

Interactive Streaming Infra with JupyterHub and Apache Flink on Kubernetes

Elkhan Dadashov, Software Engineer, Apple



Agenda

- Streaming adoption challenges
 - Steep learning curve
 - Infrastructure
 - Language
- Interactive Streaming Infra
 - JupyterHub
 - Single pod development
 - At large scale development



Steep learning curve

- Challenges:
 - Bounded vs unbounded data (reading large Iceberg tables as a source)
 - Different time characteristics
 - Out-of-order data
- Solution:
 - Lots of examples with different connectors



Steep learning curve

- Challenges:
 - Connector specific knowledge (Kafka idleness, Iceberg commits)
- Solution:
 - SDK & connectors (Kafka, Schema Registry, Iceberg)
 - Good defaults
 - Extra documentation



Steep learning curve

- Challenges:
 - Too many knobs to tune for Flink pipeline (memory, task slots, Python, Java, Flink SQL)
- Solution:
 - Flink Kubernetes Autoscaler
 - Sources: Kafka source, Iceberg Source
 - Tuning parallelism & memory

Infrastructure challenges



Running on Clouds

- Challenge:
 - Handling pod/task failures
- Solution:
 - Flink Kubernetes Operator



Running on Clouds

- Challenge:
 - Object storage limits on PUTs
- Solution:
 - Cooperation with vendors on entropy-based partitioning



Running on Clouds

- Challenge:
 - Network costs
- Solution:
 - Co-locating Kafka and Flink in the same region/AZ

Language challenges

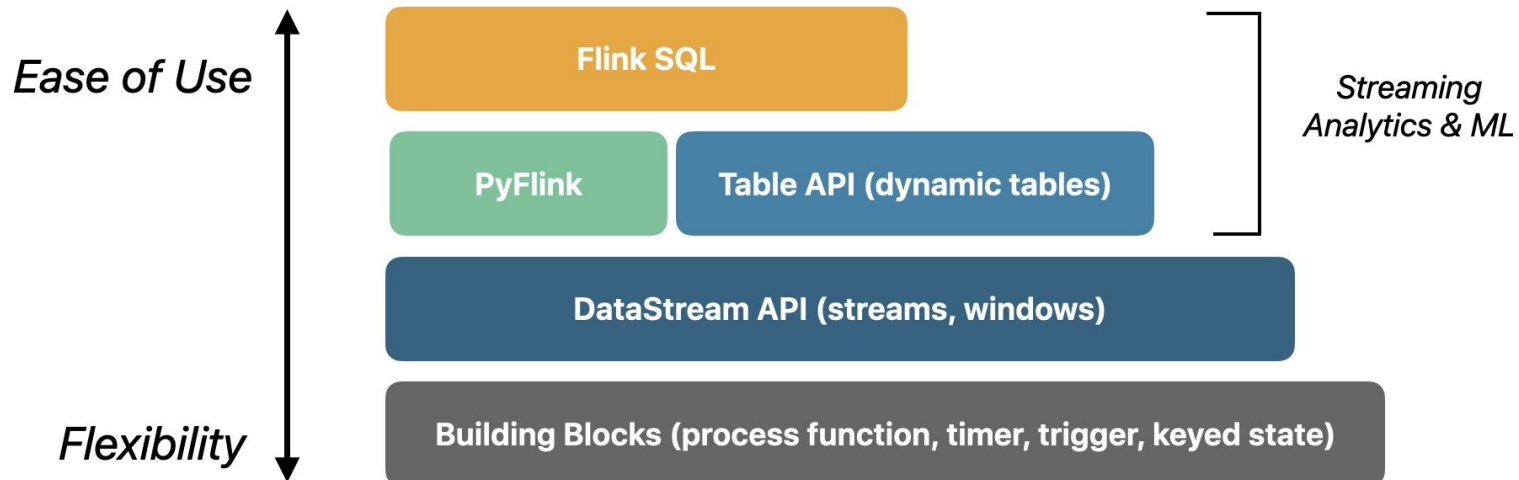


Language Challenges

- Challenge:
 - ML and Data science people use Python and SQL
 - Data engineers use Java/Scala
- Solution:
 - Support of different language APIs: Java, Flink SQL, PyFlink



Flink APIS for Application Development





Flink SQL

- Easier to use: what vs how
- NoCode: Well-known SQL standard: ANSI SQL + Temporal operators
- No additional cost for Infra or users
- Attract new users: ML/DS engineers/analysts



