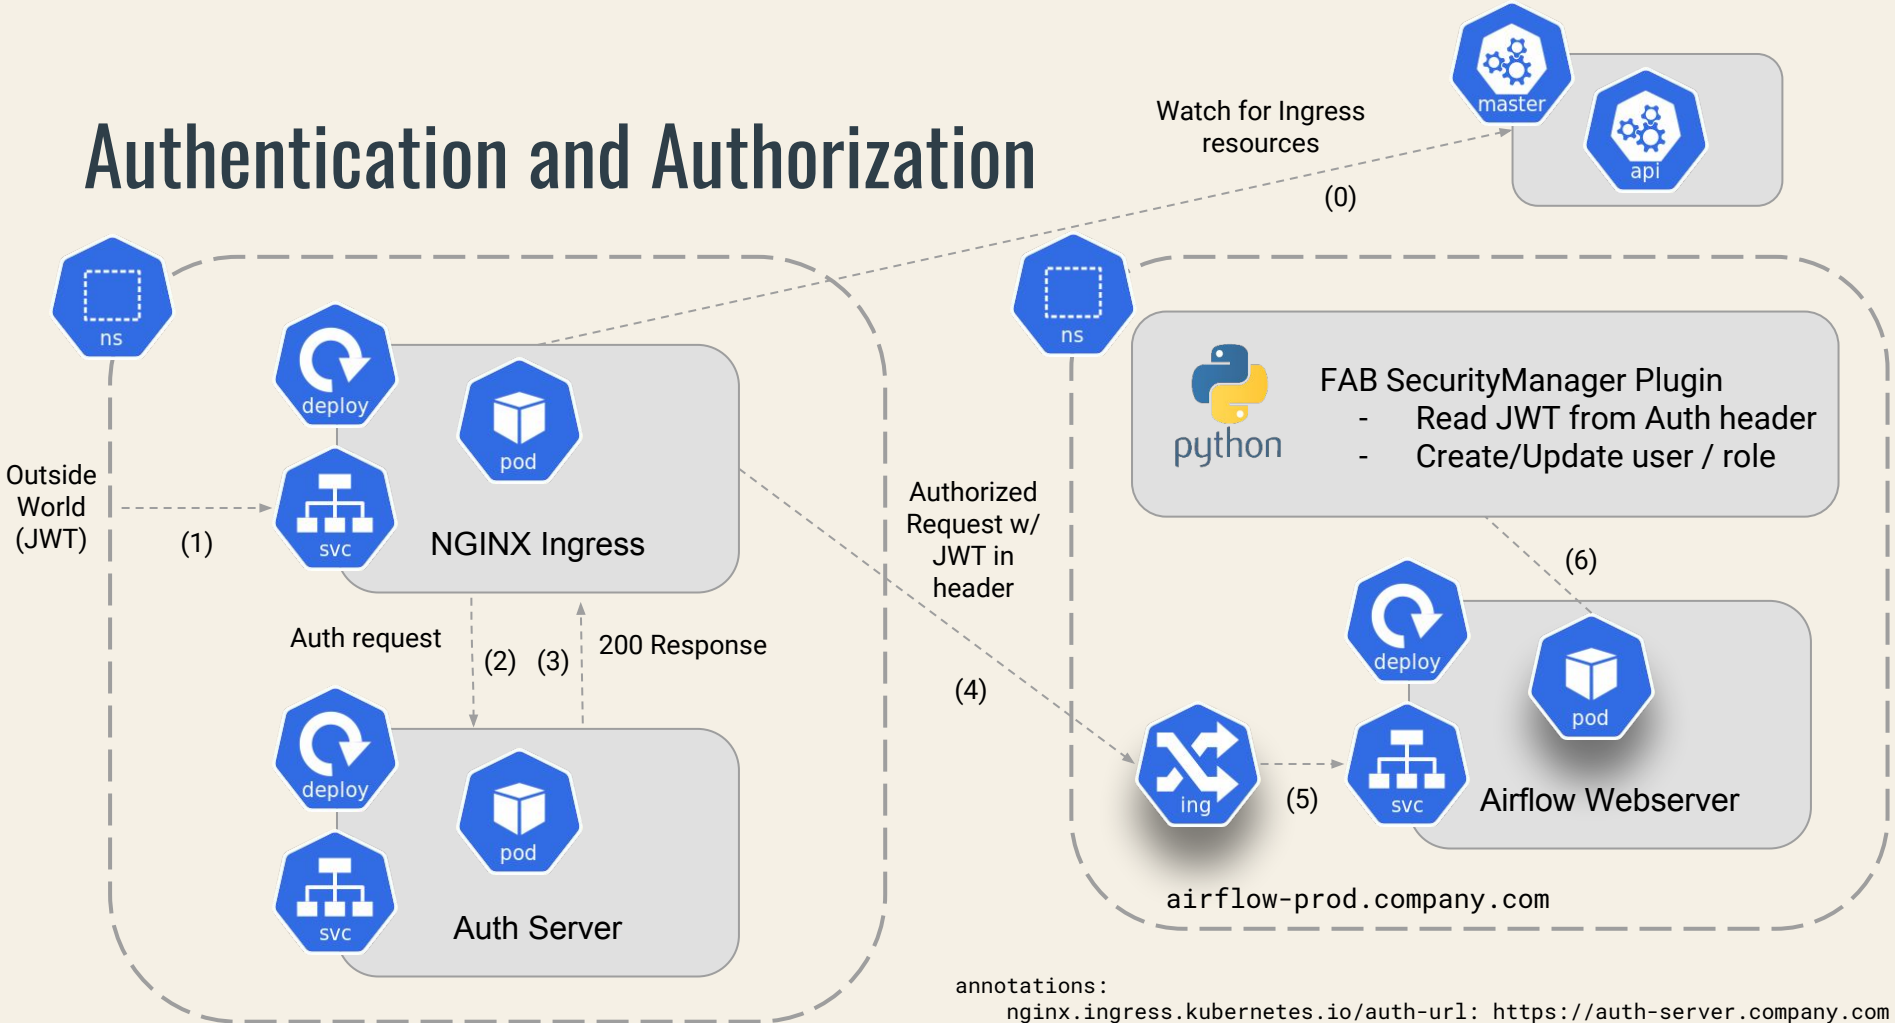
# Authentication and Authorization

- Ingress Controllers
  - Exposes a Kubernetes service to the outside world
  - Fulfills Kubernetes Ingress resources

