

# Stream Processing As You've Never Seen Before (Seriously)

Apache Flink for Java Developers



Apache Flink



X/Bluesky: @gamussa





# Stream Processing As You've Never Seen Before (Seriously)

Apache Flink for Java Developers



Apache Flink



X/Bluesky: @gamussa

# Viktor GAMOV

---

Principal Developer Advocate | **Confluent**  
co-author of Manning's **Kafka in Action**  
**Java Champion**



# Slides and Video

<https://speaking.gamov.io/>



X/Bluesky: @gamussa

A pixelated squirrel is positioned in the center of the slide, facing right. It is surrounded by a background of a dense forest with various colored trees (blue, green, yellow, orange) and a blue sky with white clouds.

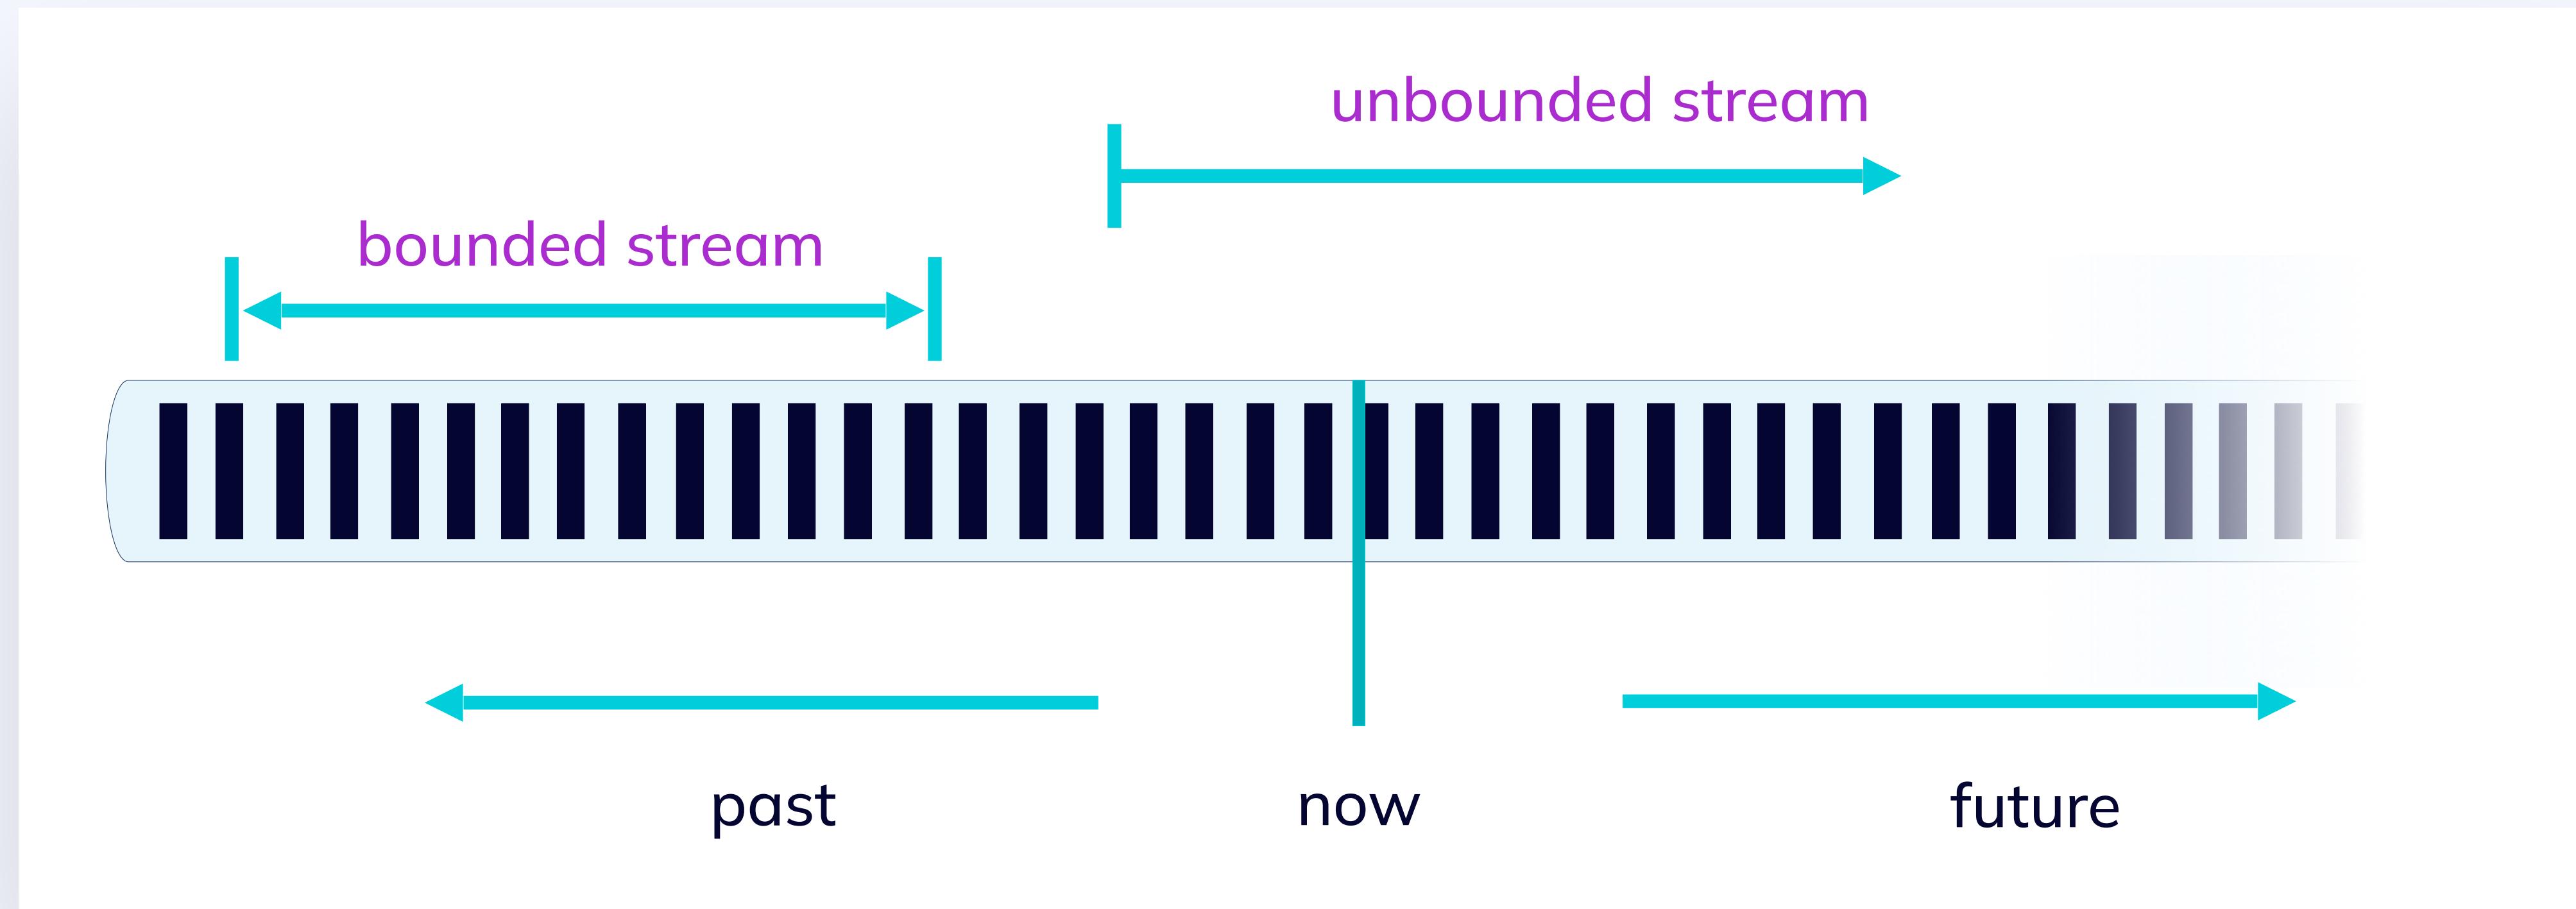
What is  
Apache Flink?

# What is Flink?



**Apache Flink is a framework and distributed processing engine for stateful computations over unbounded and bounded data streams.**

# Streaming

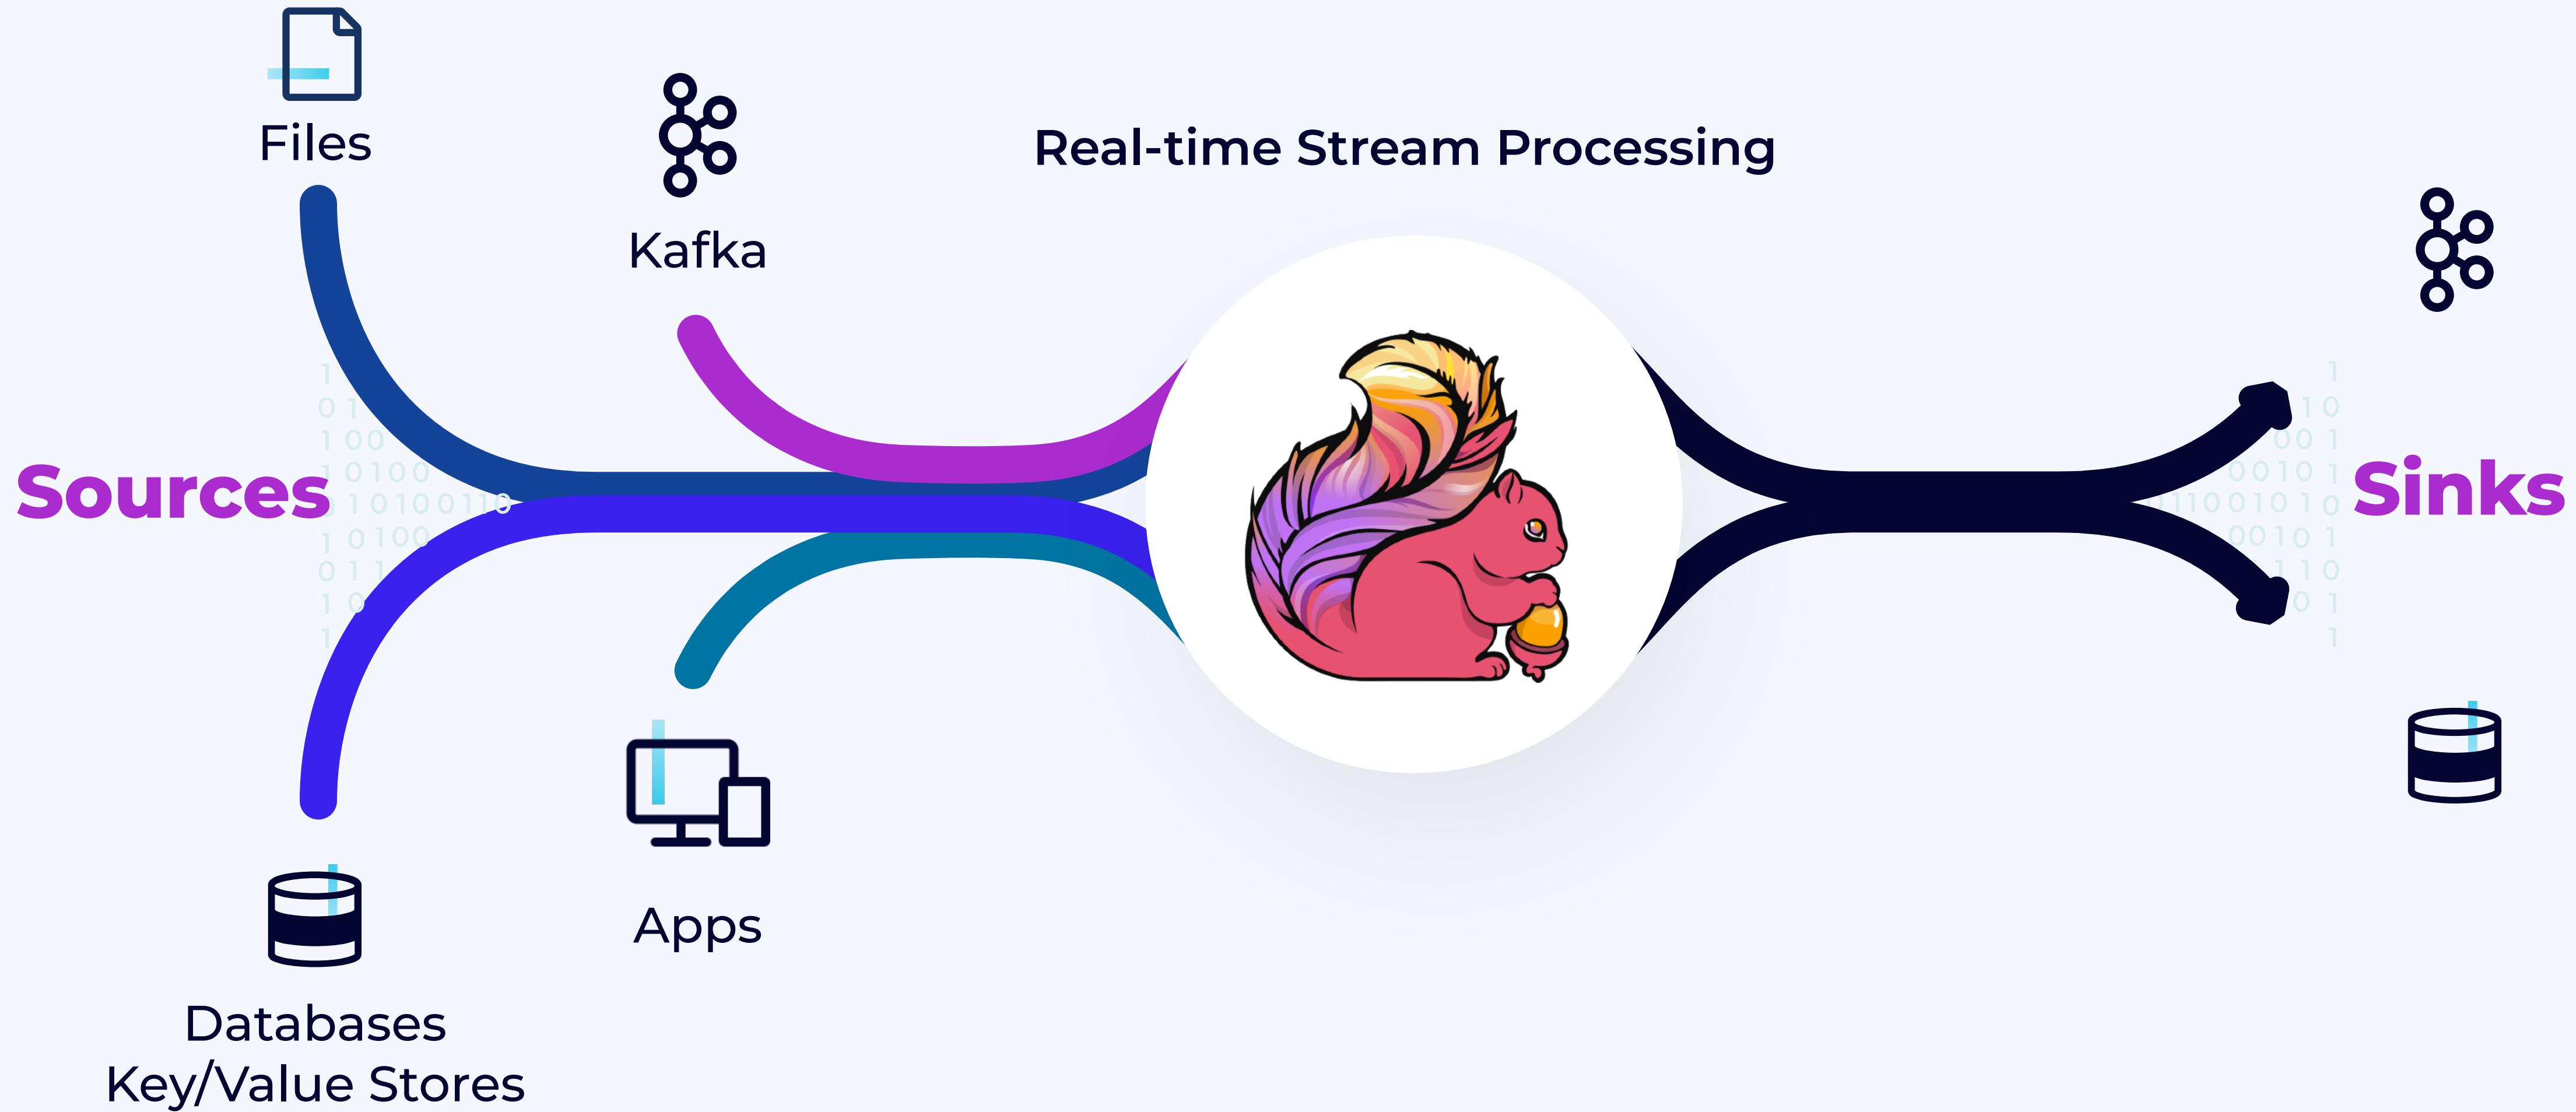


- A stream is a **sequence** of events
- Business data is always a stream: **bounded** or **unbounded**
- For Flink, batch processing is just a **special case** in the runtime