PyFlink

- PyFlink fully supports both the Table API and the DataStream API.
- Simplifies Flink SQL/Flink pipeline testing
- Python is one of the most popular language in data science and machine learning
- Pandas (powerful Python module) can be used in PyFlink also

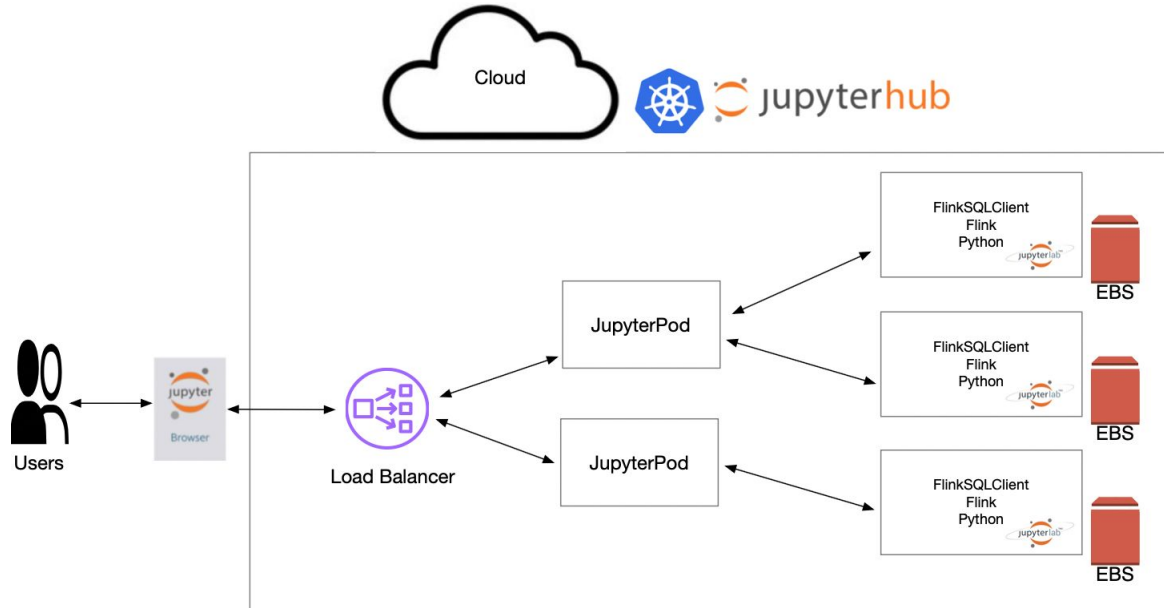
Interactive Streaming Development



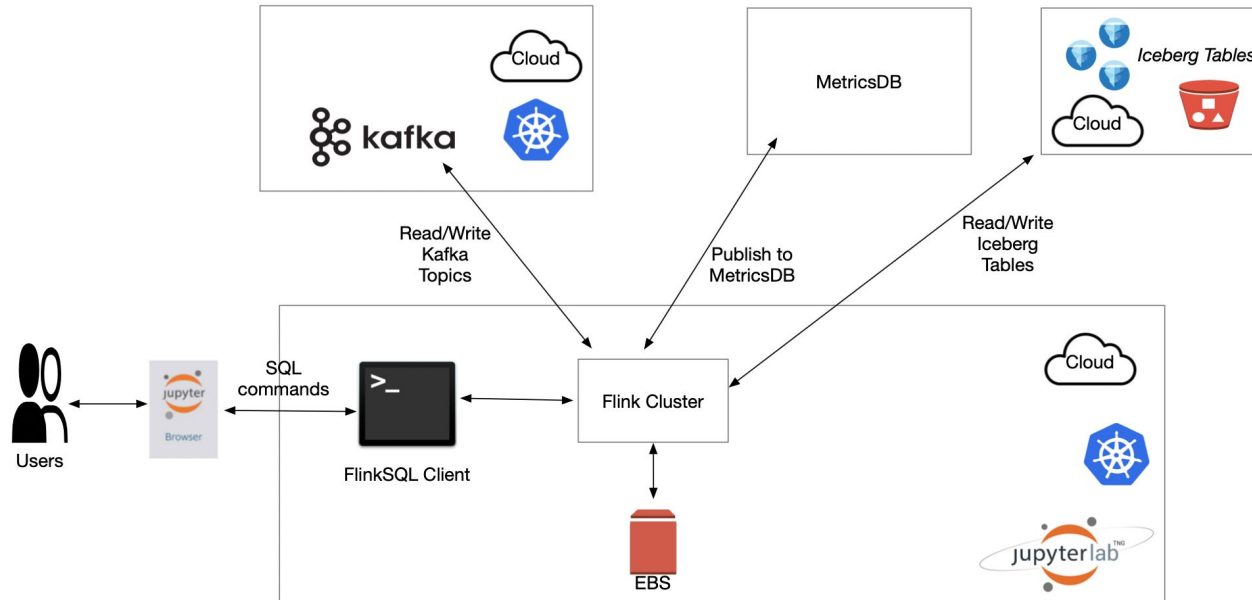
