




Select Star:

Flink SQL for Pulsar Folks

Marta Paes (@morsapaes)
Developer Advocate

Why Pulsar + Flink?



J. Doe • 1st

2w ...

Whenever I see [#ApacheFlink](#) mentioned my brain builds an image of [#ApacheSpark](#). After all these years under influence of Spark, that's not a surprise. Same with [#ApachePulsar](#) and [#ApacheKafka](#) (after years with Kafka). That made me think if I'm the only one who considers Flink x Pulsar combo an attempt to "dethrone" Spark x Kafka? And they are all from [The Apache Software Foundation](#) and fully open source! 🙌

Like | Reply



Why Pulsar + Flink?



Why Pulsar + Flink?

"Stream as a unified view on data"



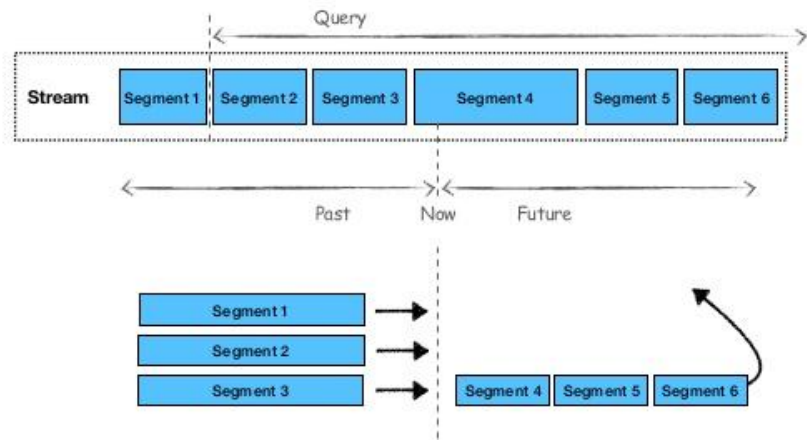
"Batch as a special case of streaming"



Pulsar: Unified Storage

“Stream as a unified view on data”

- Pub/Sub messaging layer (Streaming)
- Durable storage layer (Batch)



Credit: [StreamNative](#)



Unified Storage
(Segments / Pub/Sub)




Flink: Unified Processing Engine

"Batch as a special case of streaming"


- Reuse code and logic across batch and stream processing
- Ensure consistent semantics between processing modes
- Simplify operations
- Power applications mixing historic and real-time data



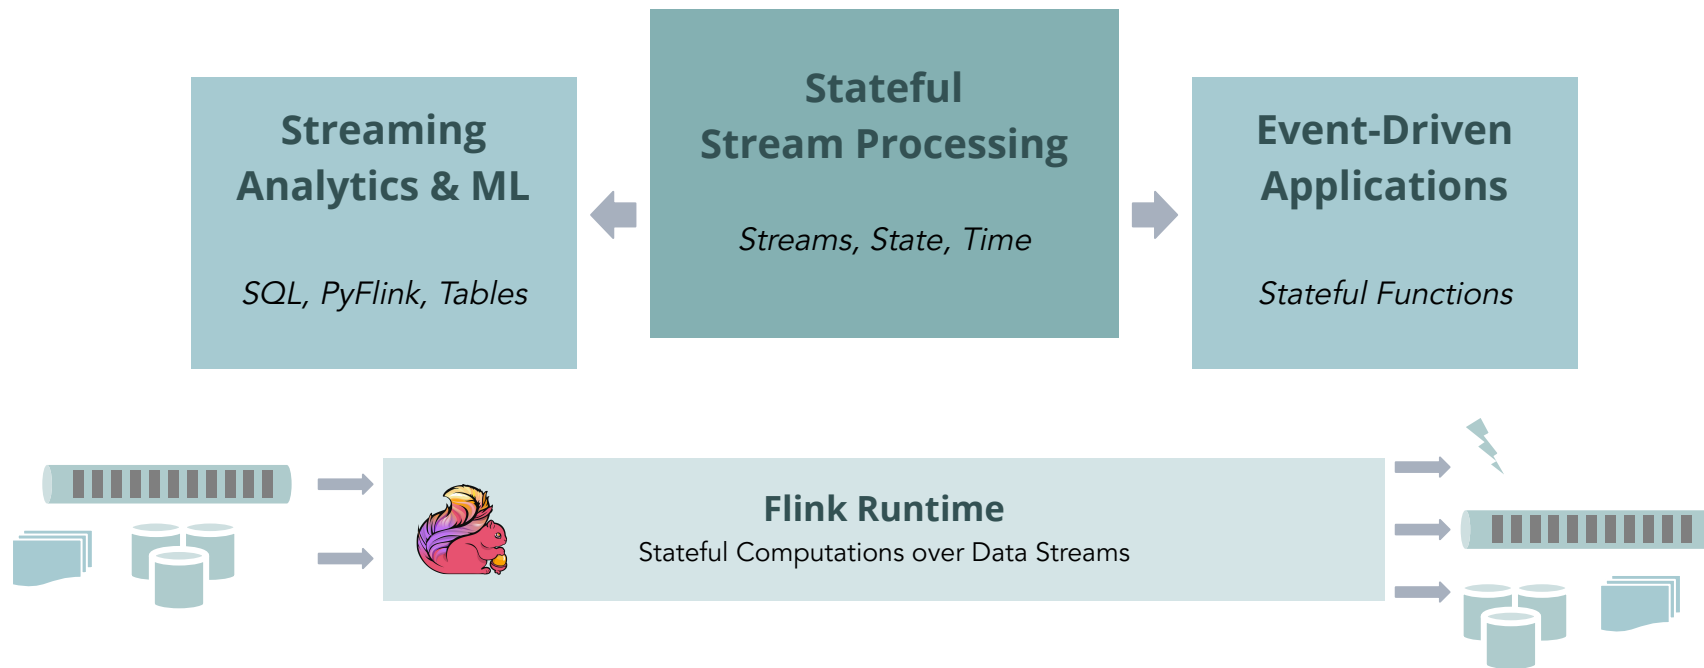
Unified Processing Engine
(Batch / Streaming)