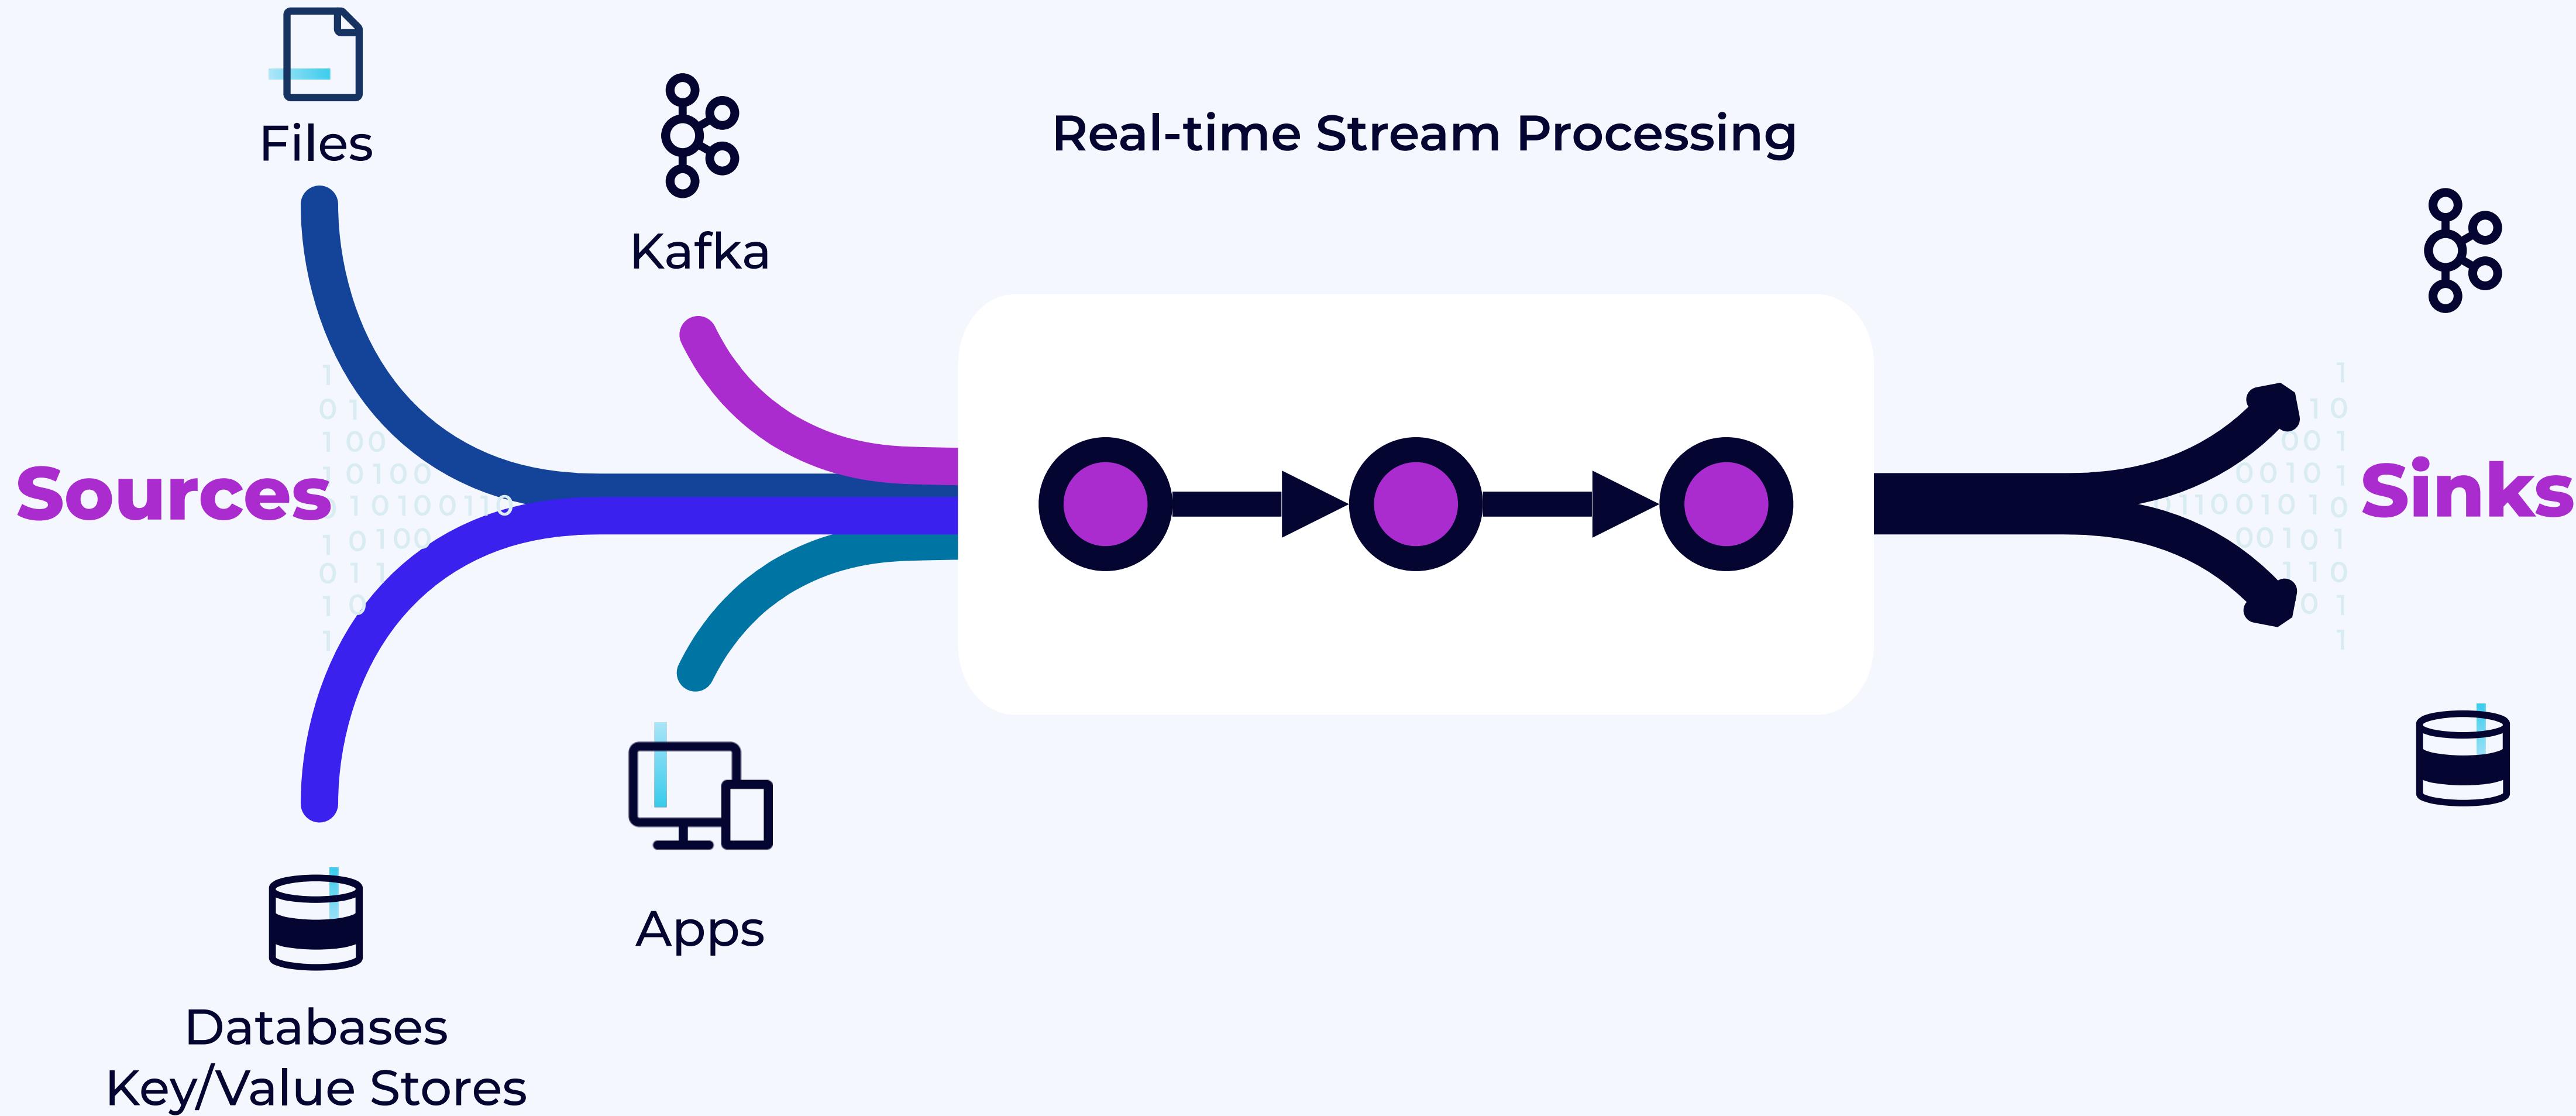
# Key Features of Flink

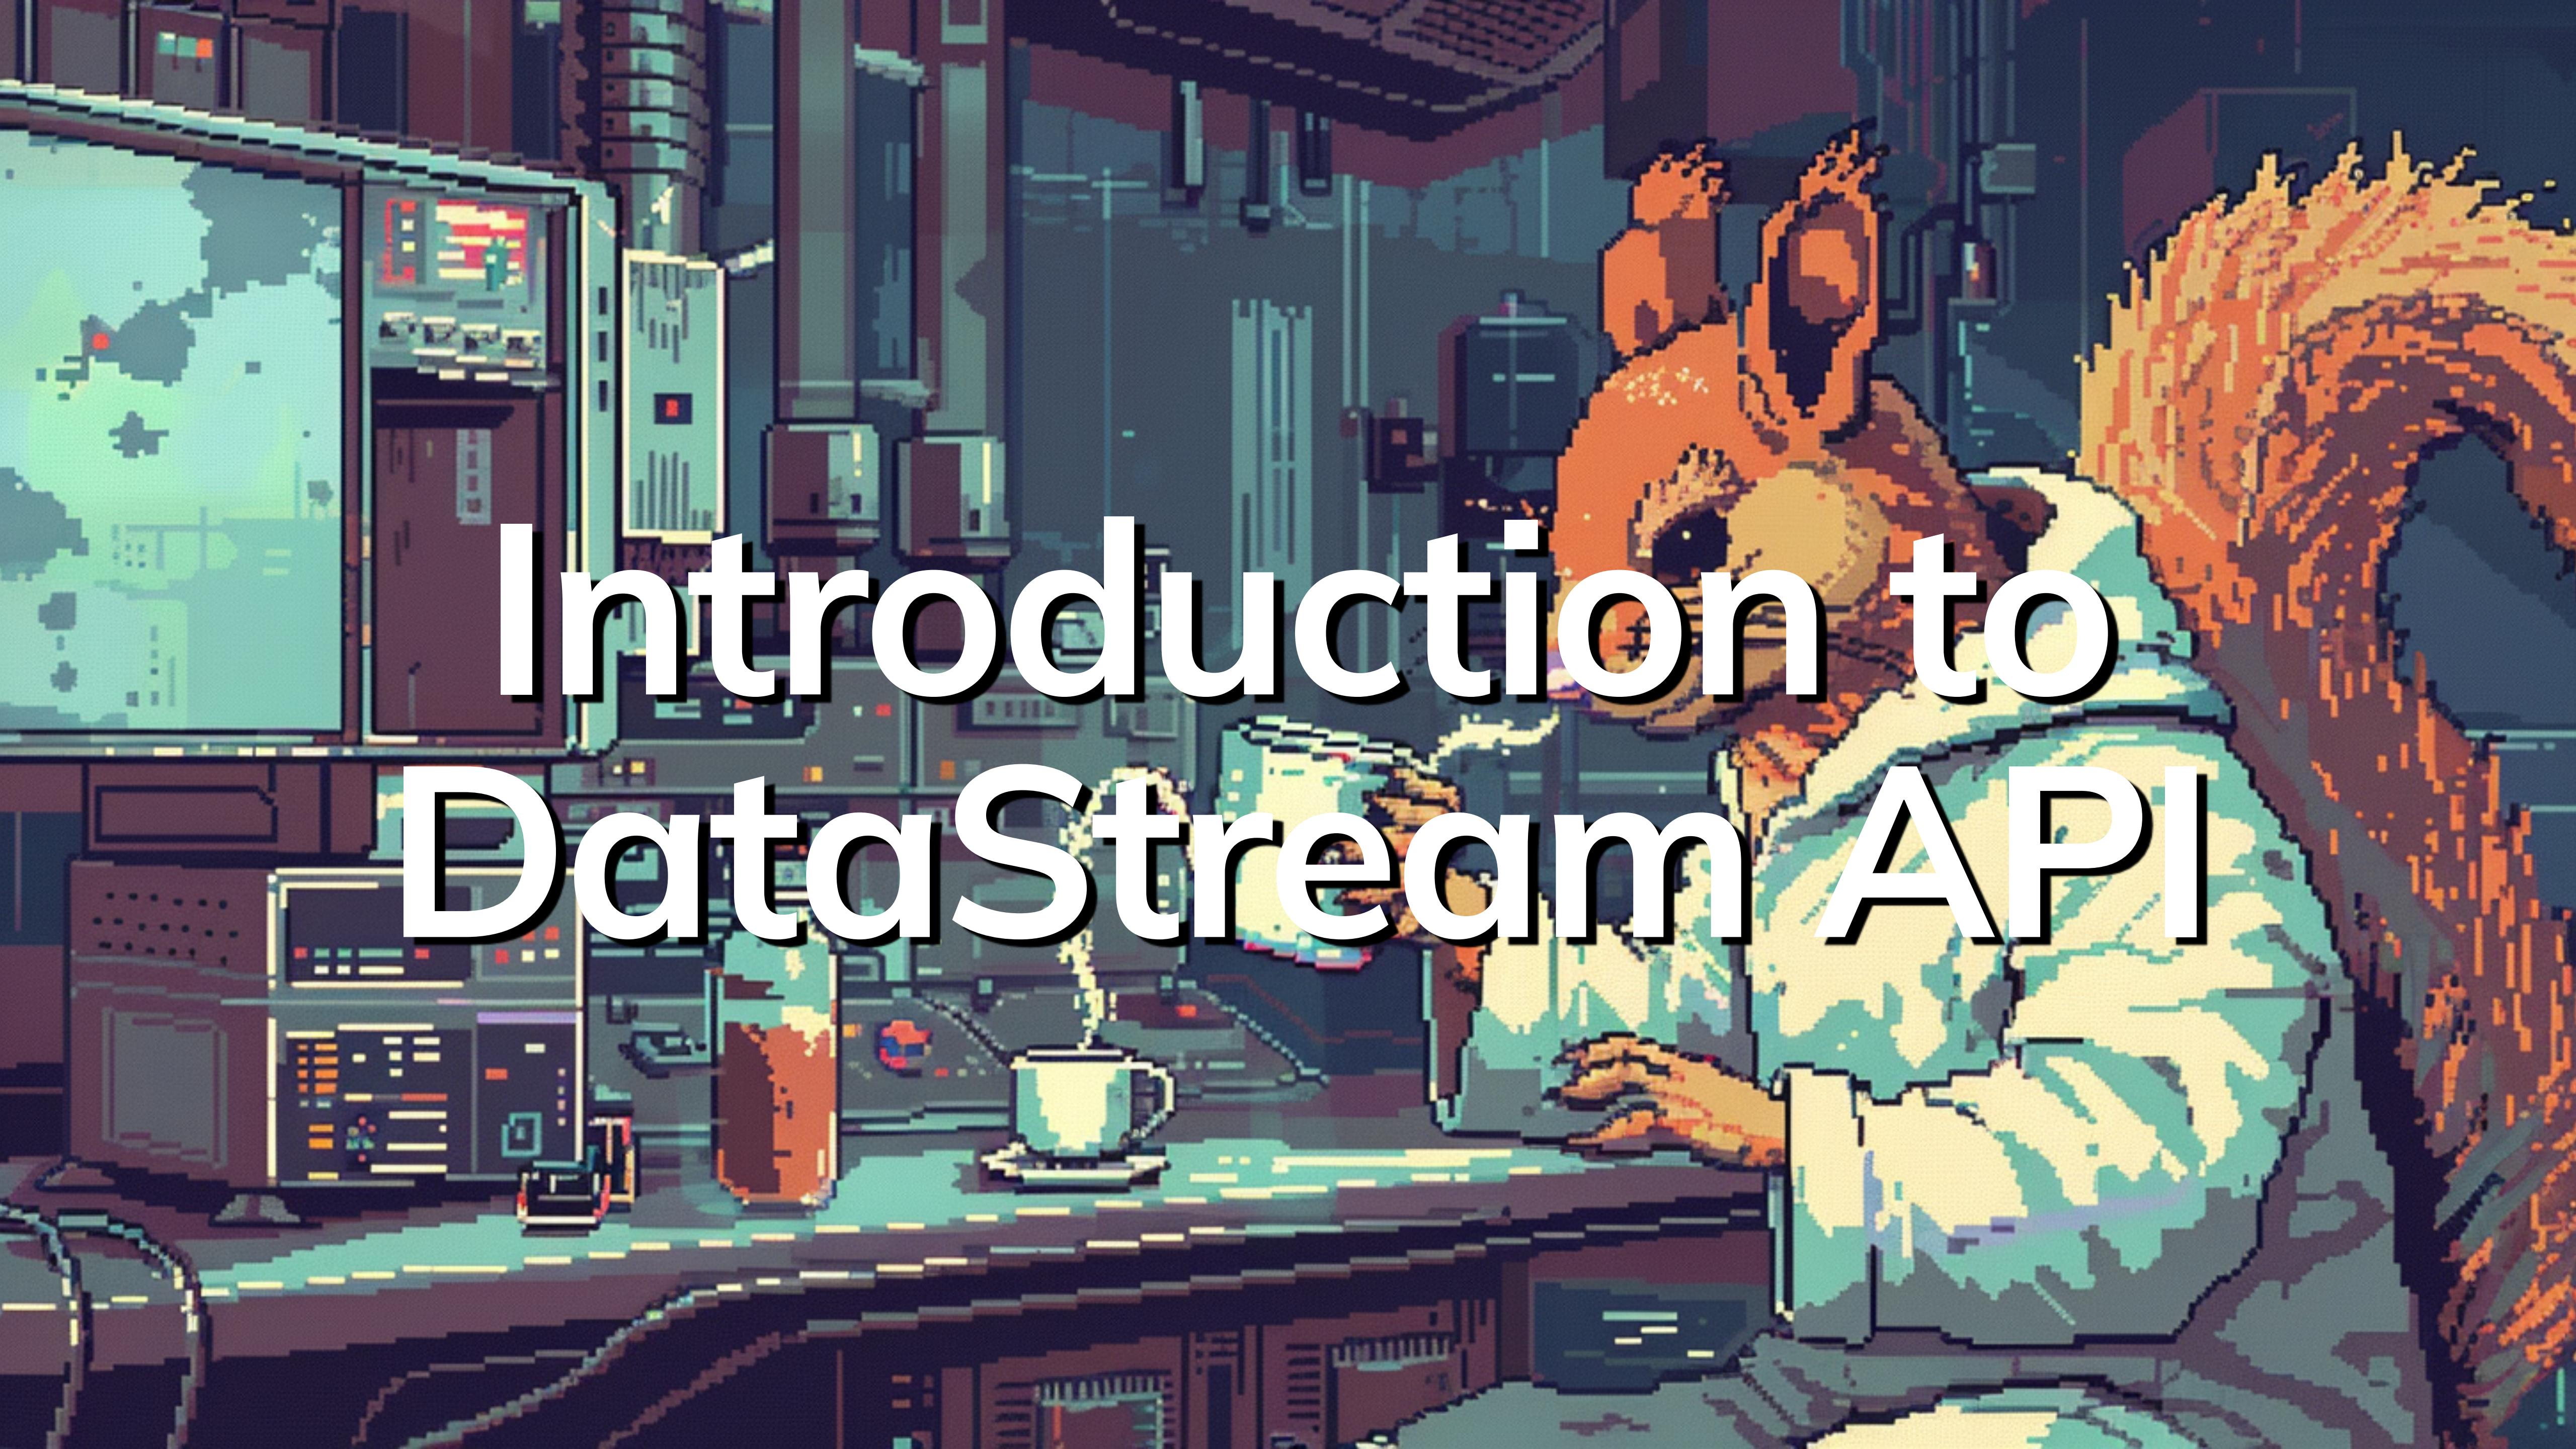


# Stream processing with Flink



# Stream processing with Flink



A pixelated, retro-futuristic background depicting a dense city skyline with various buildings, some with glowing windows. In the foreground, a person wearing a hooded jacket is seated at a desk, facing away from the viewer. On the desk, there's a computer monitor displaying a colorful interface, a keyboard, and a mouse. The overall aesthetic is reminiscent of early video games or science fiction imagery.