Interactive Streaming Development

- Use Flink SQL client
- Use Python notebooks

Shell on the Clouds: JupyterHub



Flink baked JupyterLab image



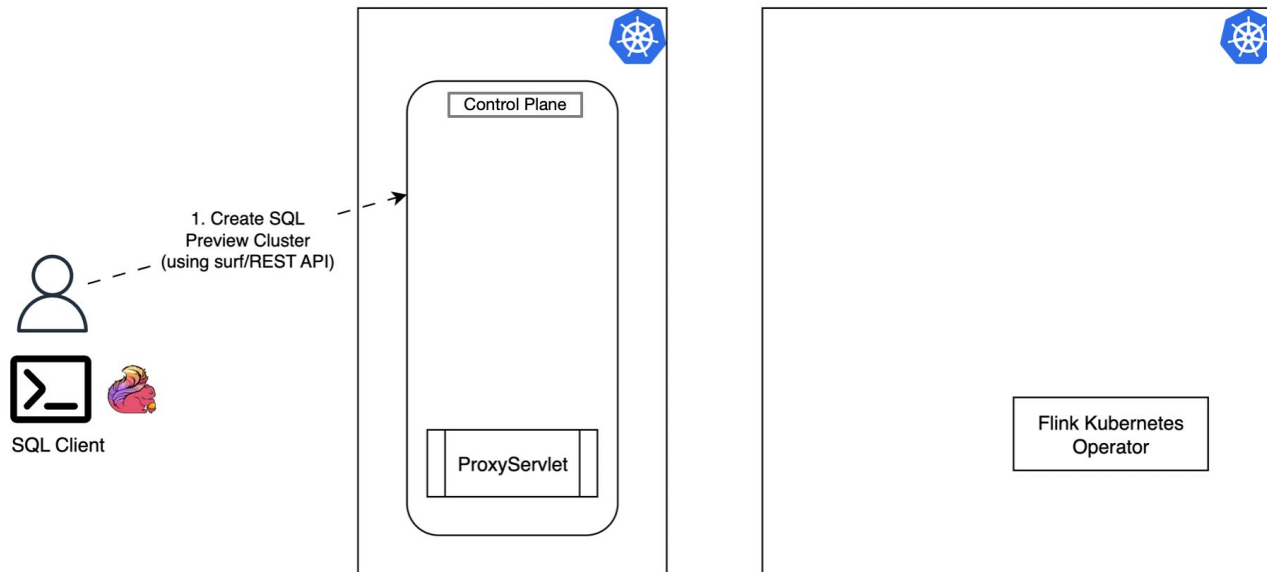


Challenges on single pod

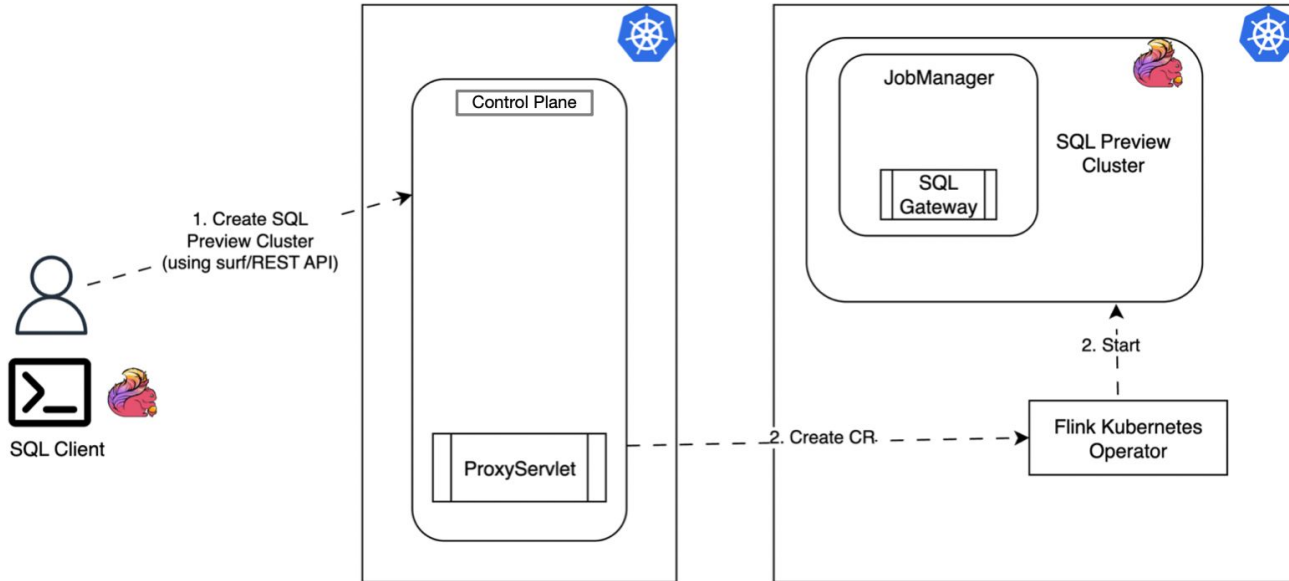
- Challenge:
 - Cannot run at large scale (reading large Iceberg table)
- Solution:
 - Remote execution of Flink pipeline for both Flink SQL & PyFlink