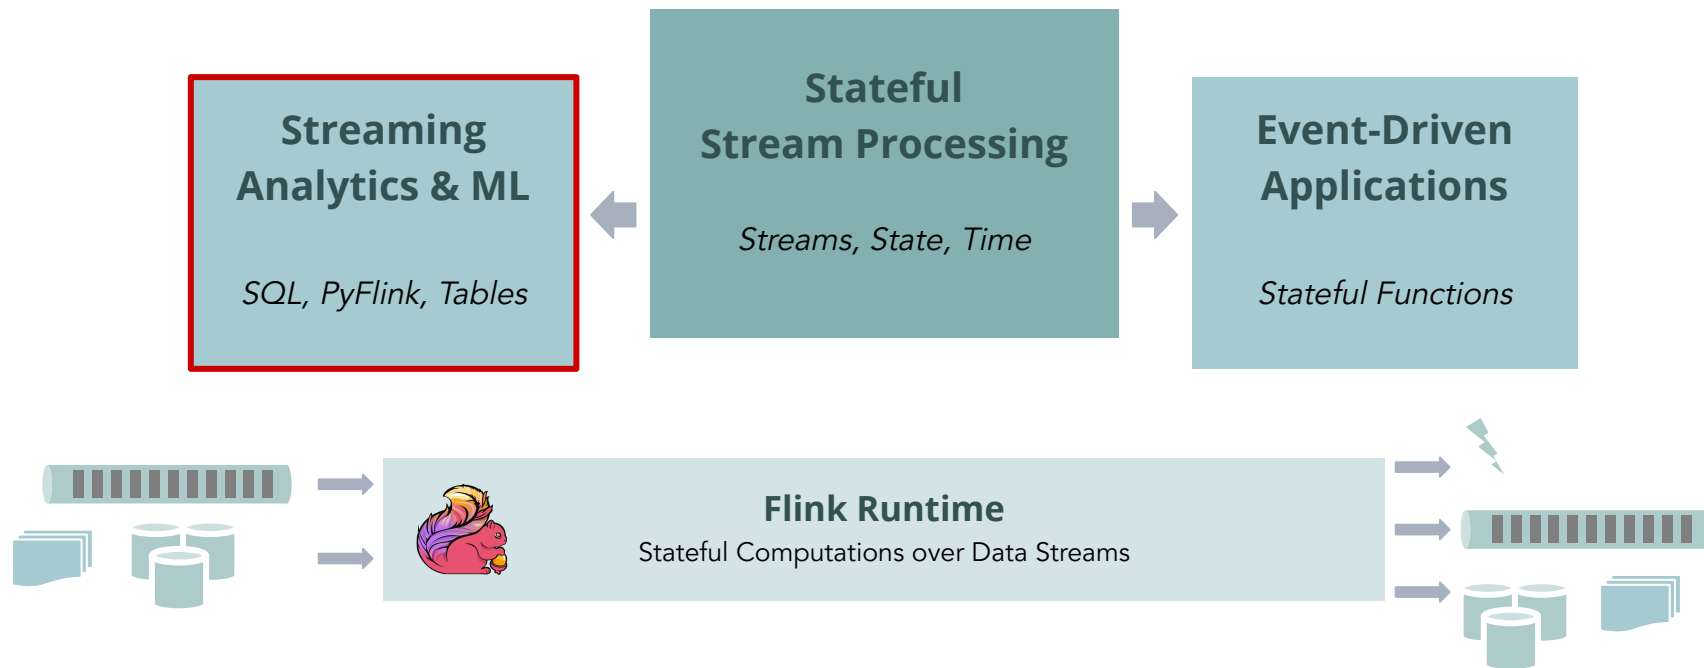
A Unified Data Stack



Flink is broad



Flink is broad



Use Cases

More high-level or domain-specific use cases can be modeled with SQL/Python and dynamic tables.

- Focus on business logic, not implementation
- Mixed workloads (batch and streaming)
- Maximize **developer speed** and **autonomy**

Examples



Uber

criteo.

[Unified Online/Offline Model Training](#)

[E2E Streaming Analytics Pipelines](#)

[ML Feature Generation](#)



Flink SQL

“Everyone knows SQL, right?”

```
SELECT user_id, COUNT(url) AS cnt  
FROM clicks  
GROUP BY user_id;
```

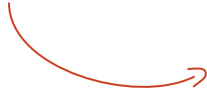
 This is standard SQL (ANSI SQL)



Flink SQL

“Everyone knows SQL, right?”

```
SELECT user_id, COUNT(url) AS cnt
FROM clicks
GROUP BY user_id;
```

 This is ~~standard SQL (ANSI SQL)~~
also Flink SQL



A Regular SQL Engine

Take a snapshot when the query starts

user	cTime	url
Mary	12:00:00	https://...
Bob	12:00:00	https://...
Mary	12:00:02	https://...
Liz	12:00:03	https://...

```
SELECT user_id,  
       COUNT(url) AS cnt  
FROM clicks  
GROUP BY user_id;
```

A final result is produced

user	cnt
Mary	2
Bob	1

A row that was added after the query was started is not considered

The query terminates



A Streaming SQL Engine

Ingest all changes as they happen

Continuously update the result

user	cTime	url
Mary	12:00:00	https://...
Bob	12:00:00	https://...
Mary	12:00:02	https://...
Liz	12:00:03	https://...

```
SELECT user_id,  
       COUNT(url) AS cnt  
FROM clicks  
GROUP BY user_id;
```

user	cnt
Mary	2
Bob	1
Liz	1

The result is identical to the one-time query (at this point)



Flink SQL in a Nutshell


- Standard SQL syntax and semantics (i.e. not a “SQL-flavor”)
- Unified APIs for batch and streaming
- Support for advanced time handling and operations (e.g. CDC, pattern matching)

