


```
breeze k8s create-cluster
```

Use “breeze k8s shell” & “echo \$KUBECONFIG” for search config location.

Configuration

Open <http://localhost:28080/connection/add> & configure k8s connection. For configuration use configuration from \$KUBECONFIG and replace 127.0.0.1 with `https://airflow-python-3.9-v1.28.13-control-plane:6443`

Edit Connection

Connection Id *	<input type="text" value="kubernetes_default"/>
Connection Type *	<input type="text" value="Kubernetes Cluster Connection"/> <small>Connection Type missing? Make sure you've installed the corresponding Airflow Provider Package.</small>
Description	<input type="text"/>
In cluster configuration	<input type="checkbox"/>
Kube config path	<input type="text" value="/tmp/.kubeconfig"/>
Kube config (JSON format)	<input type="text"/>
Namespace	<input type="text" value="default"/>
Cluster context	<input type="text"/>
Disable SSL	<input type="checkbox"/>
Disable TCP keepalive	<input type="checkbox"/>
XCom sidecar image	<input type="text"/>
XCom sidecar resources (JSON format)	<input type="text"/>

Create shared network

```
docker network create airflow_network
docker network connect airflow_network
airflow-python-3.9-v1.28.13-control-plane
docker network connect airflow_network
breeze-airflow-run-b923b5668db1
```

Upload dag

Create a dag as python file in "/files/dags" and restart scheduler service.

```
k8s_42498.py
```

```
from uuid import uuid4

from airflow import DAG
from airflow.providers.cncf.kubernetes.operators.pod import
KubernetesPodOperator
from airflow.utils.dates import days_ago
from kubernetes.client import models as k8s

default_args = {
    'owner': 'mandreev',
    'depends_on_past': False,
    'start_date': days_ago(1),
    'email': ['mark.andreev@gmail.com'],
    'email_on_failure': False,
    'email_on_retry': False,
    'max_active_runs': 1,
    'retries': 3
}

with DAG(
    "k8s_42498",
    schedule_interval=None,
    catchup=False,
) as dag:
    marker_from_init_container = f"{uuid4()}"
    marker_from_main_container = f"{uuid4()}"
    init_container = k8s.V1Container(
        name="init-container",
        image="busybox",
        command=["sh", "-cx"],
        args=[f"echo {marker_from_init_container}"],
    )
    k8s_pod = KubernetesPodOperator(
        task_id="k8s_42498",
        namespace="default",
        image="busybox",
        cmds=["sh", "-cx"],
        arguments=[f"echo {marker_from_main_container}"],
        in_cluster=False,
        do_xcom_push=False,
        startup_timeout_seconds=60,
        init_containers=[init_container],
        init_container_logs=True,
    )
```

k8s_pod

Testing

Positive case (init_container_logs=True)

Trigger dag http://localhost:28080/dags/k8s_42498/grid

The screenshot shows the Airflow web interface for the DAG 'k8s_42498'. The top navigation bar includes 'DAGs', 'Cluster Activity', 'Datasets', 'Security', 'Browse', 'Admin', and 'Docs'. The current time is 21:40 UTC. The DAG 'k8s_42498' is displayed with a 'Schedule: None' and 'Next Run ID: None'. A red circle highlights the 'Run' button (a play icon). Below the DAG name, there are filters for 'All Run Types' and 'All Run States', and a 'Clear Filters' button. A legend for run states is visible at the bottom, including 'deferred', 'failed', 'queued', 'removed', 'restarting', 'running', 'scheduled', 'shutdown', 'skipped', 'success', 'up_for_reschedule', 'up_for_retry', 'upstream_failed', and 'no_status'.

Open logs of latest run:

The screenshot shows the 'Logs' tab for the latest run of DAG 'k8s_42498' on 2024-12-28 at 21:40:45 UTC. The interface includes tabs for 'Details', 'Graph', 'Gantt', 'Code', 'Event Log', 'Logs', 'XCom', and 'Task Duration'. The 'Logs' tab is active, showing a list of log entries. A red box highlights the log entry for the 'init-container' task, which contains the command 'echo b89f48ea-7a48-4a21-8929-a9580aaaa83a'. The log entry shows the output of the command as 'b89f48ea-7a48-4a21-8929-a9580aaaa83a'. Other log entries show the pod building process and the task completion.

```
f09b4539474a
*** Found local files:
***   *
/root/airflow/logs/dag_id=k8s_42498/run_id=manual__2024-12-28T21:4
0:45.528266+00:00/task_id=k8s_42498/attempt=1.log
[2024-12-28, 21:40:51 UTC] {local_task_job_runner.py:123} ▶ Pre
task execution logs
[2024-12-28, 21:40:51 UTC] {pod.py:1148} INFO - Building pod
k8s-42498-8el28m7g with labels: {'dag_id': 'k8s_42498', 'task_id':
'k8s_42498', 'run_id':
'manual__2024-12-28T214045.5282660000-433a26fb2',
'kubernetes_pod_operator': 'True', 'try_number': '1'}
[2024-12-28, 21:40:51 UTC] {base.py:84} INFO - Retrieving
connection 'kubernetes_default'
```

```

[2024-12-28, 21:40:51 UTC] {pod.py:545} INFO - Found matching pod
k8s-42498-8e128m7g with labels {'airflow_kpo_in_cluster': 'False',
'airflow_version': '2.10.3', 'dag_id': 'k8s_42498',
'kubernetes_pod_operator': 'True', 'run_id':
'manual__2024-12-28T214045.5282660000-433a26fb2', 'task_id':
'k8s_42498', 'try_number': '1'}
[2024-12-28, 21:40:51 UTC] {pod.py:546} INFO - `try_number` of
task_instance: 1
[2024-12-28, 21:40:51 UTC] {pod.py:547} INFO - `try_number` of
pod: 1
[2024-12-28, 21:40:56 UTC] {pod_manager.py:477} INFO -
[init-container] + echo b89f48ea-7a48-4a21-8929-a9589aaaa83a
[2024-12-28, 21:40:56 UTC] {pod_manager.py:496} INFO -
[init-container] b89f48ea-7a48-4a21-8929-a9589aaaa83a
[2024-12-28, 21:40:56 UTC] {pod_manager.py:477} INFO - [base] +
echo 7d0b3030-6151-45b1-a9d1-33bd00ea3aca
[2024-12-28, 21:40:56 UTC] {pod_manager.py:496} INFO - [base]
7d0b3030-6151-45b1-a9d1-33bd00ea3aca
[2024-12-28, 21:40:56 UTC] {pod_manager.py:524} WARNING - Pod
k8s-42498-8e128m7g log read interrupted but container base still
running. Logs generated in the last one second might get
duplicated.
[2024-12-28, 21:40:57 UTC] {pod_manager.py:477} INFO - [base] +
echo 7d0b3030-6151-45b1-a9d1-33bd00ea3aca
[2024-12-28, 21:40:57 UTC] {pod_manager.py:496} INFO - [base]
7d0b3030-6151-45b1-a9d1-33bd00ea3aca
[2024-12-28, 21:40:57 UTC] {pod_manager.py:654} INFO - Pod
k8s-42498-8e128m7g has phase Running
[2024-12-28, 21:40:59 UTC] {pod.py:997} INFO - Deleting pod:
k8s-42498-8e128m7g
[2024-12-28, 21:40:59 UTC] {taskinstance.py:340} ▶ Post task
execution logs

```

Negative case (without init_container_logs=True)

The screenshot shows the Airflow web interface for a task named 'k8s_42498'. The interface includes a top navigation bar with 'Clear task', 'Mark state as...', and 'Filter DAG by task' buttons. Below the navigation bar, there are tabs for 'Details', 'Graph', 'Gantt', 'Code', 'Event Log', 'Logs', 'XCom', and 'Task Duration'. The 'Logs' tab is selected, and the logs are displayed in a scrollable area. The logs show the task's execution details, including the pod's creation, the container's initialization, and the task's completion. The logs are filtered to show only the 'k8s_42498' task.

```

f09b4539474a
*** Found local files:
*** /root/airflow/logs/dag_id=k8s_42498/run_id>manual__2024-12-28T21:43:55.168561+00:00/task_id=k8s_42498/attempt=1.log
[2024-12-28, 21:44:01 UTC] {local_task_job_runner.py:123} ▶ Pre task execution logs
[2024-12-28, 21:44:02 UTC] {pod.py:1148} INFO - Building pod k8s-42498-jpoxsknv with labels: {'dag_id': 'k8s_42498', 'task_id': 'k8s_42498', 'run_id': 'manual__2024-12-28T214355.1685610000-83d0d7040',
[2024-12-28, 21:44:02 UTC] {base.py:84} INFO - Retrieving connection 'kubernetes.default'
[2024-12-28, 21:44:02 UTC] {pod.py:545} INFO - Found matching pod k8s-42498-jpoxsknv with labels {'airflow_kpo_in_cluster': 'False', 'airflow_version': '2.10.3', 'dag_id': 'k8s_42498', 'kubernetes_pod_o
[2024-12-28, 21:44:02 UTC] {pod.py:546} INFO - `try_number` of task_instance: 1
[2024-12-28, 21:44:02 UTC] {pod.py:547} INFO - `try_number` of pod: 1
[2024-12-28, 21:44:02 UTC] {pod_manager.py:392} WARNING - Pod not yet started: k8s-42498-jpoxsknv
[2024-12-28, 21:44:07 UTC] {pod_manager.py:477} INFO - [base] + echo 88369997-683b-4d08-9904-59d6f63a465c
[2024-12-28, 21:44:07 UTC] {pod_manager.py:496} INFO - [base] 88369997-683b-4d08-9904-59d6f63a465c
[2024-12-28, 21:44:07 UTC] {pod_manager.py:654} INFO - Pod k8s-42498-jpoxsknv has phase Running
[2024-12-28, 21:44:09 UTC] {pod.py:997} INFO - Deleting pod: k8s-42498-jpoxsknv
[2024-12-28, 21:44:09 UTC] {taskinstance.py:340} ▶ Post task execution logs