Flink SQL Interactive Development Workflow

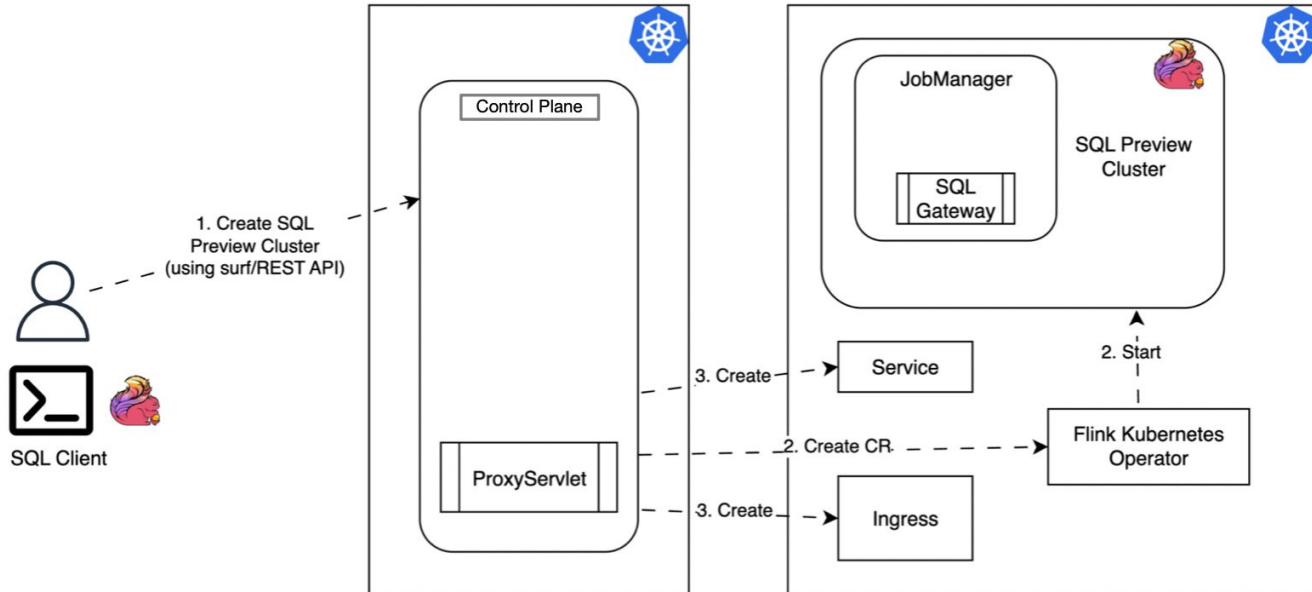
Session Cluster Initialization



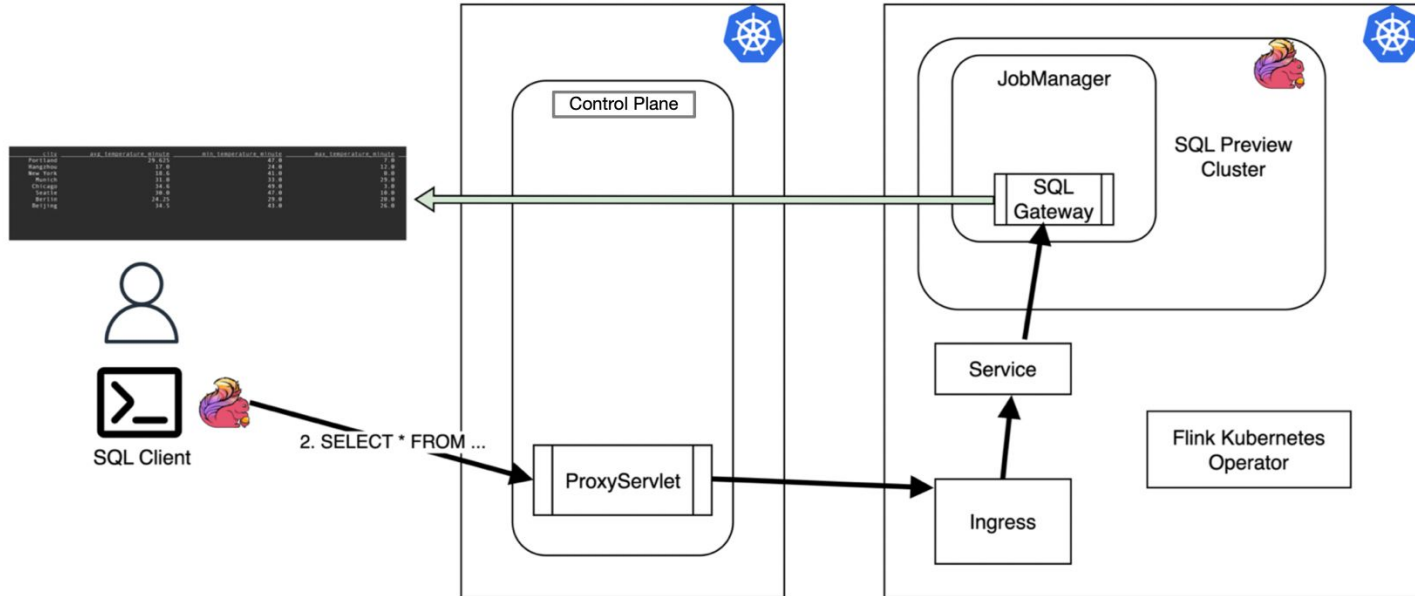