Execution




TPC-DS Coverage


UDF Support




Native Connectors



Formats



Data Catalogs



Pulsar Integration Over Time

Streaming Source/Sink Connectors

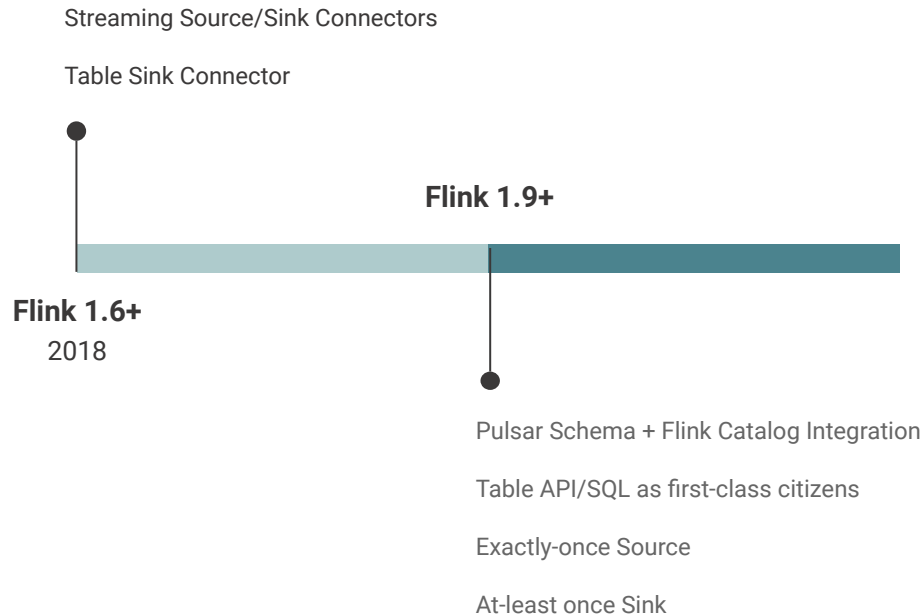
Table Sink Connector