# Introduction to DataStream API

```
// Set up the execution environment
StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();  
  
// Create a DataStream from some elements
DataStream<String> inputStream = env.fromData("apple", "banana", "cherry", "date", "elderberry");  
  
// Perform a transformation
DataStream<Tuple2<String, Integer>> resultStream = inputStream
    .map(value -> new Tuple2<>(value, value.length()))
    .returns(Types.TUPLE(Types.STRING, Types.INT));  
  
// Print the results to the console
resultStream.print();  
  
// Execute the Flink job
env.execute("Simple Flink Job");
```

```
// Set up the execution environment
StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();

// Create a DataStream from some elements
DataStream<String> inputStream = env.fromData("apple", "banana", "cherry", "date", "elderberry");

// Perform a transformation
DataStream<Tuple2<String, Integer>> resultStream = inputStream
    .map(value -> new Tuple2<>(value, value.length()))
    .returns(Types.TUPLE(Types.STRING, Types.INT));

// Print the results to the console
resultStream.print();

// Execute the Flink job
env.execute("Simple Flink Job");
```

```
// Set up the execution environment
StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();

// Create a DataStream from some elements
DataStream<String> inputStream = env.fromData("apple", "banana", "cherry", "date", "elderberry");

// Perform a transformation
DataStream<Tuple2<String, Integer>> resultStream = inputStream
    .map(value -> new Tuple2<>(value, value.length()))
    .returns(Types.TUPLE(Types.STRING, Types.INT));

// Print the results to the console
resultStream.print();

// Execute the Flink job
env.execute("Simple Flink Job");
```

```
// Set up the execution environment
StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();

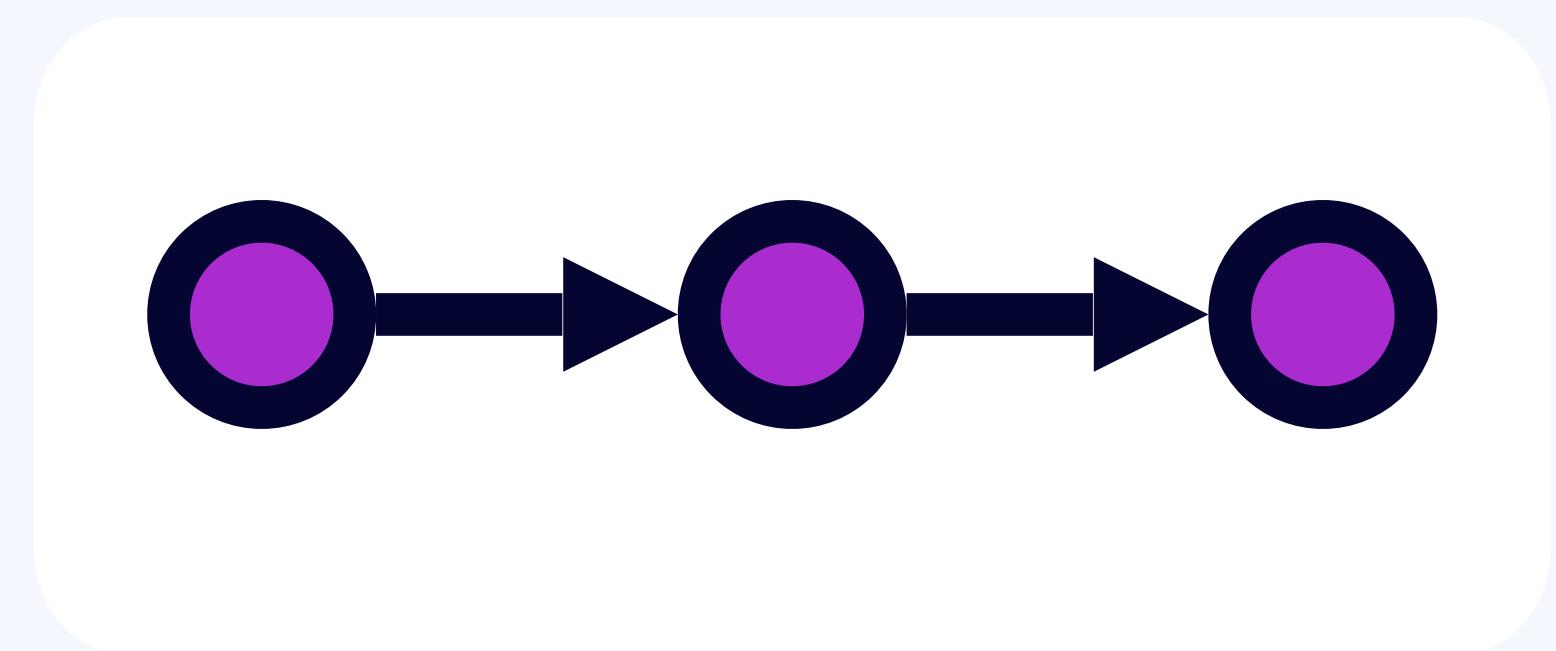
// Create a DataStream from some elements
DataStream<String> inputStream = env.fromData("apple", "banana", "cherry", "date", "elderberry");

// Perform a transformation
DataStream<Tuple2<String, Integer>> resultStream = inputStream
    .map(value -> new Tuple2<>(value, value.length()))
    .returns(Types.TUPLE(Types.STRING, Types.INT));

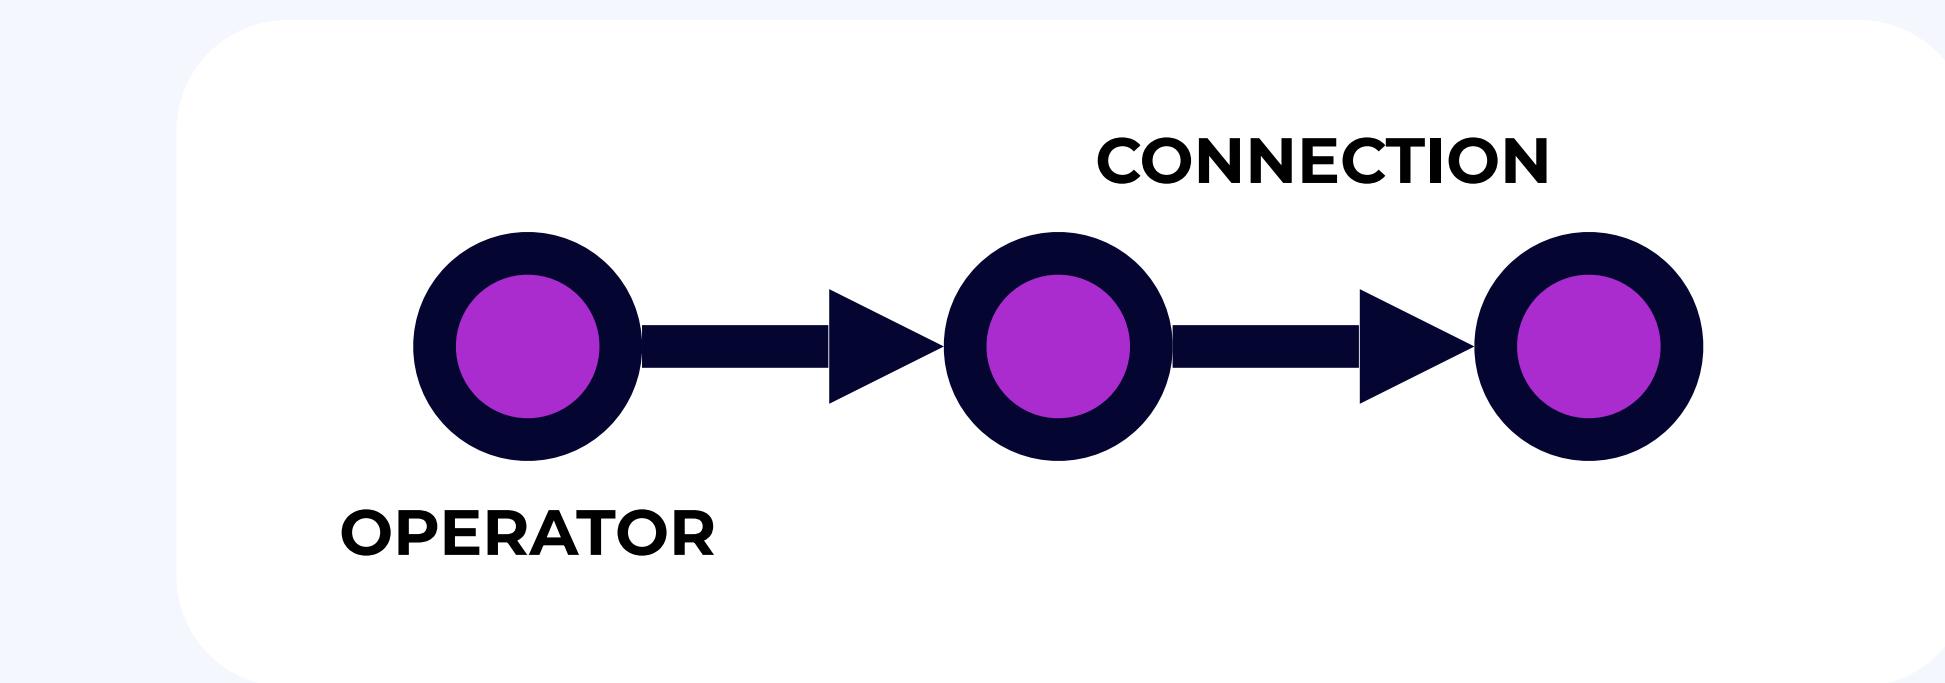
// Print the results to the console
resultStream.print();

// Execute the Flink job
env.execute("Simple Flink Job");
```

# The JobGraph (or topology)



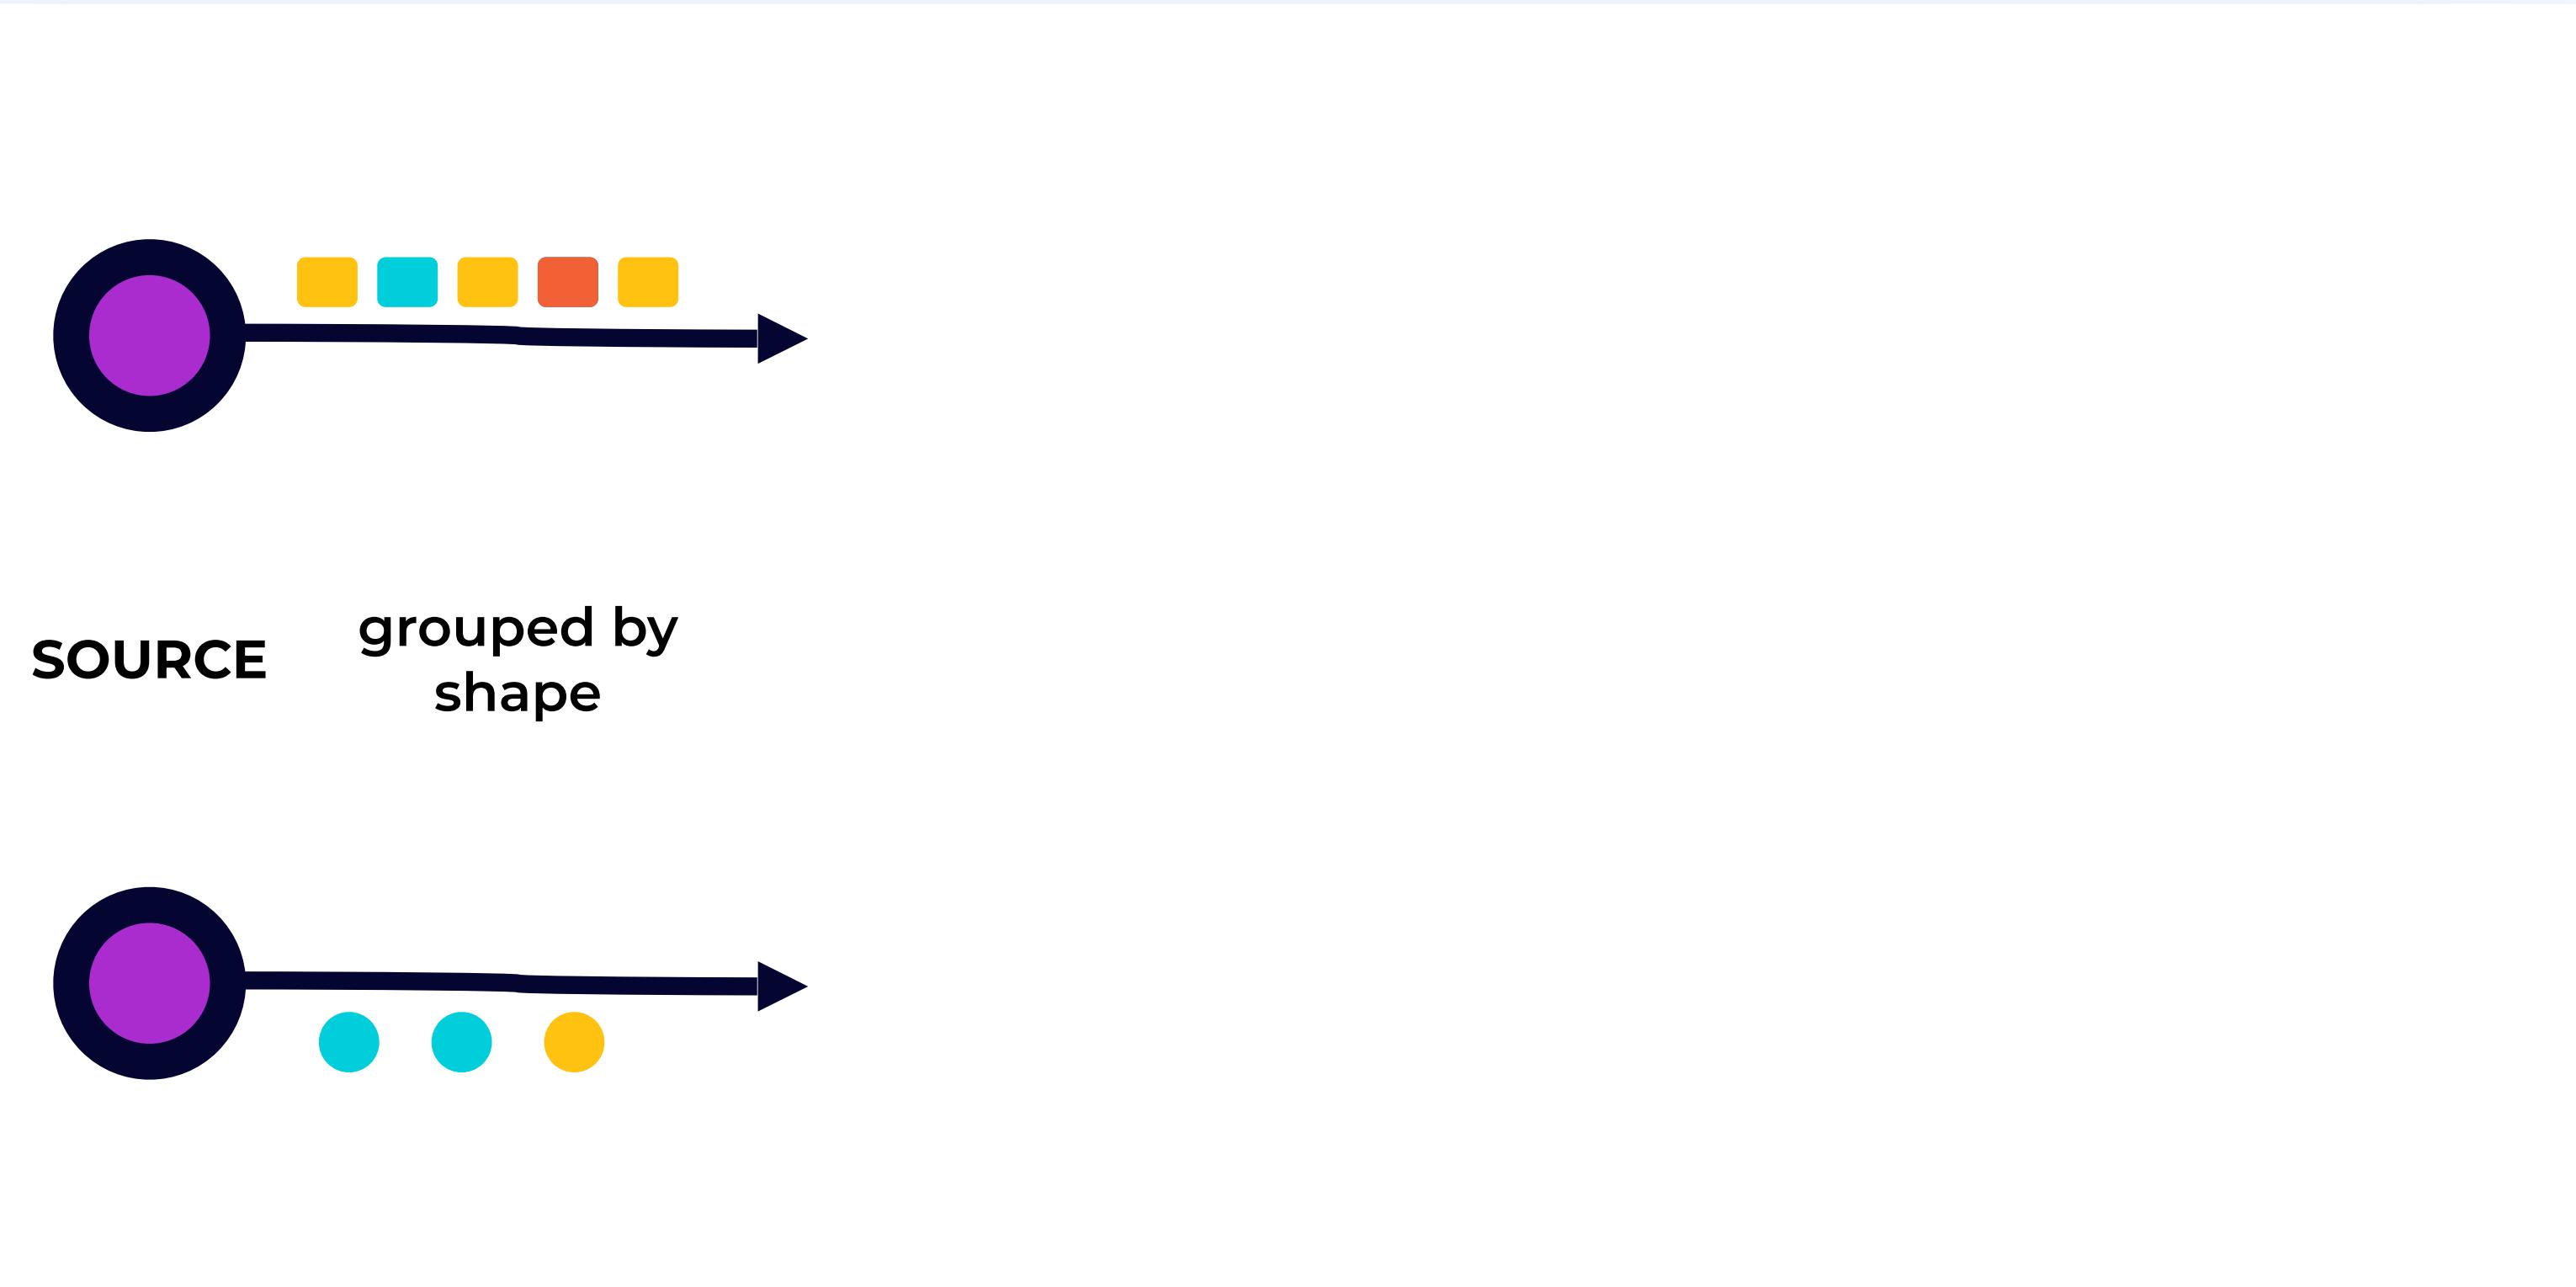
# The JobGraph (or topology)