Session Cluster Initialization



Session Cluster Initialization

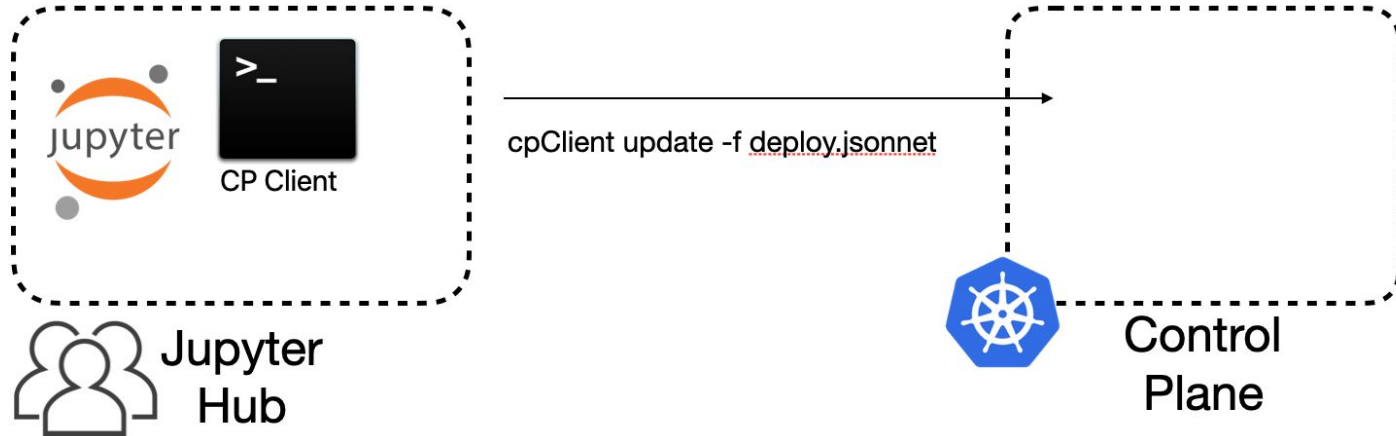


Interactive Flink SQL

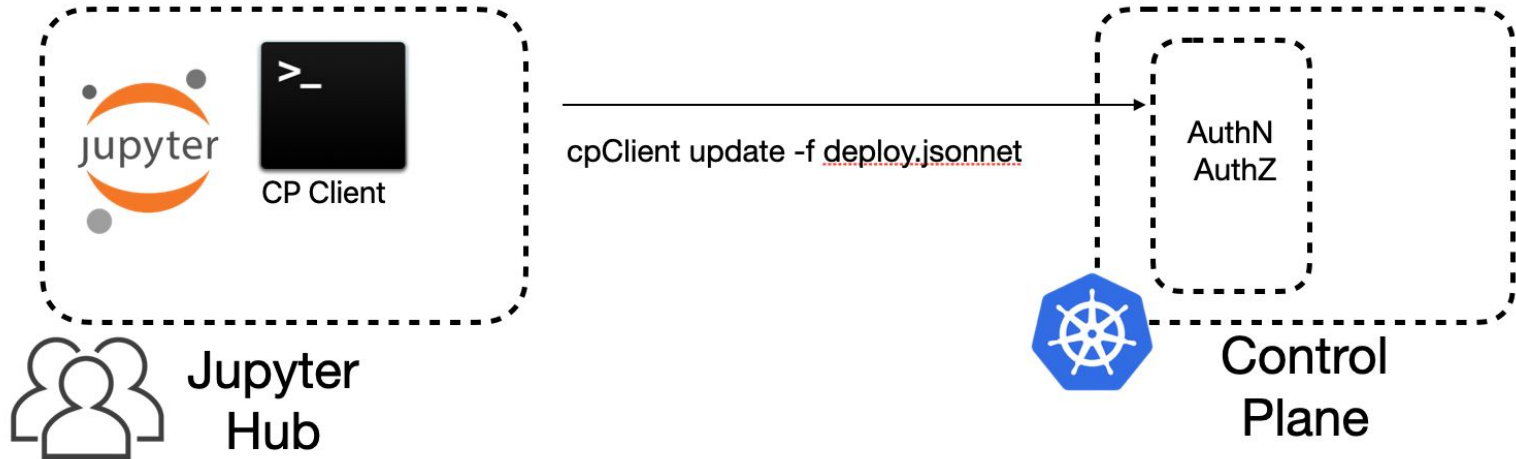