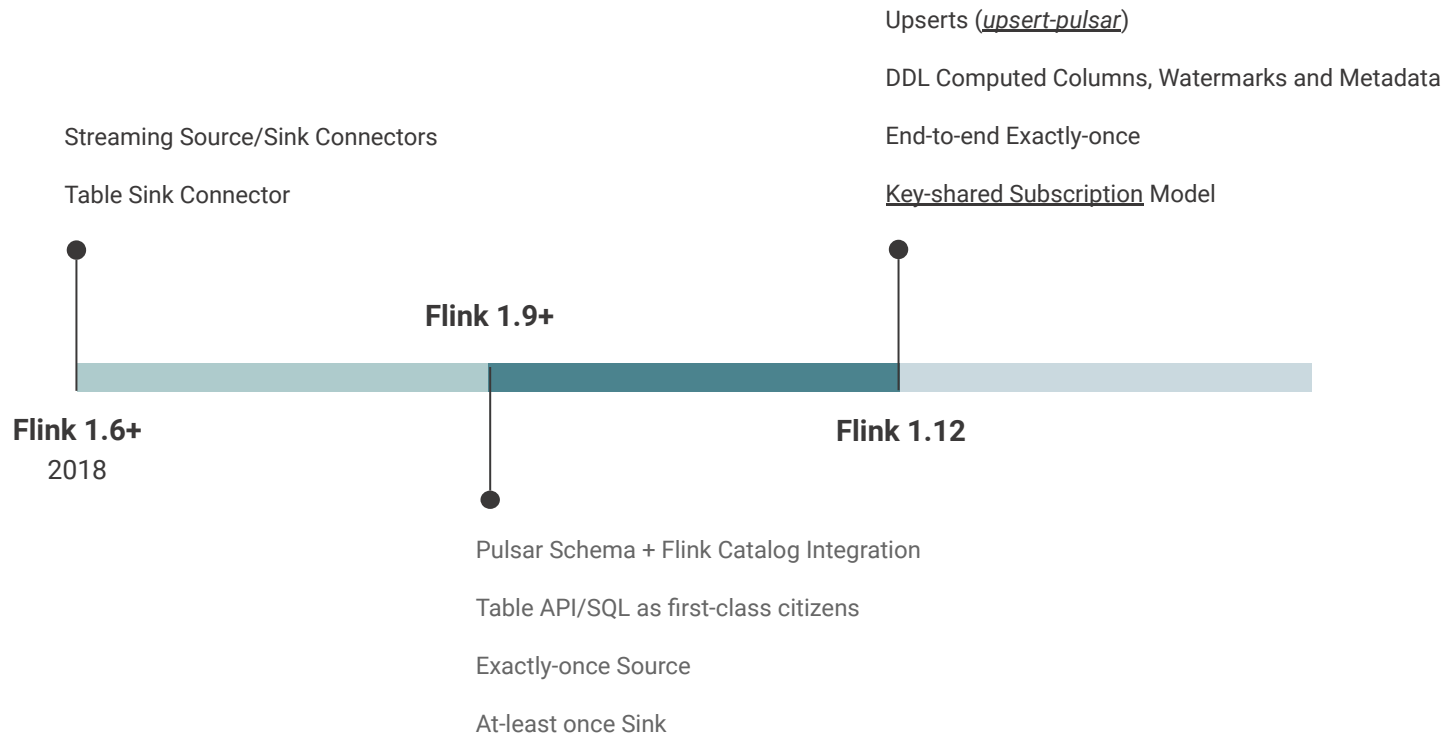
Flink 1.6+
2018



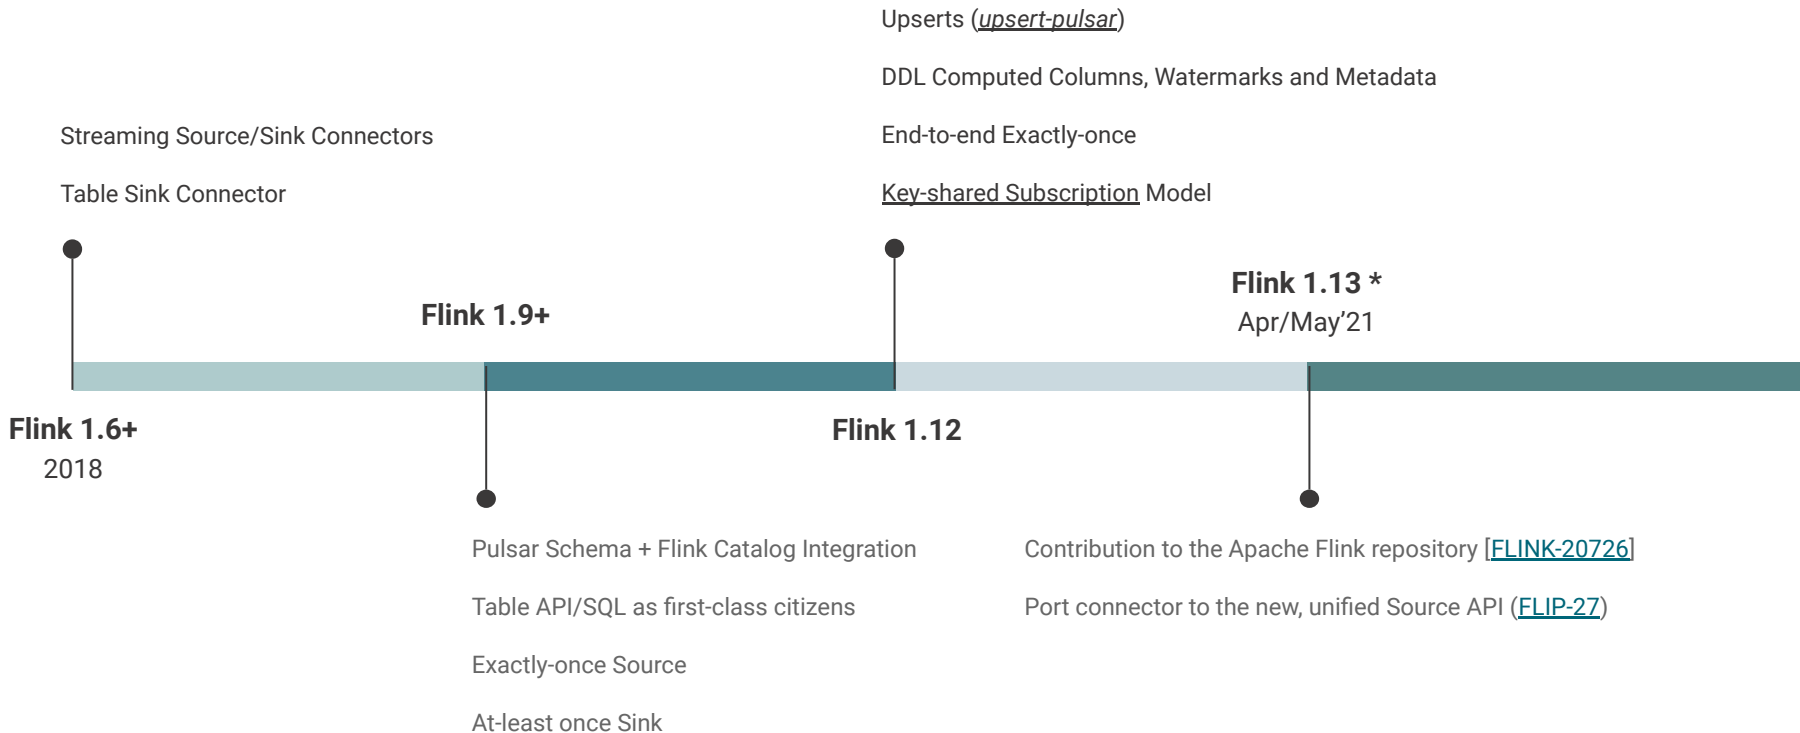
Pulsar Integration Over Time



Pulsar Integration Over Time



Pulsar Integration Over Time



How does this look like, **in practice**?