# Stream processing



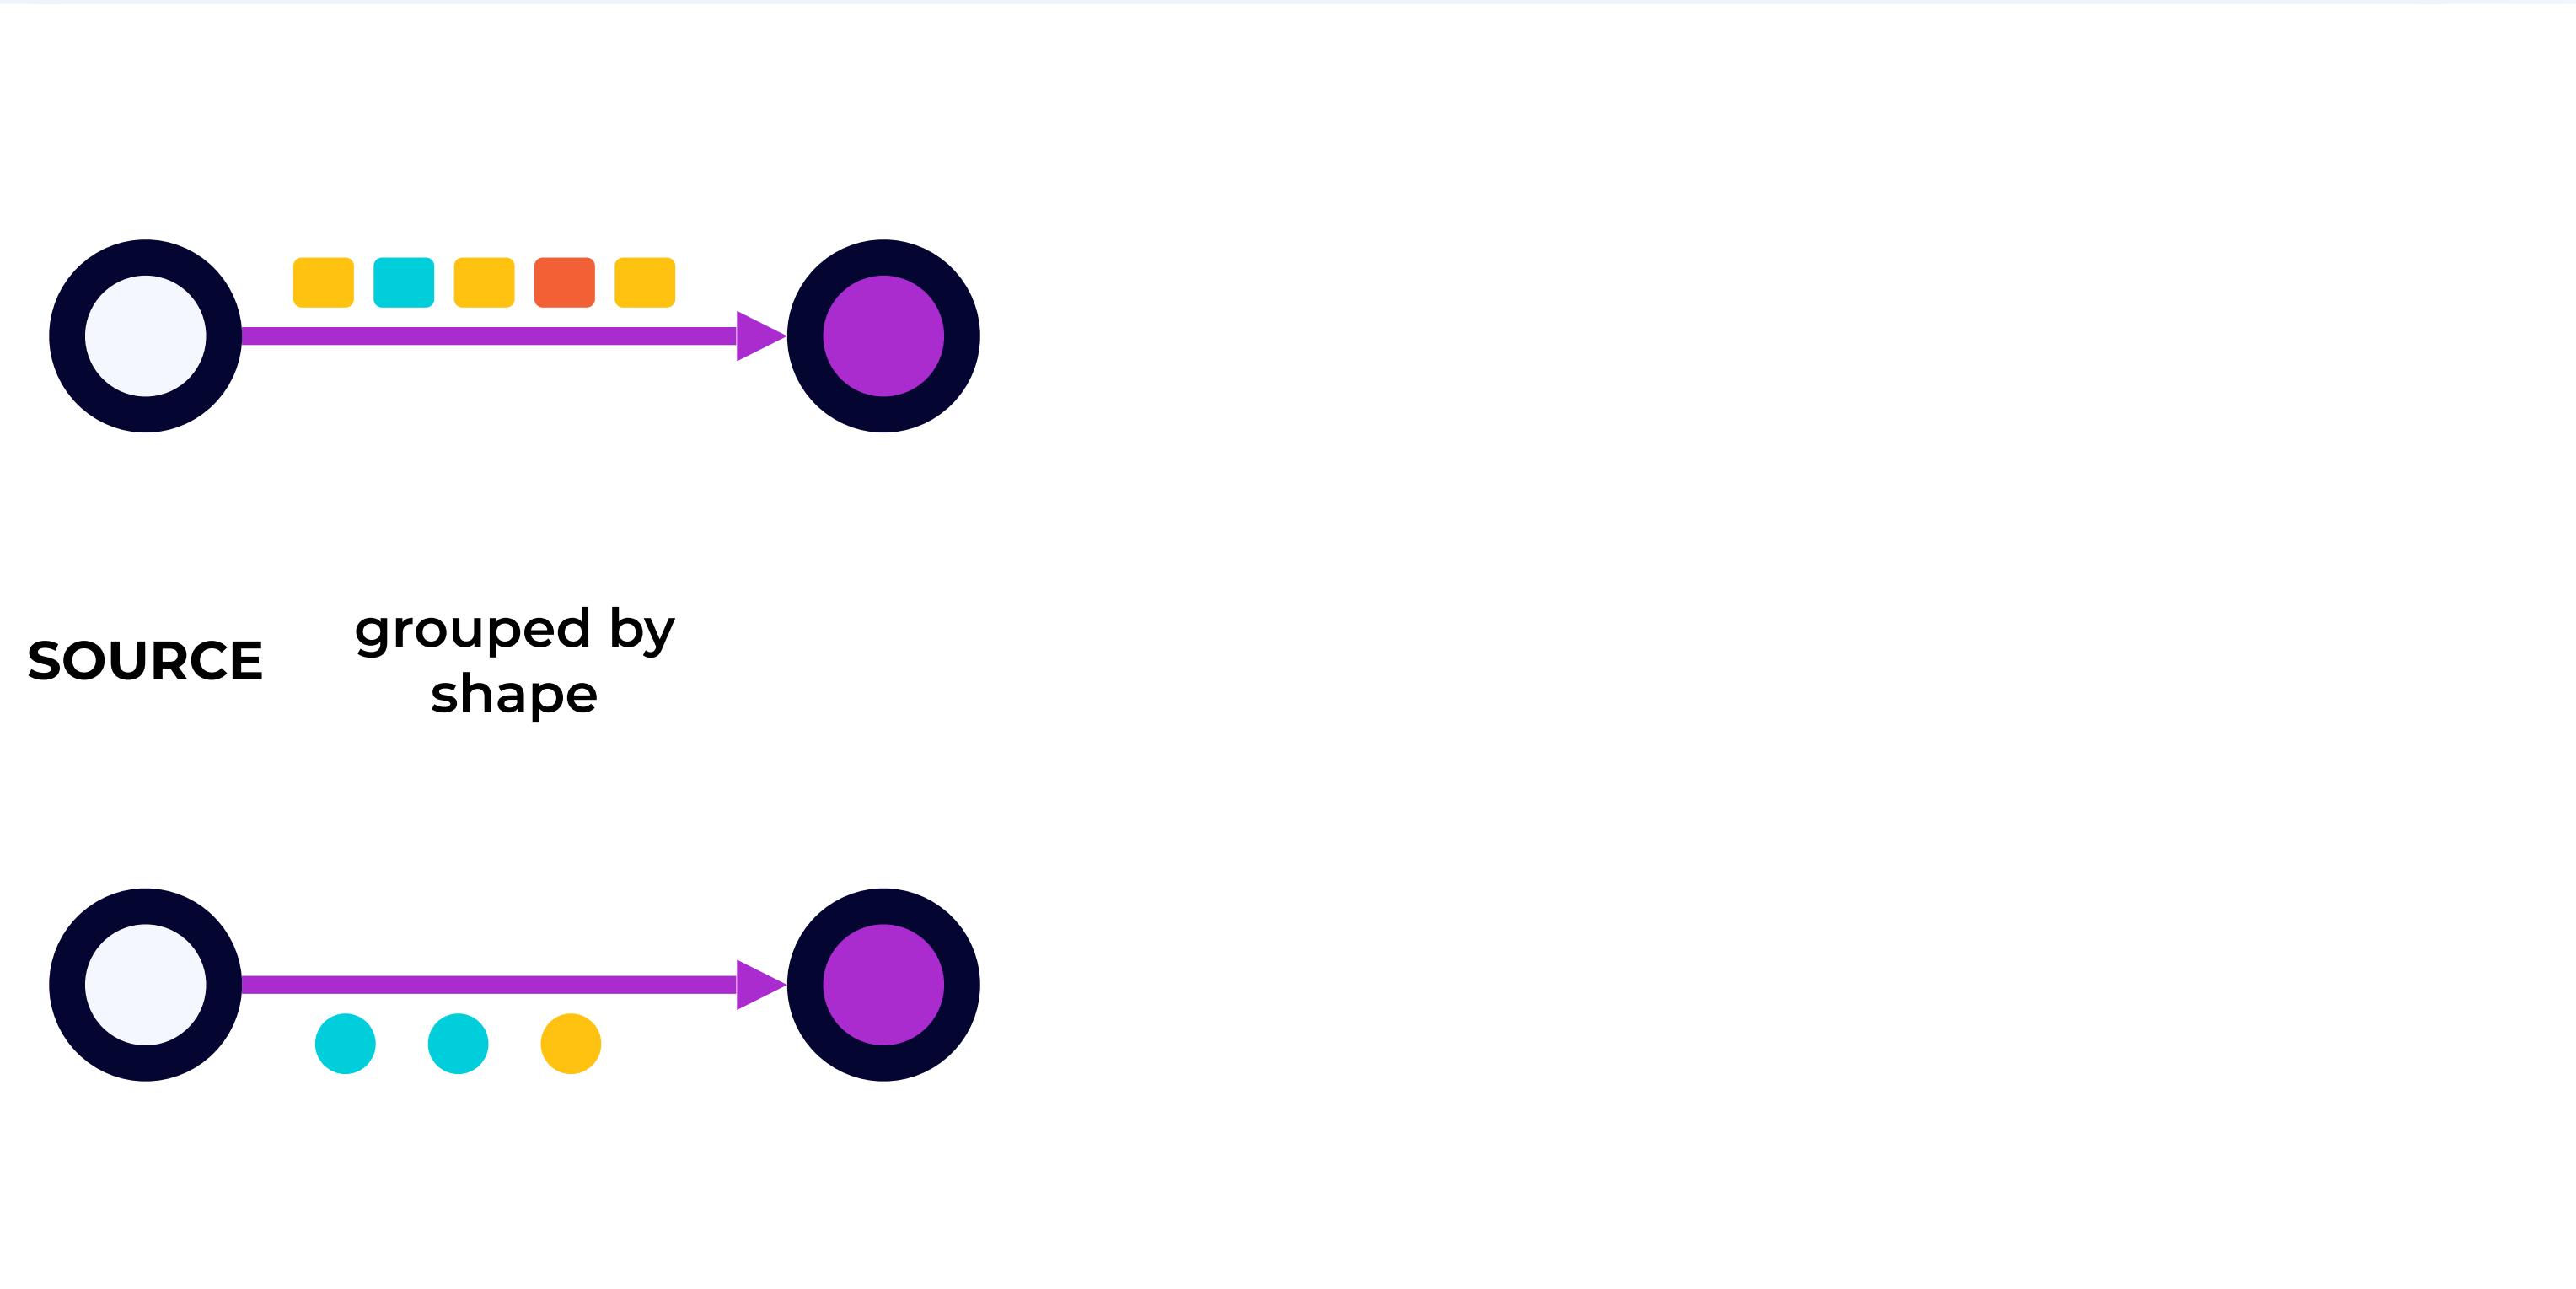
- **Parallel**
- **Forward**
- **Repartition**
- **Rebalance**



# Stream processing



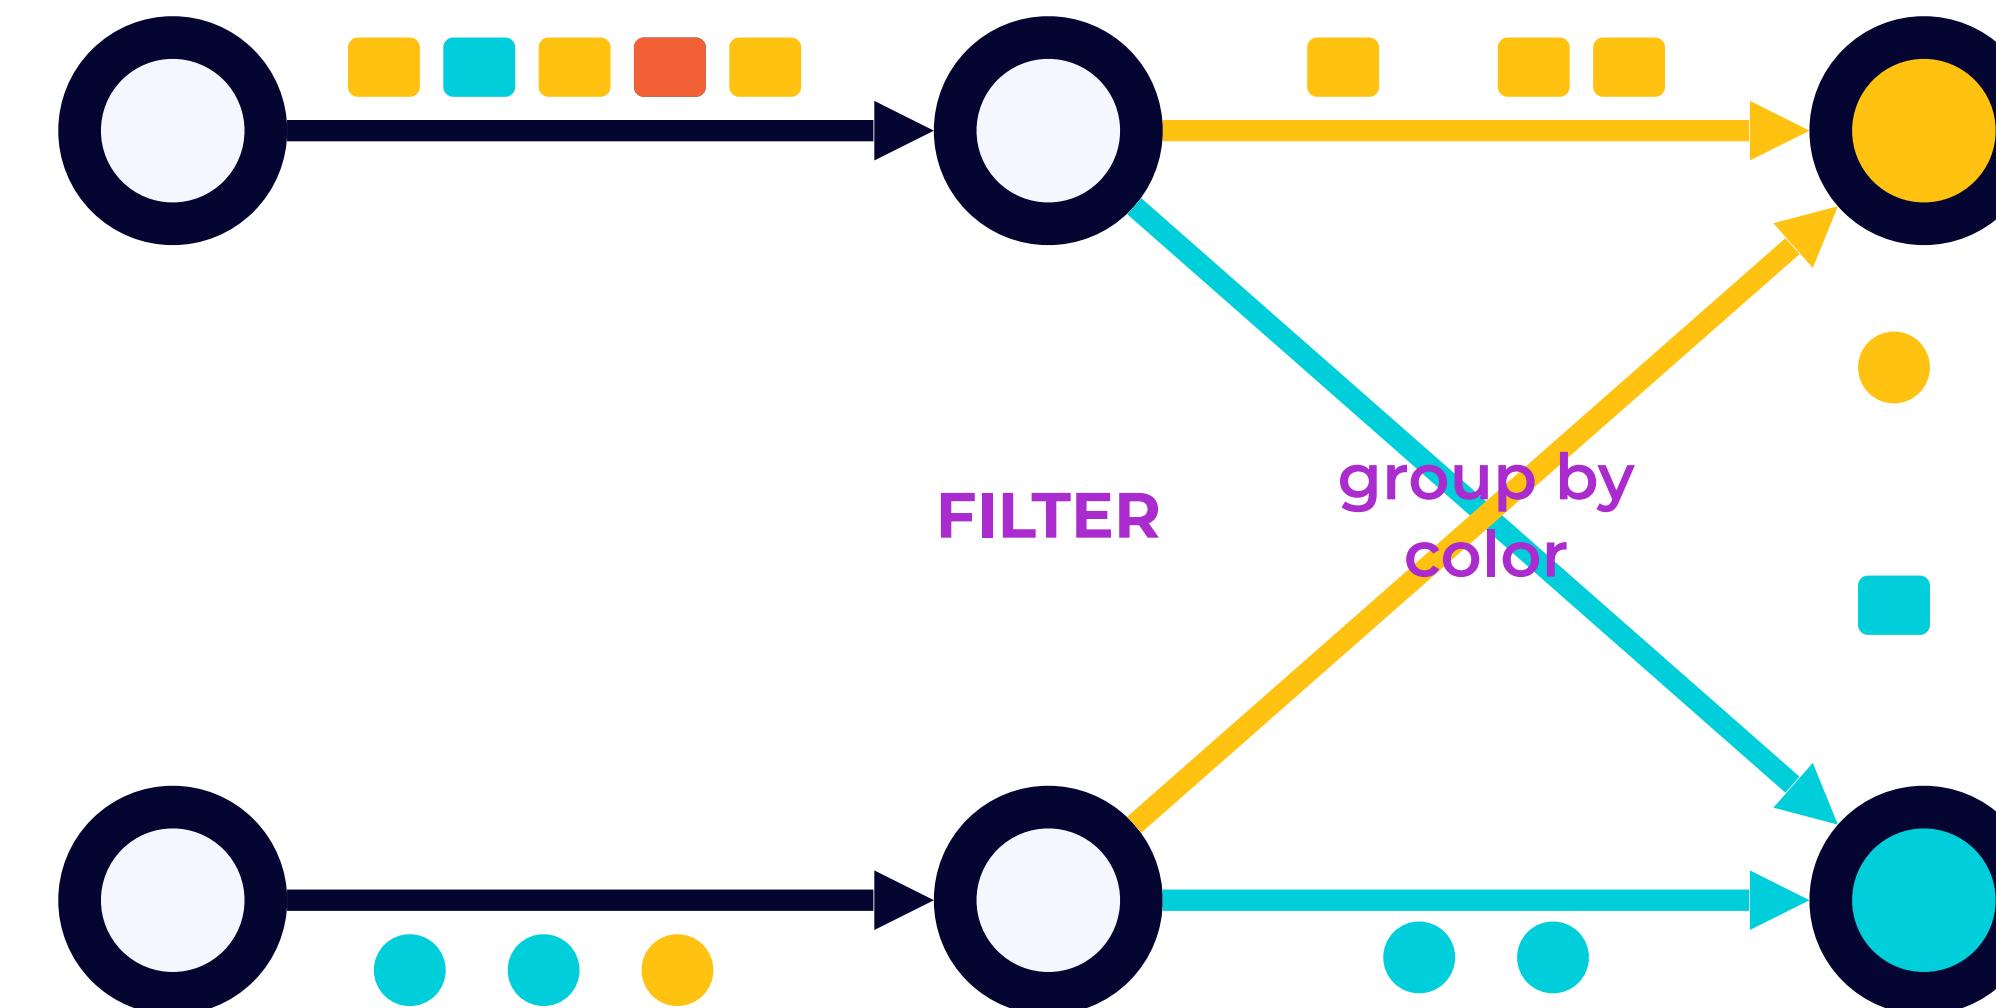
- **Parallel**
- **Forward**
- **Repartition**
- **Rebalance**