PyFlink Interactive Development Workflow

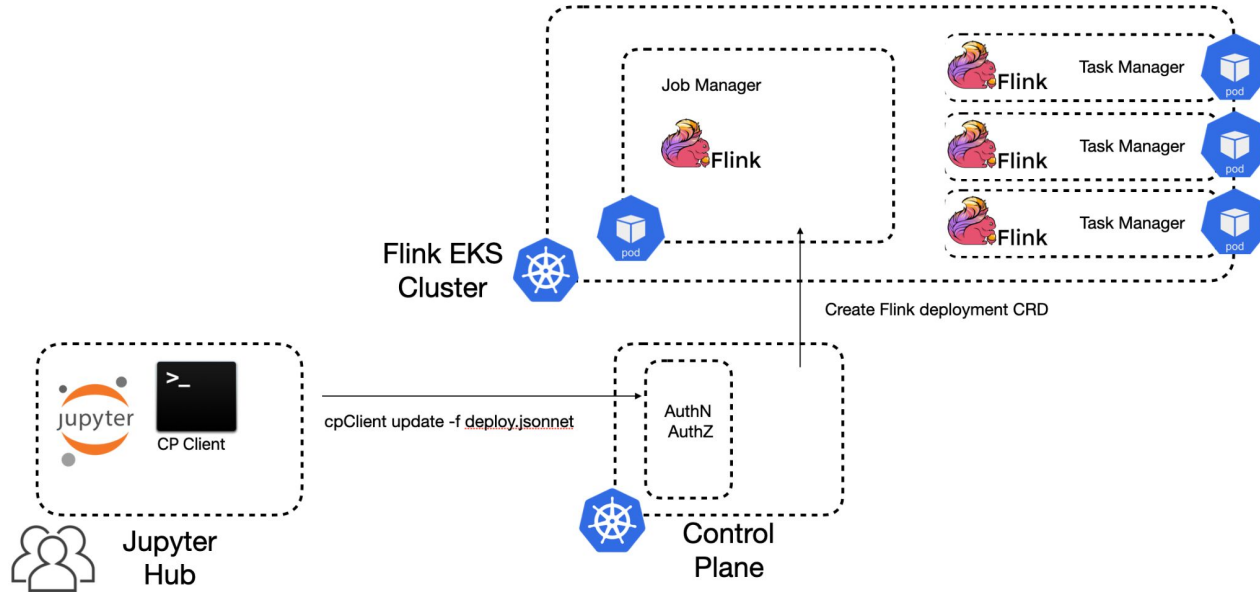
Step 1: Create Standalone cluster on remote Flink EKS cluster