1. Use the Twitter Firehose [built-in connector](#) to consume tweets about gardening 🌿 into a Pulsar topic (tweets).

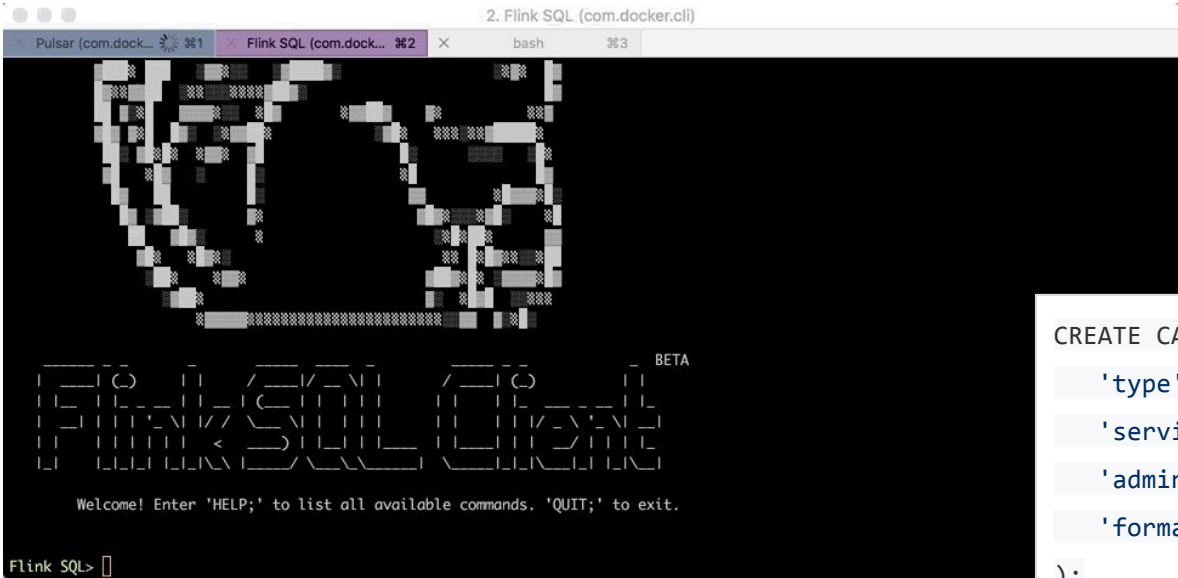
```
2. Pulsar (bash)
Pulsar (bash) %1 Flink SQL (bash) %2 bash %3
morsapaes@local:flink-sql-pulsar (main *)$
```

```
2. Pulsar (com.docker.cli)
Pulsar (com.docker.cli) %1 Flink SQL (bash) %2 bash %3
16:37:45.473 [pulsar-client-io-1-1] INFO org.apache.pulsar.client.impl.ConsumerImpl - [tweets][test] Subscribed to topic on localhost/127.0.0.1:6650 -- consumer: 0
16:37:52.186 [pulsar-client-io-1-1] INFO com.scurrilous.circe.checksum.Crc32CIntChecksum - SSE4.2 CRC32C provider initialized
----- got message -----
key:[null], properties:[], content:{"id":"1372225659821568000","text":"@radishroot2002 Soooo envious! We dug a smaller version of this last autumn but so far have only beetles. Really ho_ https://t.co/gPHu0YKw3","source":{"<a href="http://twitter.com/#/download/ipad" rel="nofollow">Twitter for iPad</a>"},"truncated":true,"user":{"id":1477092500,"name":"crow's feet \u003D\u00C99#FBPE #FBPPR #FBPA","description":"atheist, centre-left tosser, art, #Rejoin \u003C\u00DEA\u003C\u00DDFA","verified":false},"favorited":false,"retweeted":false,"lang":"en"}}
----- got message -----
key:[null], properties:[], content:{"id":"1372225666356113412","text":"@meltrue7 This is so cool! I've been doing some gardening but I have killed about half of my plants. I'm still lear_ https://t.co/OKc6D07Eyz","source":{"<a href="https://mobile.twitter.com" rel="nofollow">Twitter Web App</a>"},"truncated":true,"user":{"id":1308792227821281281,"name":"BGF_12","description":"she/her, banner by @vanilla_dipped","verified":false},"favorited":false,"retweeted":false,"lang":"en"}}
----- got message -----
key:[null], properties:[], content:{"id":"1372225720869646347","text":"It's sad. Mines like that because they need watering - in pots!","source":{"<a href="http://twitter.com/download/iphone" rel="nofollow">Twitter for iPhone</a>"},"truncated":false,"user":{"id":1146067339512692736,"name":"soulbag50grvs","location":"Cambride","description":"Caring sharing listener","verified":false},"favorited":false,"retweeted":false,"lang":"en"}}
----- got message -----
key:[null], properties:[], content:{"id":"1372225750775005192","text":"RT @LotteTruelsen: We're having some very cold mornings...Hellebore Cinderella and I don't like it...&#x26#flowers #WinterGarden #gardening h_","source":{"<a href="http://twitter.com/download/iphone" rel="nofollow">Twitter for iPhone</a>"},"truncated":false,"user":{"id":1146067339512692736,"name":"soulbag50grvs","location":"Cambride","description":"Caring sharing listener","verified":false},"favorited":false,"retweeted":false,"lang":"en"}}
```