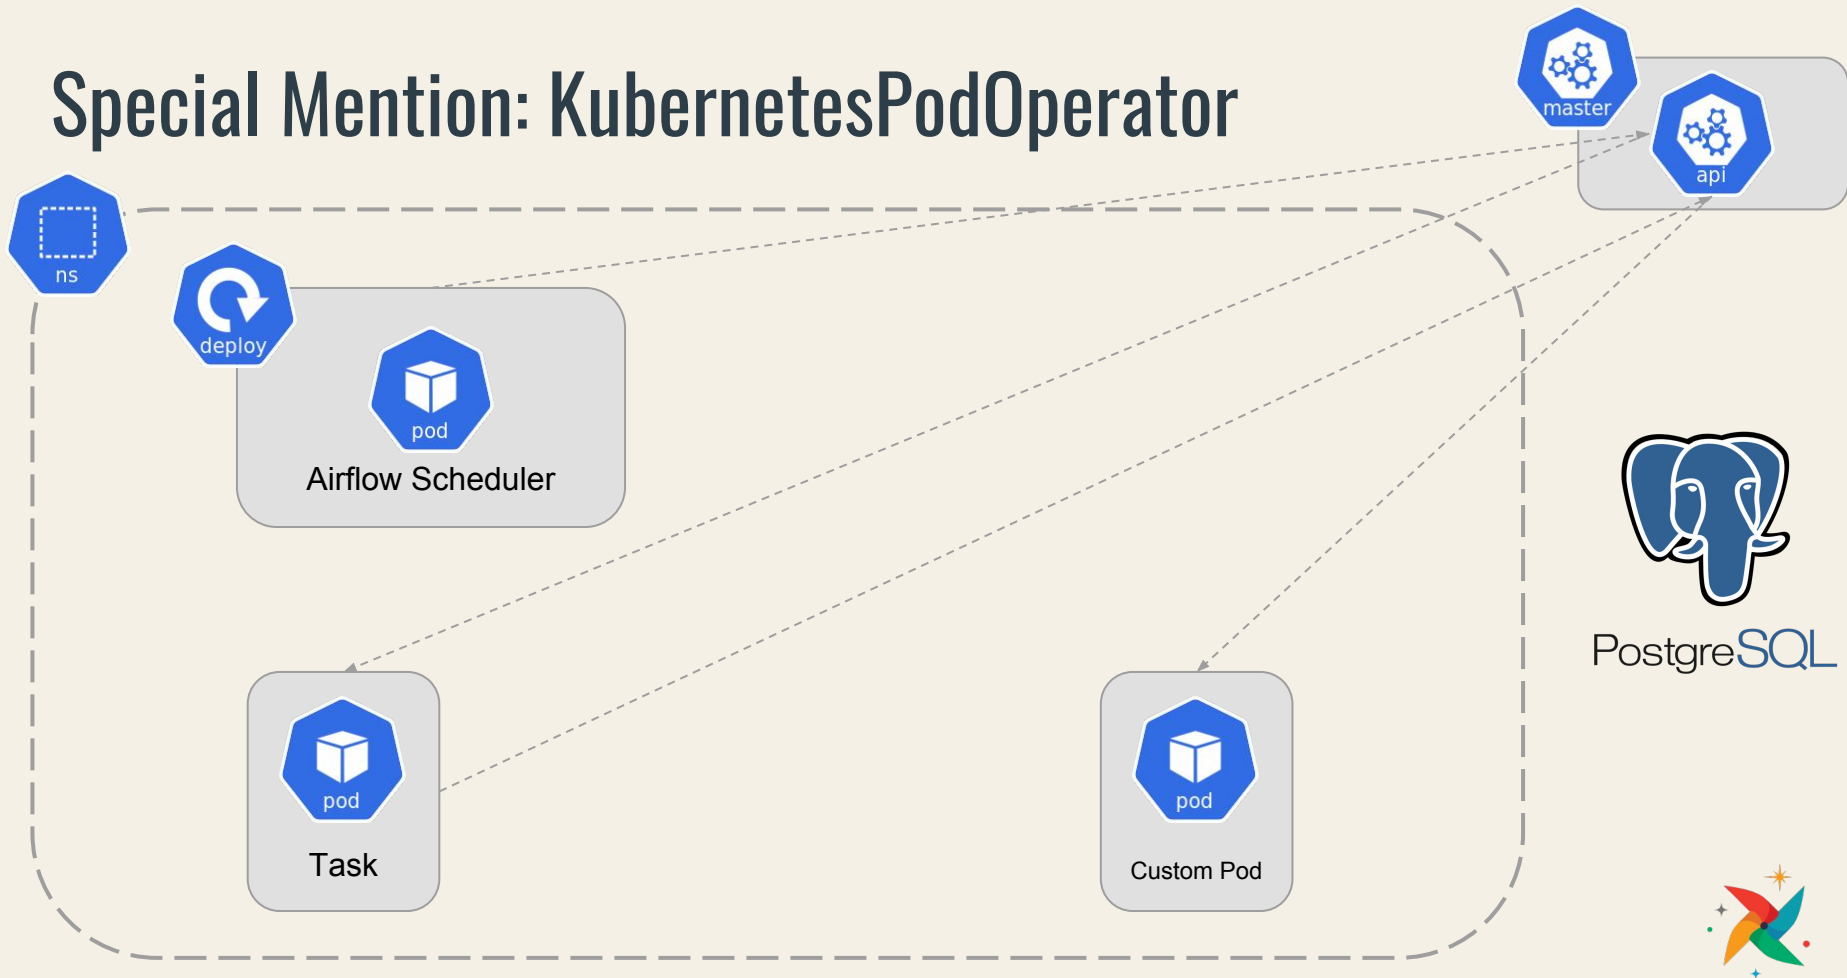
```
helm install stable/nginx-ingress
```



# Authentication and Authorization



# Special Mention: KubernetesPodOperator



Thank you!

greg@astronomer.io