Step 1: Create Standalone cluster on remote Flink EKS cluster



Step 1: Create Standalone cluster on remote Flink EKS cluster



Step 2: PyFlink job submission




Table API PyFlink Examples for Apache Flink

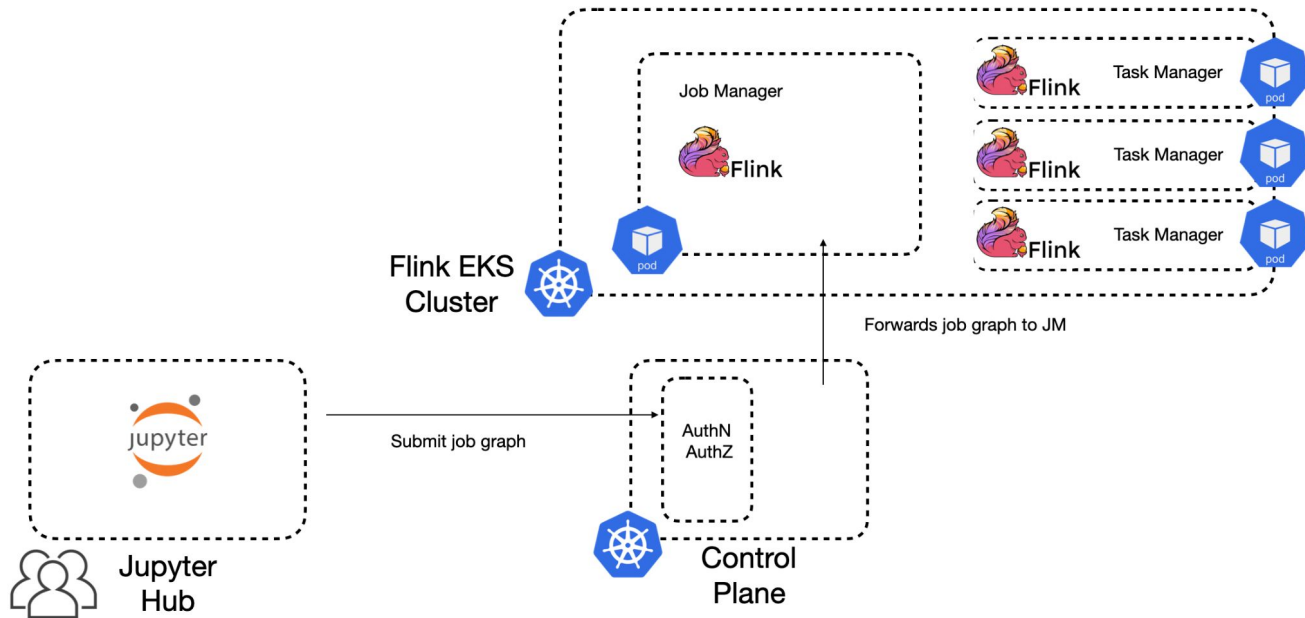
```
[01] 1 import argparse
2 import json
3 import logging
4 import random
5 import sys
6
7 import numpy as np
8 import pandas as pd
9 from pyflink.common import Row, Type
10 from pyflink.common.time import Instant
11 from pyflink.datastream import StreamExecutionEnvironment
12 from pyflink.table import (
13     DataTypes,
14     EnvironmentSettings,
15     FormatDescriptor,
16     Schema,
17     StreamTableEnvironment,
18     TableDescriptor,
19     TableEnvironment,
20 )
21 from pyflink.table.expressions import col, lit
22 from pyflink.table.sdf import Sdf, SdfSink
23 from pyflink.table.window import Tumble
24
25 [02] 1 env = StreamExecutionEnvironment.get_execution_env(environment())
26     env.set_parallelism(1)
27     t_env = StreamTableEnvironment.create(stream_execution_environment=env)
28     # start a checkpoint every 2000 ms
29     env.enable_checkpointing(2000)
30     # advanced options:
31     # # make sure 500 ms of progress happen between checkpoints
32     env.get_checkpoint_config().set_min_pause_between_checkpoints(500)
33     # # checkpoints have to complete within one minute, or are discarded
34     env.get_checkpoint_config().set_checkpoint_timeout(60000)
35     # # allow only one checkpoint to be in progress at the same time
36     env.get_checkpoint_config().set_max_concurrent_checkpoints(1)
37     # # enables the unaligned checkpoints
38     env.get_checkpoint_config().enable_unaligned_checkpoints()
39     t_env.get_config().set("parallelism.default", "1")
40
41 [03] -pyflink.table.config.TableConfig at 8c7f6d26c318>
42
43 Lab: Ingest data into Iceberg from Kafka
44
45 • Create local table with @$Tagen connector (controlled random data)
46 • Create Kafka topic and ingest data into Kafka topic from local table
47 • Create Iceberg table and ingest data into Iceberg table from Kafka topic
48
49 Create engagement_datagen_table table with datagen connector for random data generation
50
51 [04]
```