```

```
from uuid import uuid4
```

```

from airflow import DAG
from airflow.providers.cncf.kubernetes.operators.pod import
KubernetesPodOperator
from airflow.utils.dates import days_ago
from kubernetes.client import models as k8s

default_args = {
    'owner': 'mandreev',
    'depends_on_past': False,
    'start_date': days_ago(1),
    'email': ['mark.andreev@gmail.com'],
    'email_on_failure': False,
    'email_on_retry': False,
    'max_active_runs': 1,
    'retries': 3
}

with DAG(
    "k8s_42498",
    schedule_interval=None,
    catchup=False,
) as dag:
    marker_from_init_container = f"{uuid4()}"
    marker_from_main_container = f"{uuid4()}"
    init_container = k8s.V1Container(
        name="init-container",
        image="busybox",
        command=["sh", "-cx"],
        args=[f"echo {marker_from_init_container}"],
    )
    k8s_pod = KubernetesPodOperator(
        task_id="k8s_42498",
        namespace="default",
        image="busybox",
        cmds=["sh", "-cx"],
        arguments=[f"echo {marker_from_main_container}"],
        in_cluster=False,
        do_xcom_push=False,
        startup_timeout_seconds=60,
        init_containers=[init_container],
        #init_container_logs=True,
    )

    k8s_pod

```

f09b4539474a

*** Found local files:

*** *

/root/airflow/logs/dag_id=k8s_42498/run_id=manual__2024-12-28T21:43:55.168561+00:00/task_id=k8s_42498/attempt=1.log

[2024-12-28, 21:44:01 UTC] {local_task_job_runner.py:123} ▶ Pre

```
task execution logs
[2024-12-28, 21:44:02 UTC] {pod.py:1148} INFO - Building pod
k8s-42498-jpoxsknv with labels: {'dag_id': 'k8s_42498', 'task_id':
'k8s_42498', 'run_id':
'manual__2024-12-28T214355.1685610000-83d0d7040',
'kubernetes_pod_operator': 'True', 'try_number': '1'}
[2024-12-28, 21:44:02 UTC] {base.py:84} INFO - Retrieving
connection 'kubernetes_default'
[2024-12-28, 21:44:02 UTC] {pod.py:545} INFO - Found matching pod
k8s-42498-jpoxsknv with labels {'airflow_kpo_in_cluster': 'False',
'airflow_version': '2.10.3', 'dag_id': 'k8s_42498',
'kubernetes_pod_operator': 'True', 'run_id':
'manual__2024-12-28T214355.1685610000-83d0d7040', 'task_id':
'k8s_42498', 'try_number': '1'}
[2024-12-28, 21:44:02 UTC] {pod.py:546} INFO - `try_number` of
task_instance: 1
[2024-12-28, 21:44:02 UTC] {pod.py:547} INFO - `try_number` of
pod: 1
[2024-12-28, 21:44:02 UTC] {pod_manager.py:392} WARNING - Pod not
yet started: k8s-42498-jpoxsknv
[2024-12-28, 21:44:07 UTC] {pod_manager.py:477} INFO - [base] +
echo 88369997-683b-4d08-9904-59d6f63a465c
[2024-12-28, 21:44:07 UTC] {pod_manager.py:496} INFO - [base]
88369997-683b-4d08-9904-59d6f63a465c
[2024-12-28, 21:44:07 UTC] {pod_manager.py:654} INFO - Pod
k8s-42498-jpoxsknv has phase Running
[2024-12-28, 21:44:09 UTC] {pod.py:997} INFO - Deleting pod:
k8s-42498-jpoxsknv
[2024-12-28, 21:44:09 UTC] {taskinstance.py:340} ► Post task
execution logs
```