# Stream processing



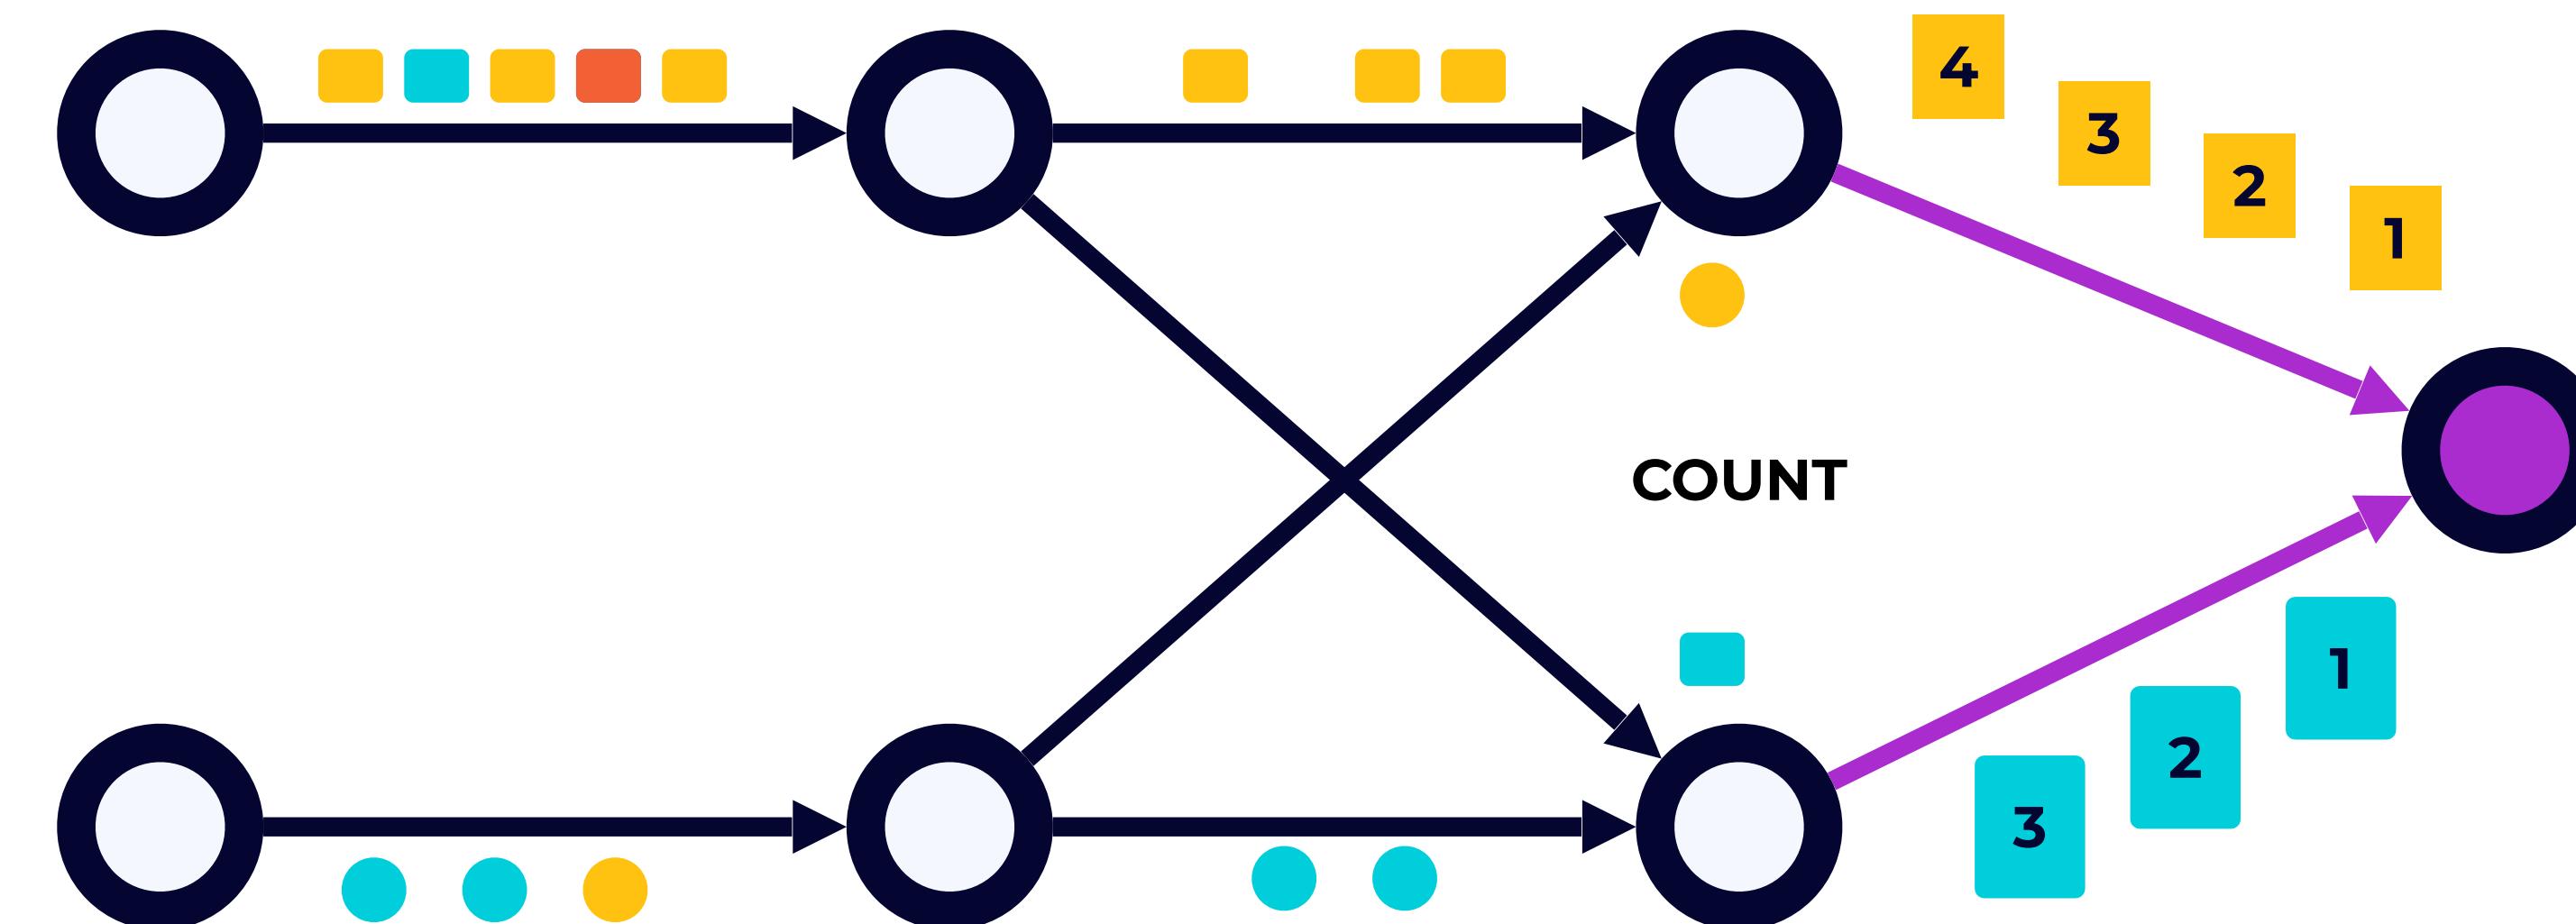
- Parallel
- Forward
- Repartition
- Rebalance



# Stream processing



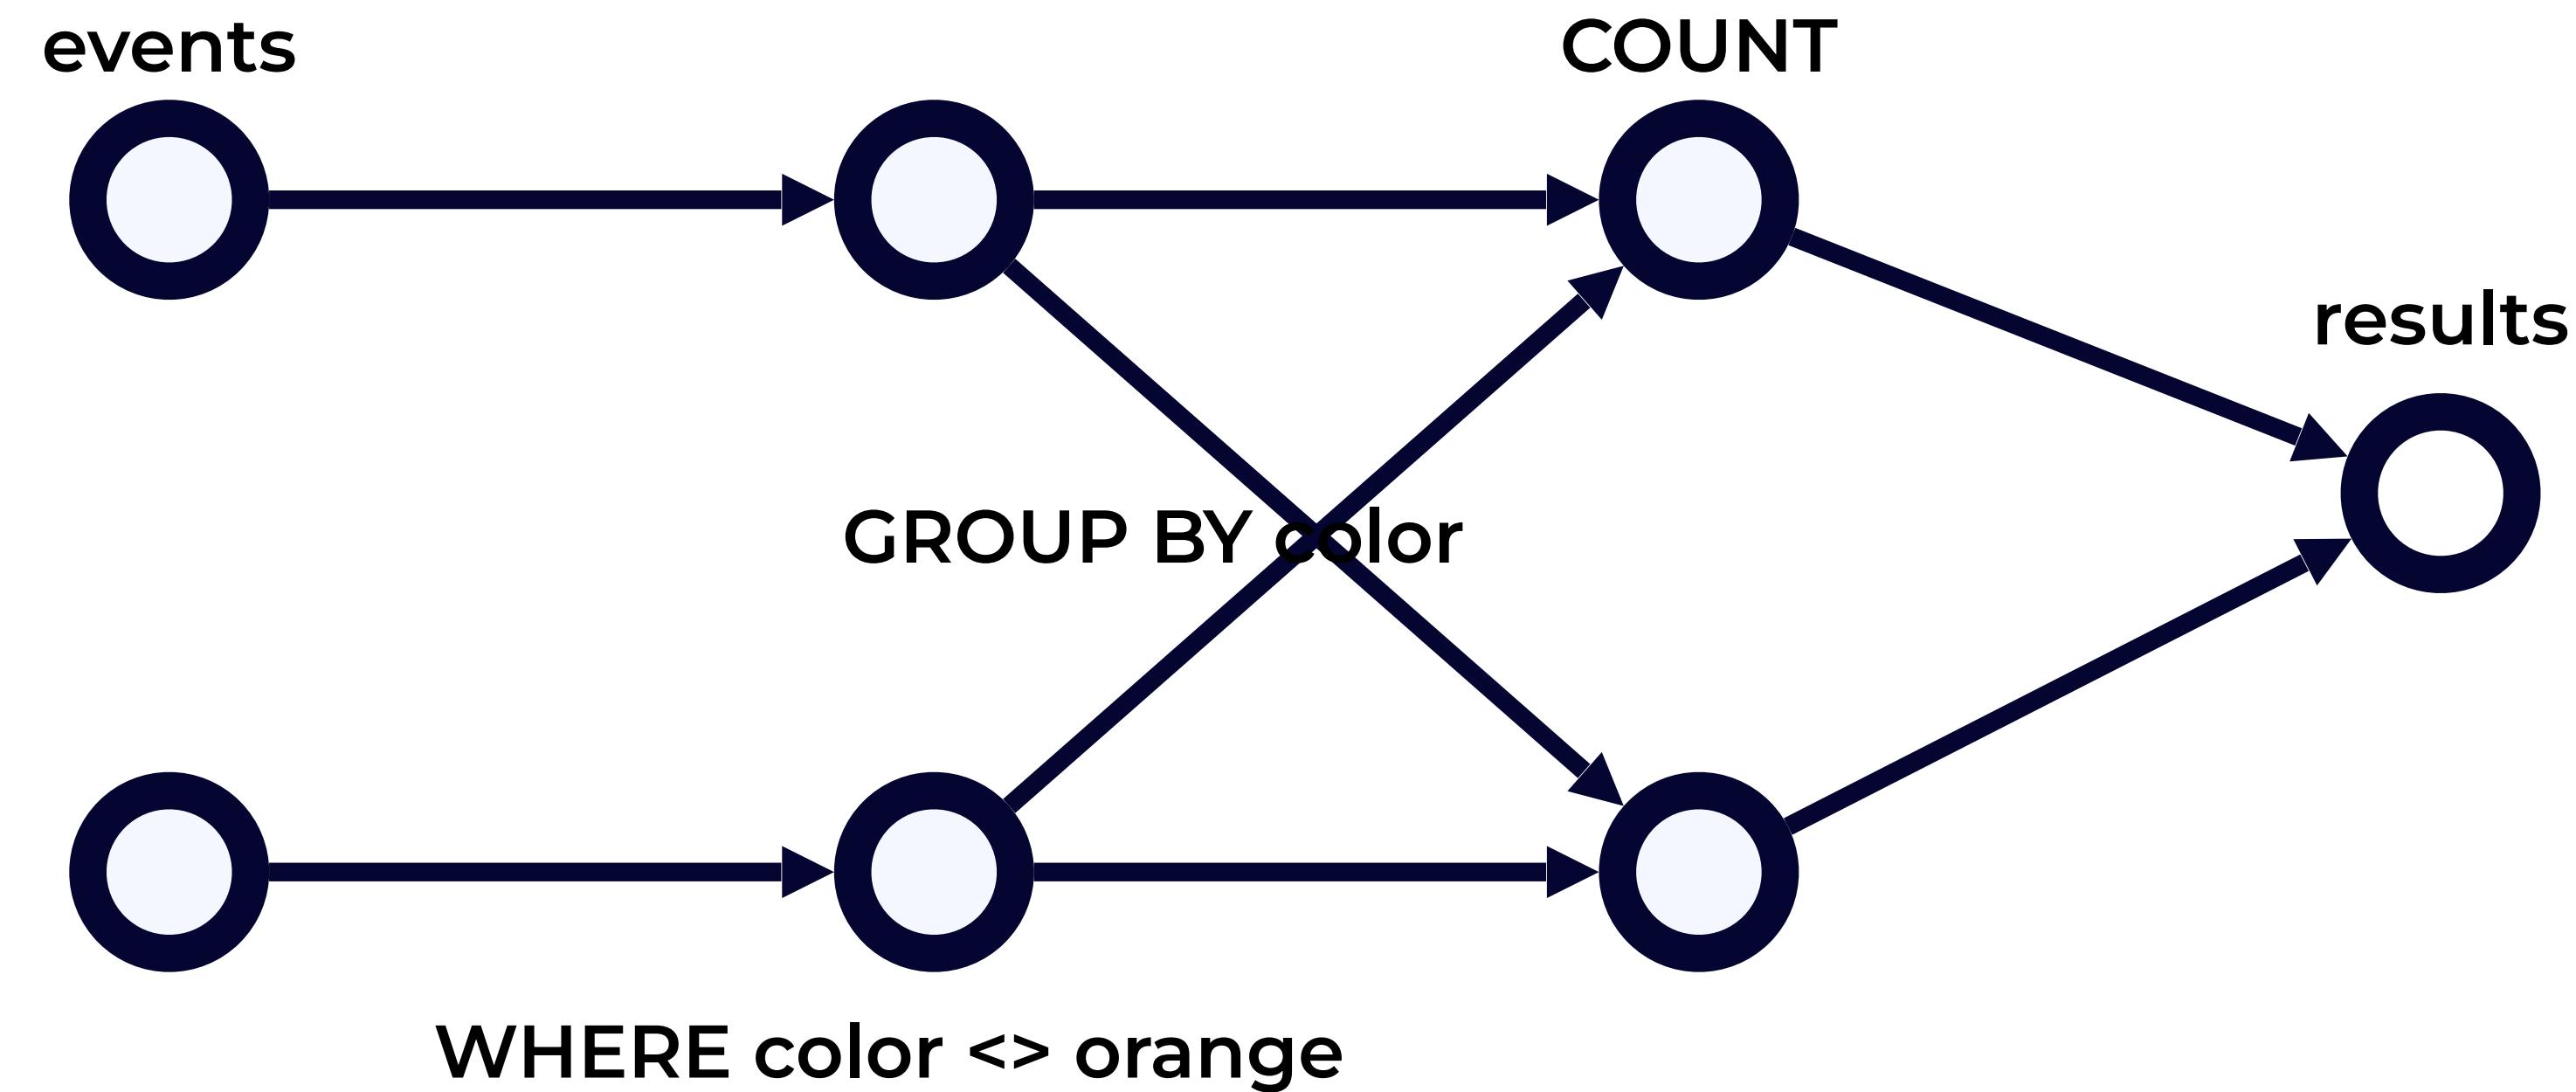
- Parallel
- Forward
- Repartition
- Rebalance