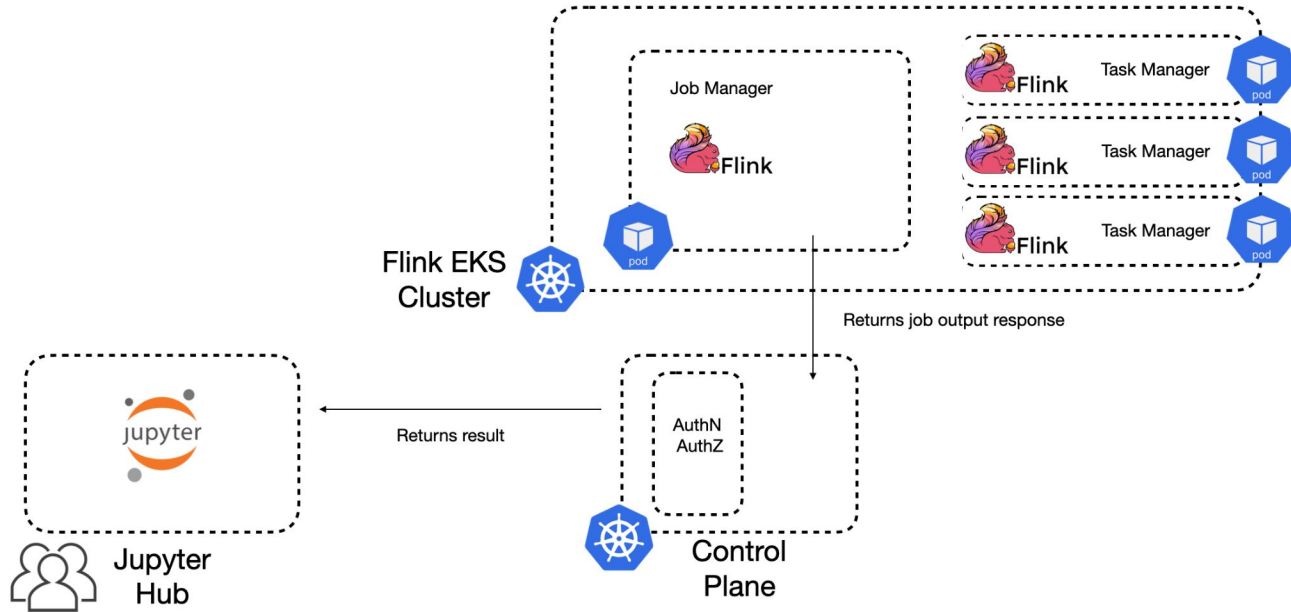
Step 2: PyFlink job submission



Step 2: PyFlink job submission



Step 2: PyFlink job submission





Thank you.