2. Start the Flink SQL Client and use a Pulsar catalog to access the topic directly as a table in Flink.

SQL Client



Catalog DDL

```
CREATE CATALOG pulsar WITH (  
  'type' = 'pulsar',  
  'service-url' = 'pulsar://pulsar:6650',  
  'admin-url' = 'http://pulsar:8080',  
  'format' = 'json'  
);
```



2.1. You can query the tweets topic off-the-bat using a simple SELECT statement — it's treated as a Flink table!

SQL Client

The screenshot shows a terminal window titled "2. Flink SQL (com.docker.cli)". The terminal displays the output of a SQL query, which is a table with the following columns: createdAt, id, idStr, text, and source. The data is as follows:

createdAt	id	idStr	text	source
(NULL)	1372229856721784834	(NULL)	RT @thefreckledrose: Her~	<a href="https://mobile.~
(NULL)	137222987459777411	(NULL)	RT @PicturedImage: Good ~	<a href="http://twitter.~

At the bottom of the terminal, there are several keyboard shortcuts: Quit Refresh, Inc Refresh Dec Refresh, Goto Page Last Page, Next Page Prev Page, and Open Row.



2.2. But then you find out that most Firehose events have a null createdAt. What now?

SQL Client

Refresh: 1 s SQL Query Result (Table) Updated: 16:54:38.846
Page: Last of 1

createdAt	id	idStr	text	source
(NULL)	1372229856721784834	(NULL)	RT @thefreckledrose: Her~	<a href="https://mobile.~
(NULL)	137222987459777411	(NULL)	RT @PicturedImage: Good ~	<a href="http://twitter.~

Not cool. 😞

Quit Refresh Inc Refresh Dec Refresh Goto Page Last Page Next Page Prev Page Open Row



3. One way to get a relevant timestamp is to use Pulsar metadata to get the `publishTime` (i.e. ingestion time).

Source Table DDL

```
CREATE TABLE pulsar_tweets (  
    publishTime TIMESTAMP(3) METADATA,  
    WATERMARK FOR publishTime AS publishTime - INTERVAL '5' SECOND  
) WITH (  
    'connector' = 'pulsar',  
    'topic' = 'persistent://public/default/tweets',  
    'value.format' = 'json',  
    'service-url' = 'pulsar://pulsar:6650',  
    'admin-url' = 'http://pulsar:8080',  
    'scan.startup.mode' = 'earliest-offset'  
)  
LIKE tweets;
```

Read and use Pulsar message metadata

Define the source connector (Pulsar)

Derive schema from the original topic



4. Perform a simple windowed aggregation (count), and insert results into a new pulsar topic (tweets_agg).

Sink Table DDL

```
CREATE TABLE pulsar_tweets_agg (  
    tmstamp TIMESTAMP(3),  
    tweet_cnt BIGINT  
) WITH (  
    'connector'='pulsar',  
    'topic'='persistent://public/default/tweets_agg',  
    'value.format'='json',  
    'service-url'='pulsar://pulsar:6650',  
    'admin-url'='http://pulsar:8080'  
);
```