# Stream processing with SQL



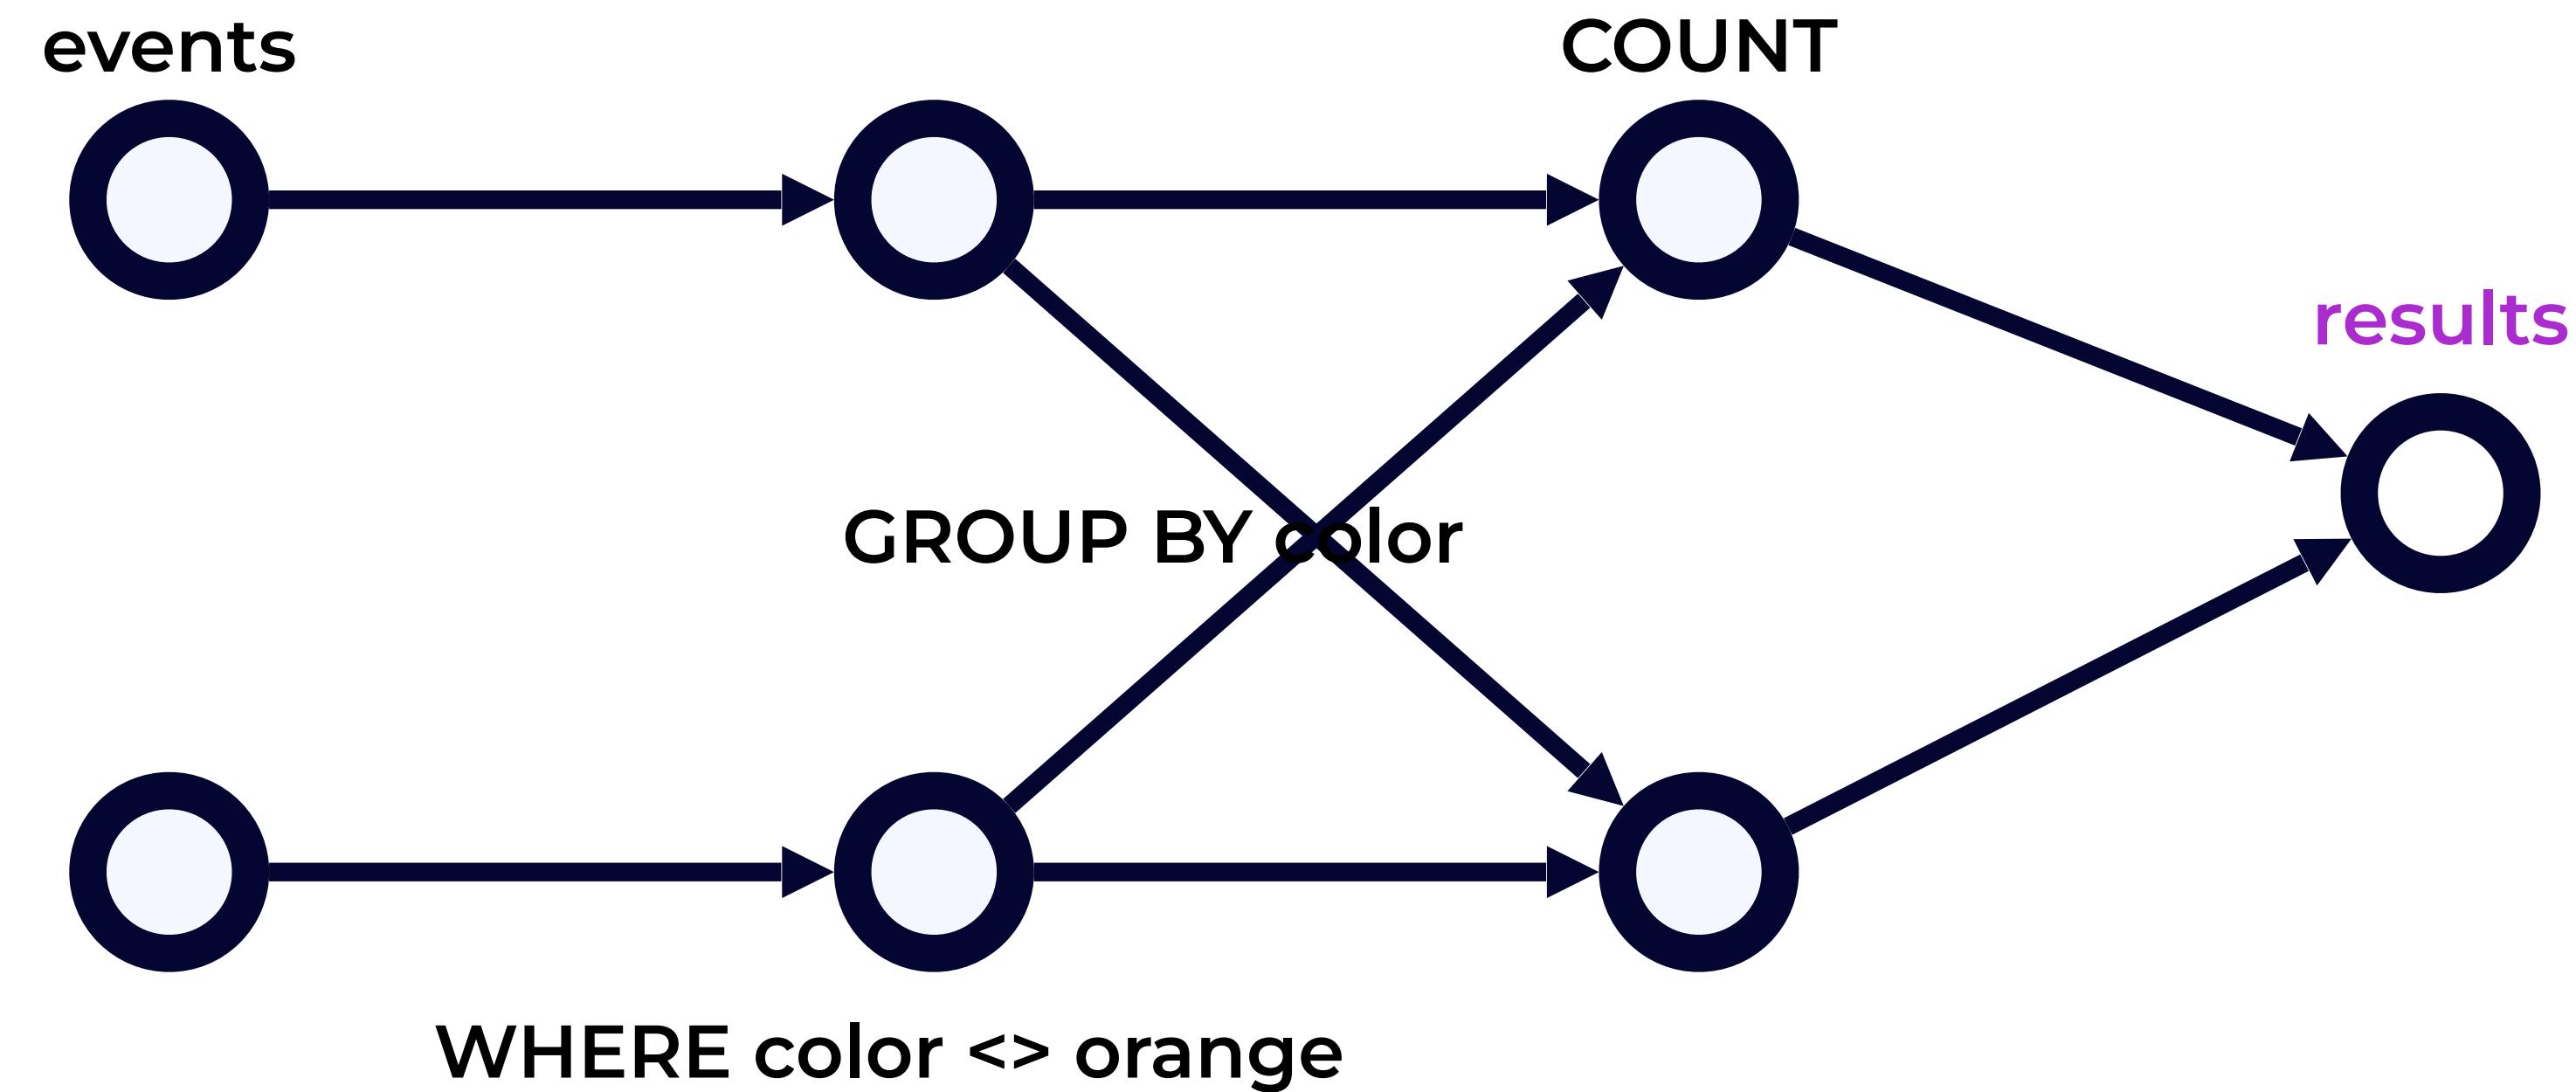
```
INSERT INTO results
SELECT color, COUNT(*)
FROM events
WHERE color <> orange
GROUP BY color;
```



# Stream processing with SQL



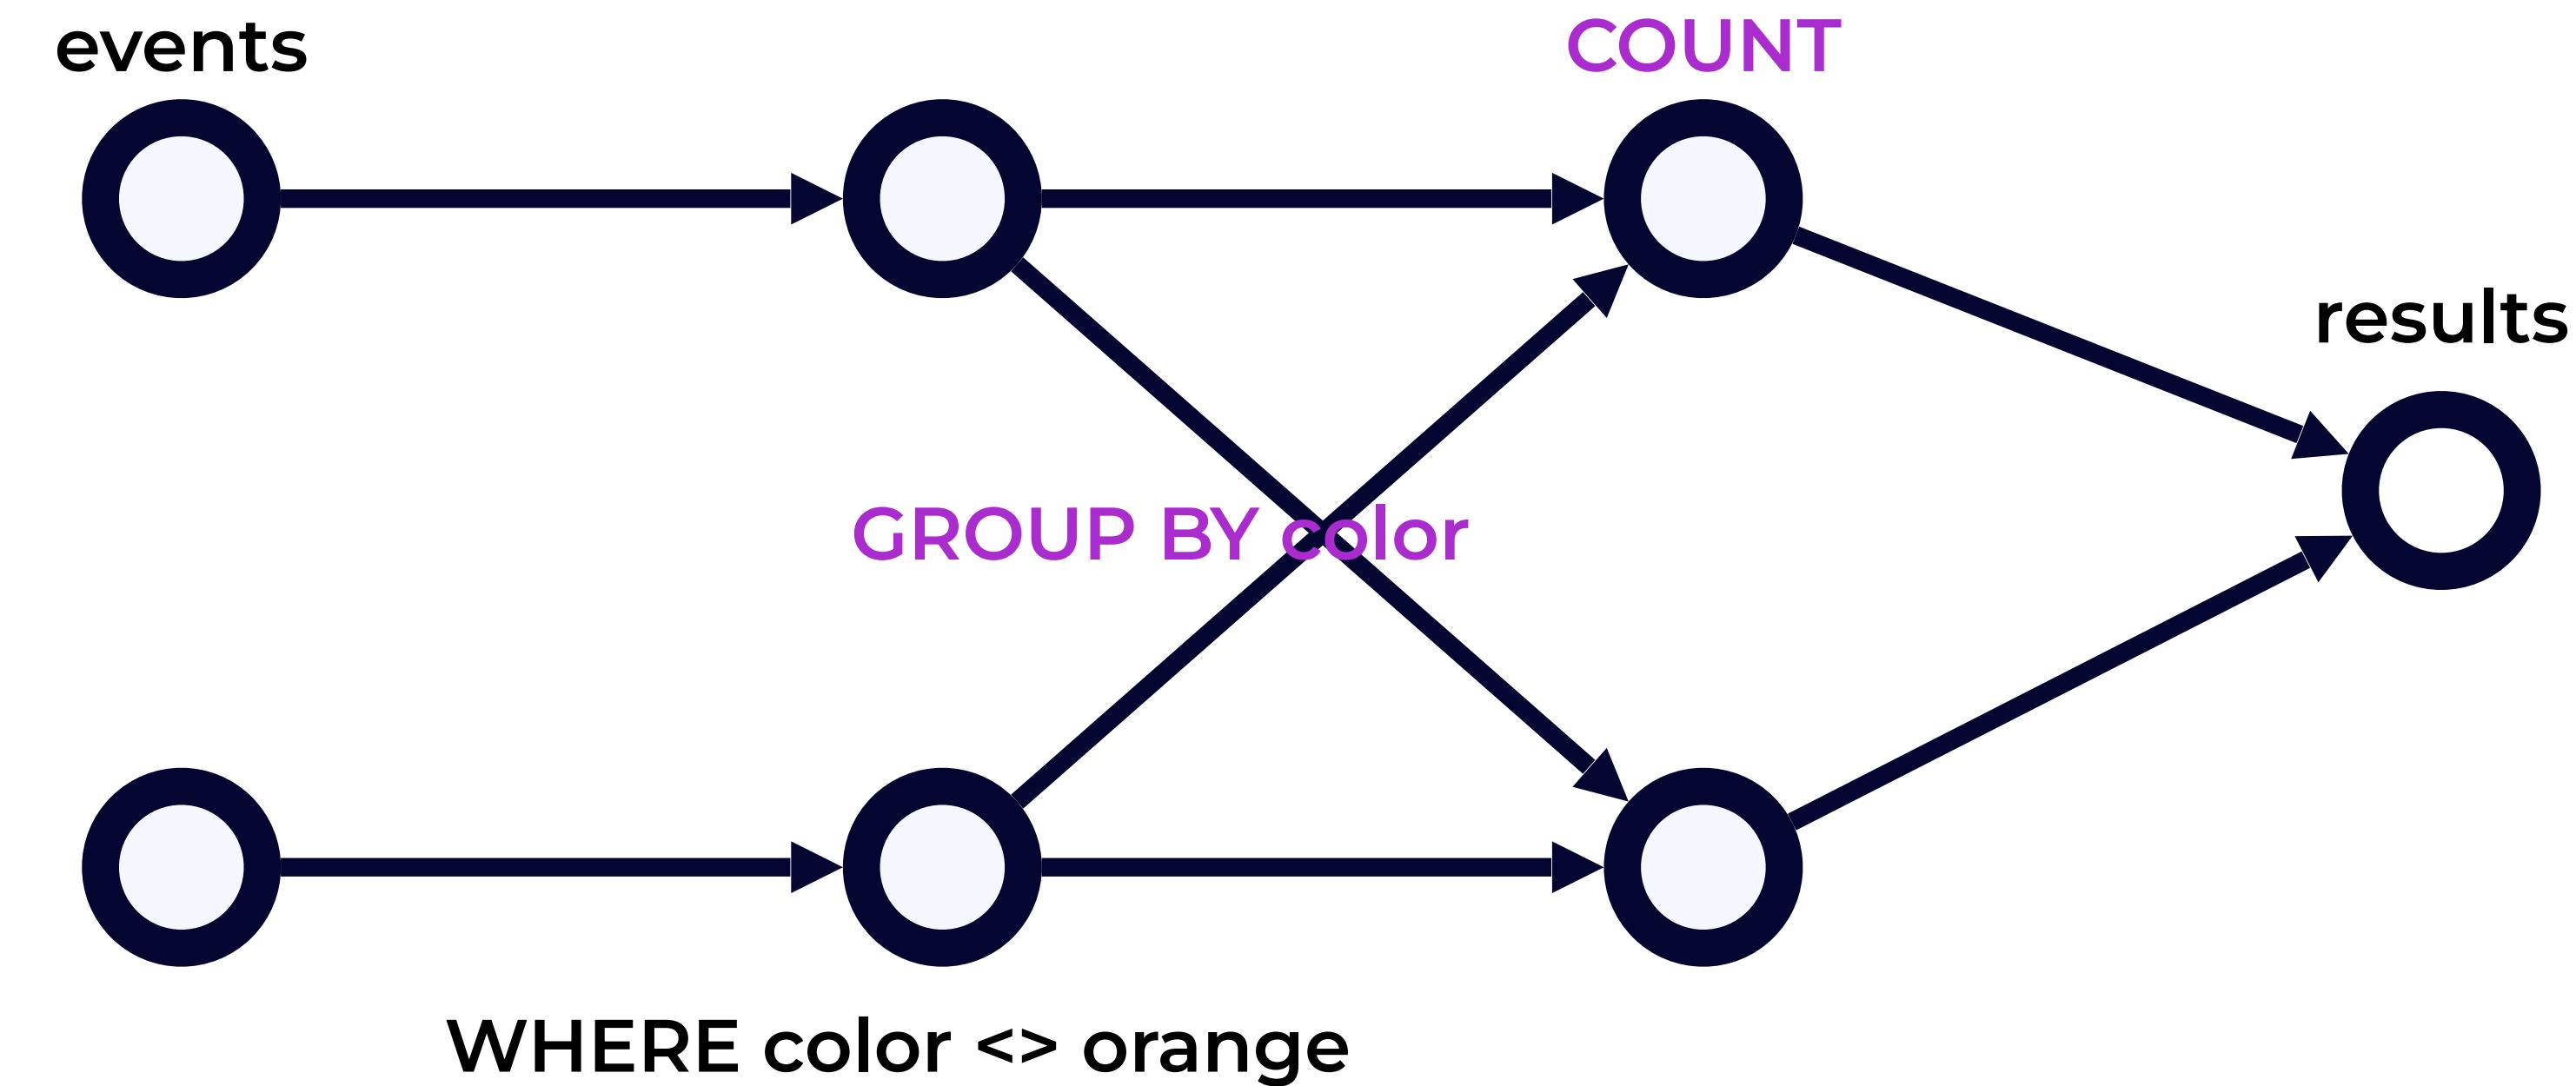
```
INSERT INTO results
SELECT color, COUNT(*)
FROM events
WHERE color <> orange
GROUP BY color;
```



# Stream processing with SQL



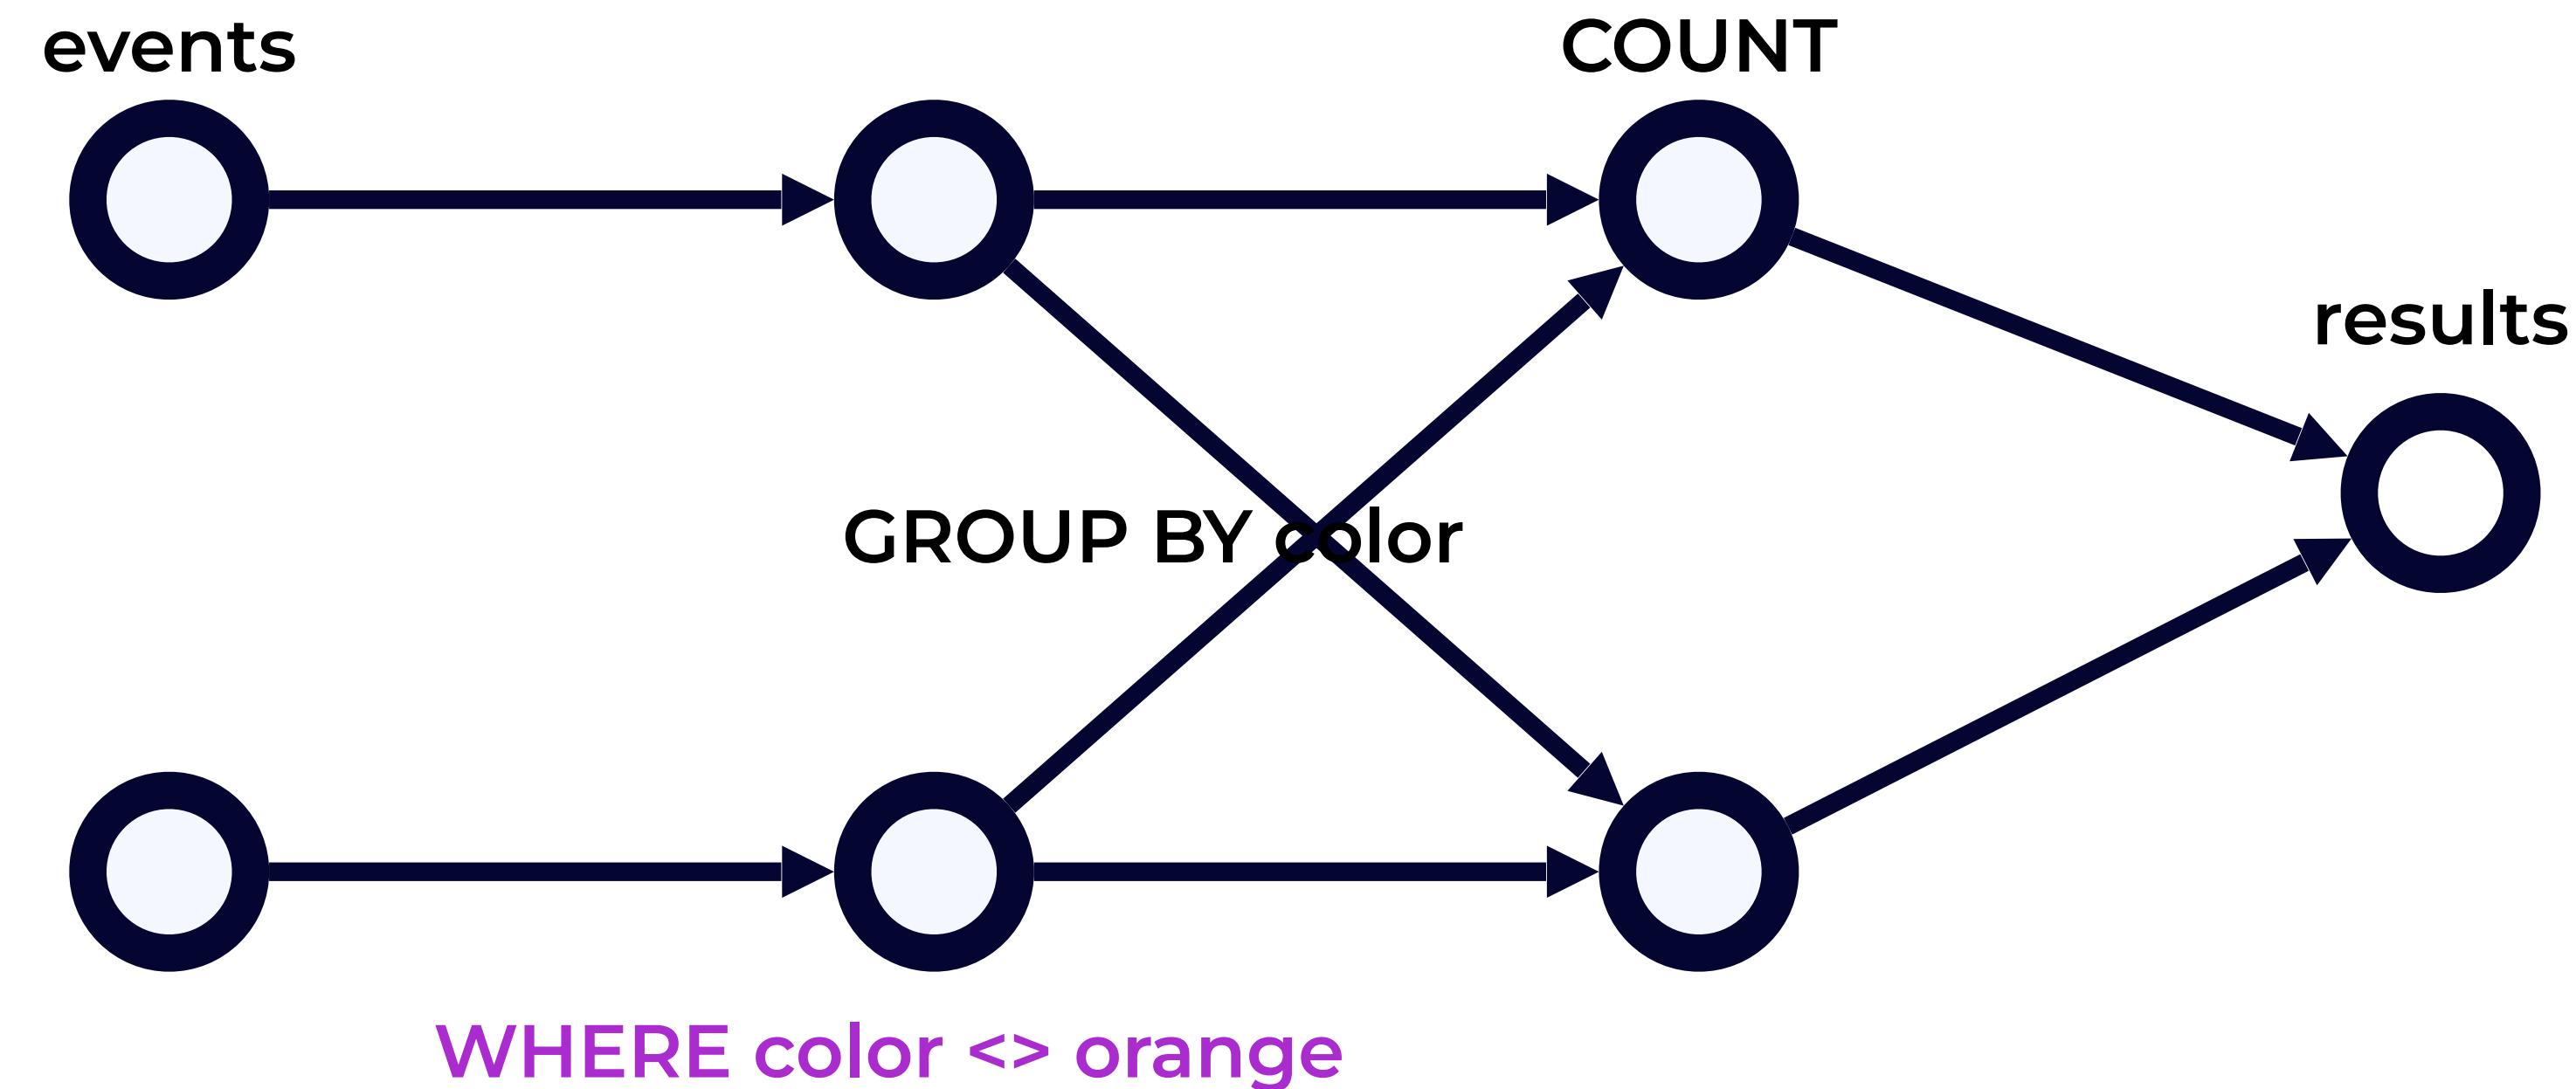
```
INSERT INTO results
SELECT color, COUNT(*)
FROM events
WHERE color <> orange
GROUP BY color;
```



# Stream processing with SQL



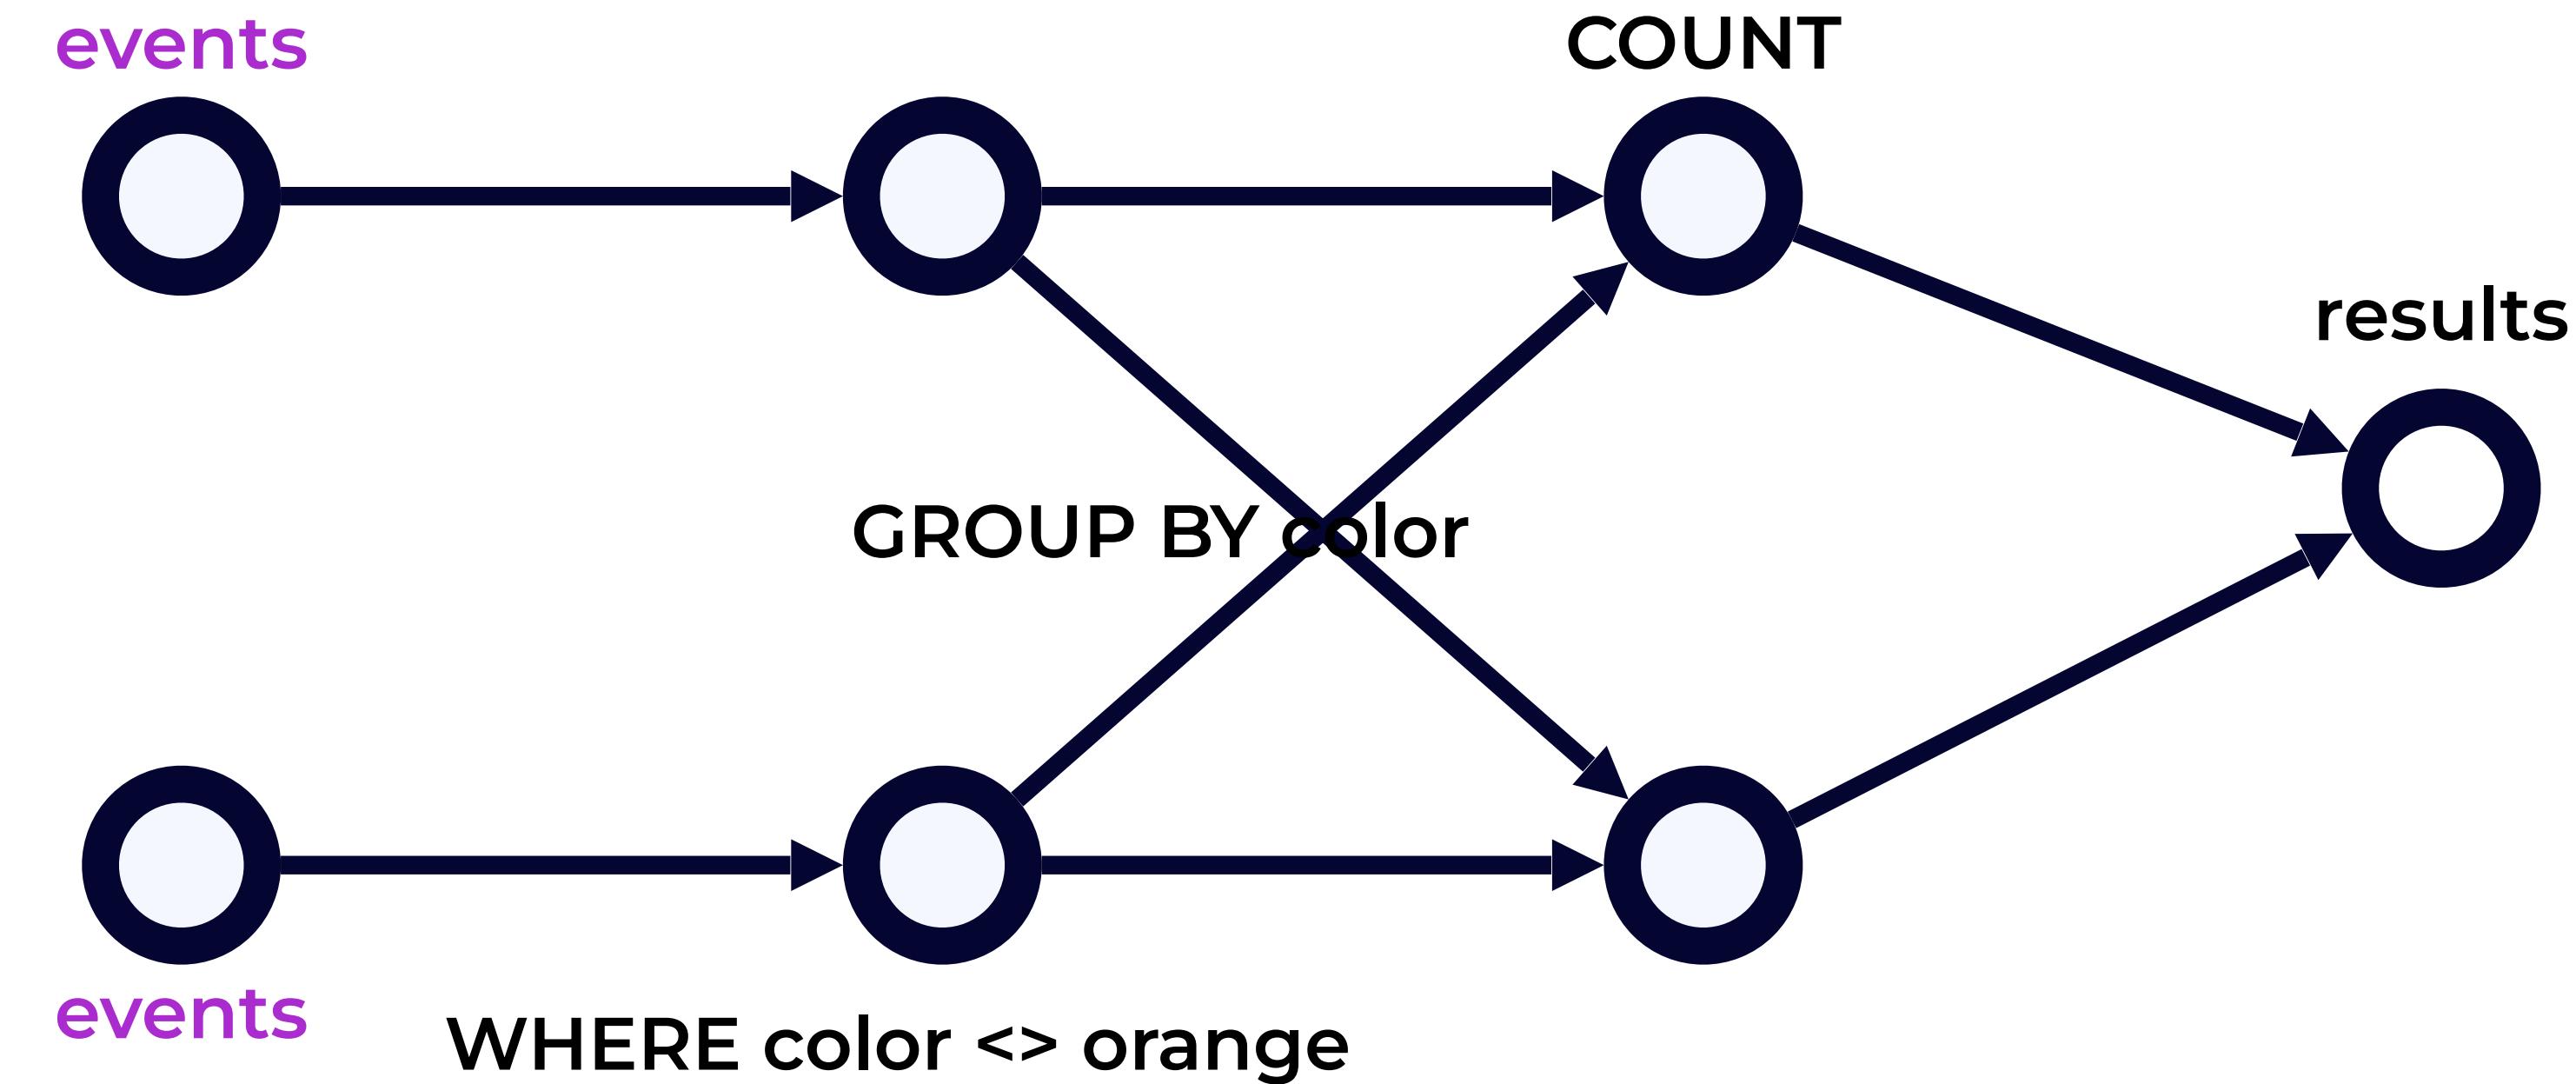
```
INSERT INTO results
SELECT color, COUNT(*)
FROM events
WHERE color <> orange
GROUP BY color;
```