Continuous SQL Query

```
INSERT INTO pulsar_tweets_agg  
SELECT TUMBLE_START(publishTime, INTERVAL '10' SECOND) AS  
wStart,  
    COUNT(id) AS tweet_cnt  
FROM pulsar_tweets  
GROUP BY TUMBLE(publishTime, INTERVAL '10' SECOND);
```



5. We'll get a count of the # of tweets in windows of 10 seconds (based on event time!).

```
2. Flink SQL (com.docker.cli)
Pulsar (com.docker.cli) %1  Flink SQL (com.dock... %2  bash %3
Flink SQL> SELECT TUMBLE_START(publishTime, INTERVAL '10' SECOND) AS wStart,
>     COUNT(id) AS tweet_cnt
> FROM pulsar_tweets
> GROUP BY TUMBLE(publishTime, INTERVAL '10' SECOND);
>
```



There's a lot more to it!

Check out the [Flink SQL Cookbook](#), where we share hands-on examples, patterns, and use cases for Flink SQL.

09 Maintaining Materialized Views with Change Data Capture (CDC) and Debezium

Flink Version 1.11+

This example will show how you can use Flink SQL and Debezium to monitor changes in a database. In the world of analytics, databases are still mostly seen as static sources waiting to be queried. The reality is that most of the data stored in these databases is so...why not stream it?

Change Data Capture (CDC) allows you to do just that: track and propagate changes in a database (e.g. [Ahead-Log](#) in Postgres) to downstream consumers. [Debezium](#) is a popular CDC tool that can be used to capture changes in a database and stream them to a Kafka connector and 2) a set of "standalone" [Flink CDC Connectors](#).

Let's get to it!

In this example, you'll monitor a table with insurance claim data related to aggregated **materialized view** that is **incrementally updated** with the latest data by deploying Debezium, Kafka and Kafka Connect in [this repository](#).

Pre-requisites

05 Real Time Star Schema Denormalization (N-Way Join)

In this recipe, we will de-normalize a simple star schema with an n-way temporal table join.

Star schemas are a popular way of normalizing data within a data warehouse. At the center of a star schema is a **fact table** whose rows contain metrics, measurements, and other facts about the world. Surrounding fact tables contain metadata useful for enriching facts when computing queries. You are running a small data warehouse for a railroad company which consists of a fact table (`stations` , `booking_channels` , and `passengers`). All inserts to the fact table, and all upserts to the `stations` and `booking_channels` tables are interpreted as inserts only, and so the table is configured with `connector = kafka` ;. In contrast, the records in the dimensional tables are upserts because they are updated frequently. The `fact` table is configured with `connector = upsert-kafka` .

With Flink SQL you can now easily join all dimensions to our fact table using a 5-way temporal join (left input/probe site) and correlate each row to the corresponding row's right input/build side). Flink uses the SQL syntax of `FOR SYSTEM_TIME AS OF` to perform this operation. This ensures consistent, reproducible results when joining a fact table with more (slowly) changing dimensions. The fact table is joined to its corresponding value of each dimension based on when the event occurred.

Script

```
CREATE TEMPORARY TABLE passengers (
```

08 Detecting patterns with MATCH_RECOGNIZE

This example will show how you can use Flink SQL to detect patterns in a stream of events with `MATCH_RECOGNIZE` .

A common (but historically complex) task in SQL day-to-day work is to identify meaningful sequences of events in a data set — also known as Complex Event Processing (CEP). This becomes even more relevant when dealing with streaming data, as you want to react quickly to known patterns or changing trends to deliver up-to-date business insights. In Flink SQL, you can easily perform this kind of tasks using the standard SQL clause `MATCH_RECOGNIZE` .

Breaking down MATCH_RECOGNIZE

In this example, you want to find users that downgraded their service subscription from one of the premium tiers (`type IN ('premium', 'platinum')`) to the basic tier.

Input

The input argument of `MATCH_RECOGNIZE` will be a row pattern table based on `subscriptions` . As a first step, logical partitioning and ordering must be applied to the input row pattern table to ensure that event processing is correct and deterministic:

```
PARTITION BY user_id
ORDER BY proc_time
```

Output





Thank you!

Follow me on Twitter: @morsapaes

Learn more about Flink: <https://flink.apache.org/>