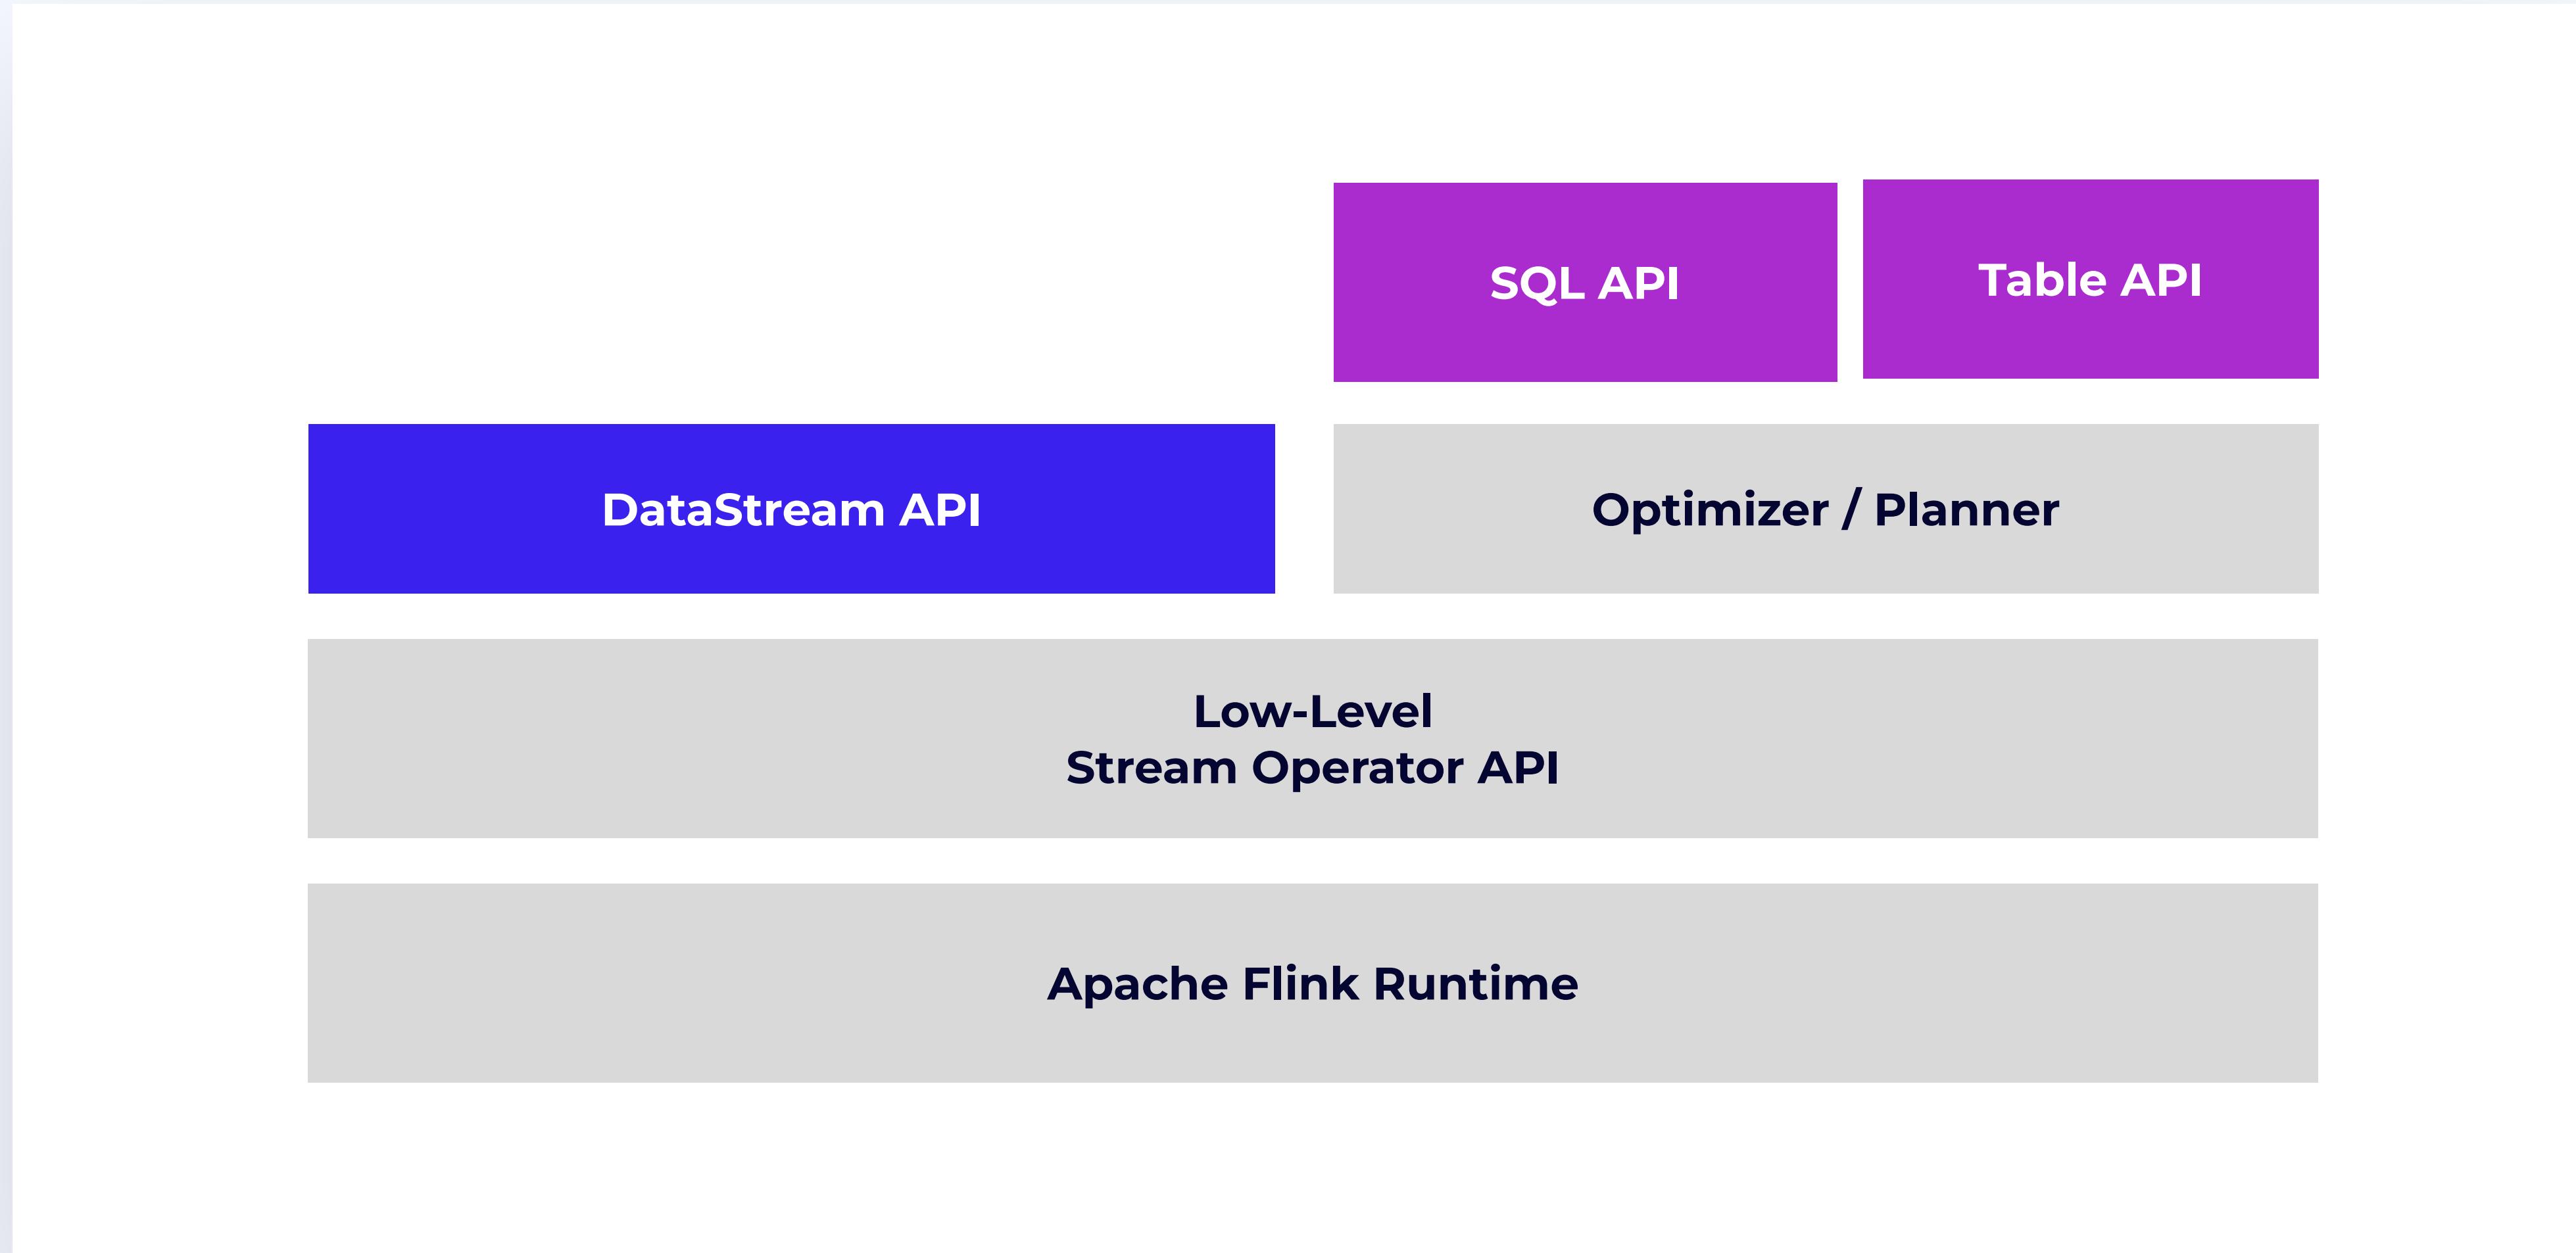
# Stream processing with SQL



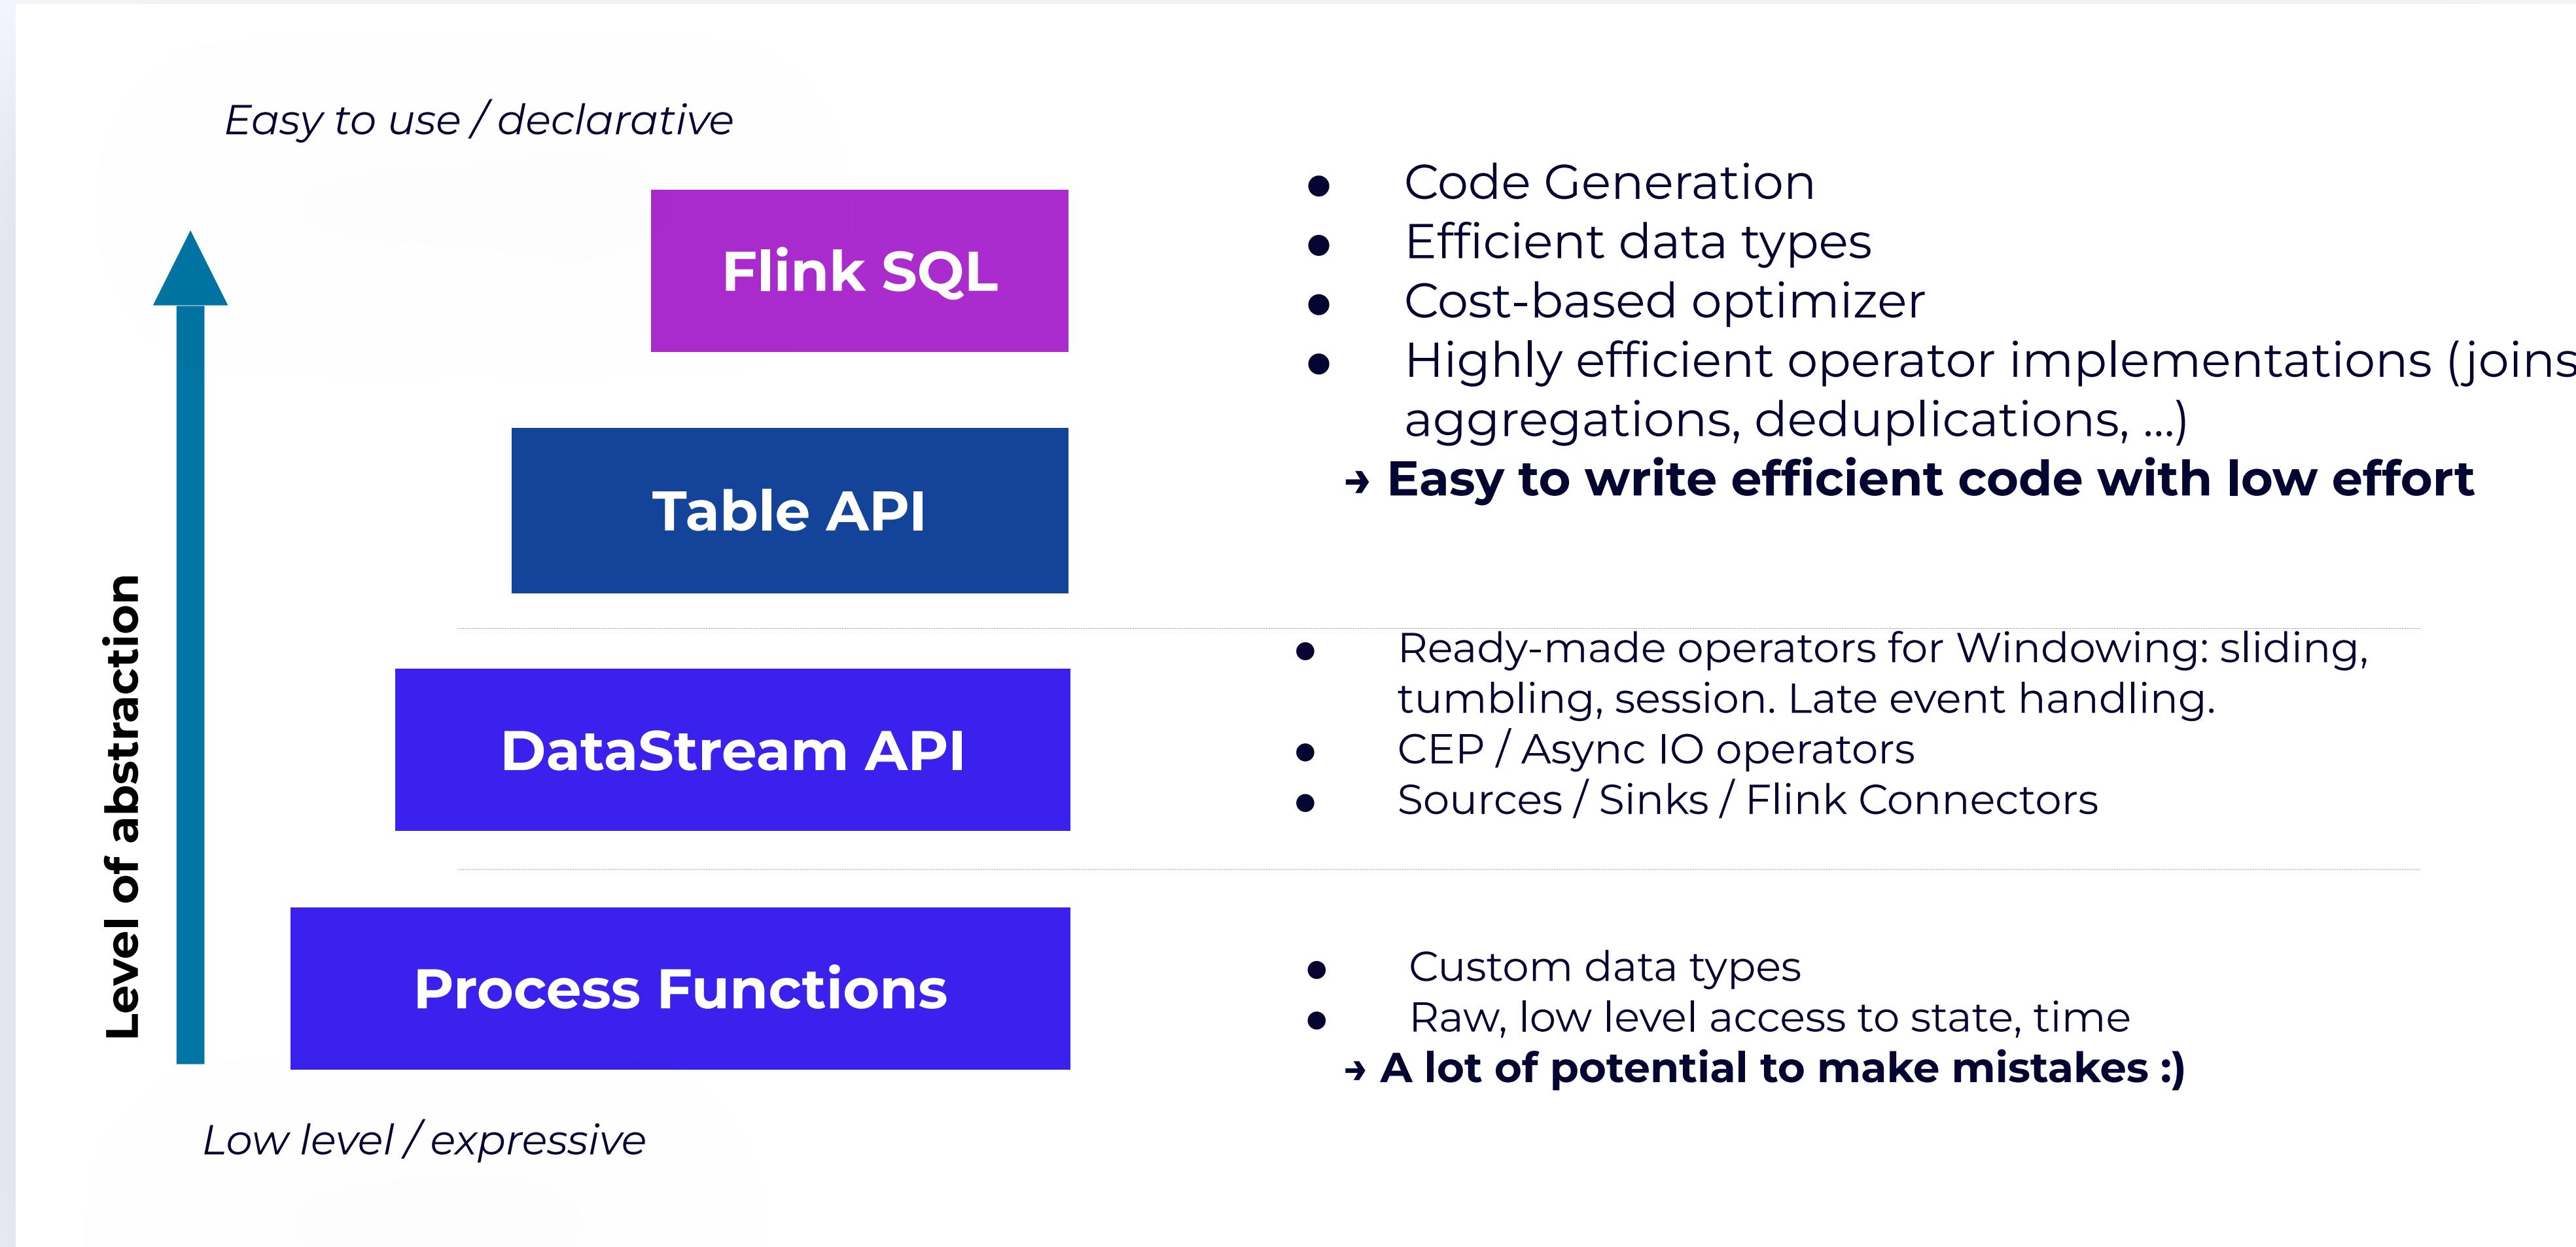
```
INSERT INTO results
SELECT color, COUNT(*)
FROM events
WHERE color != orange
GROUP BY color;
```



# Flink's APIs

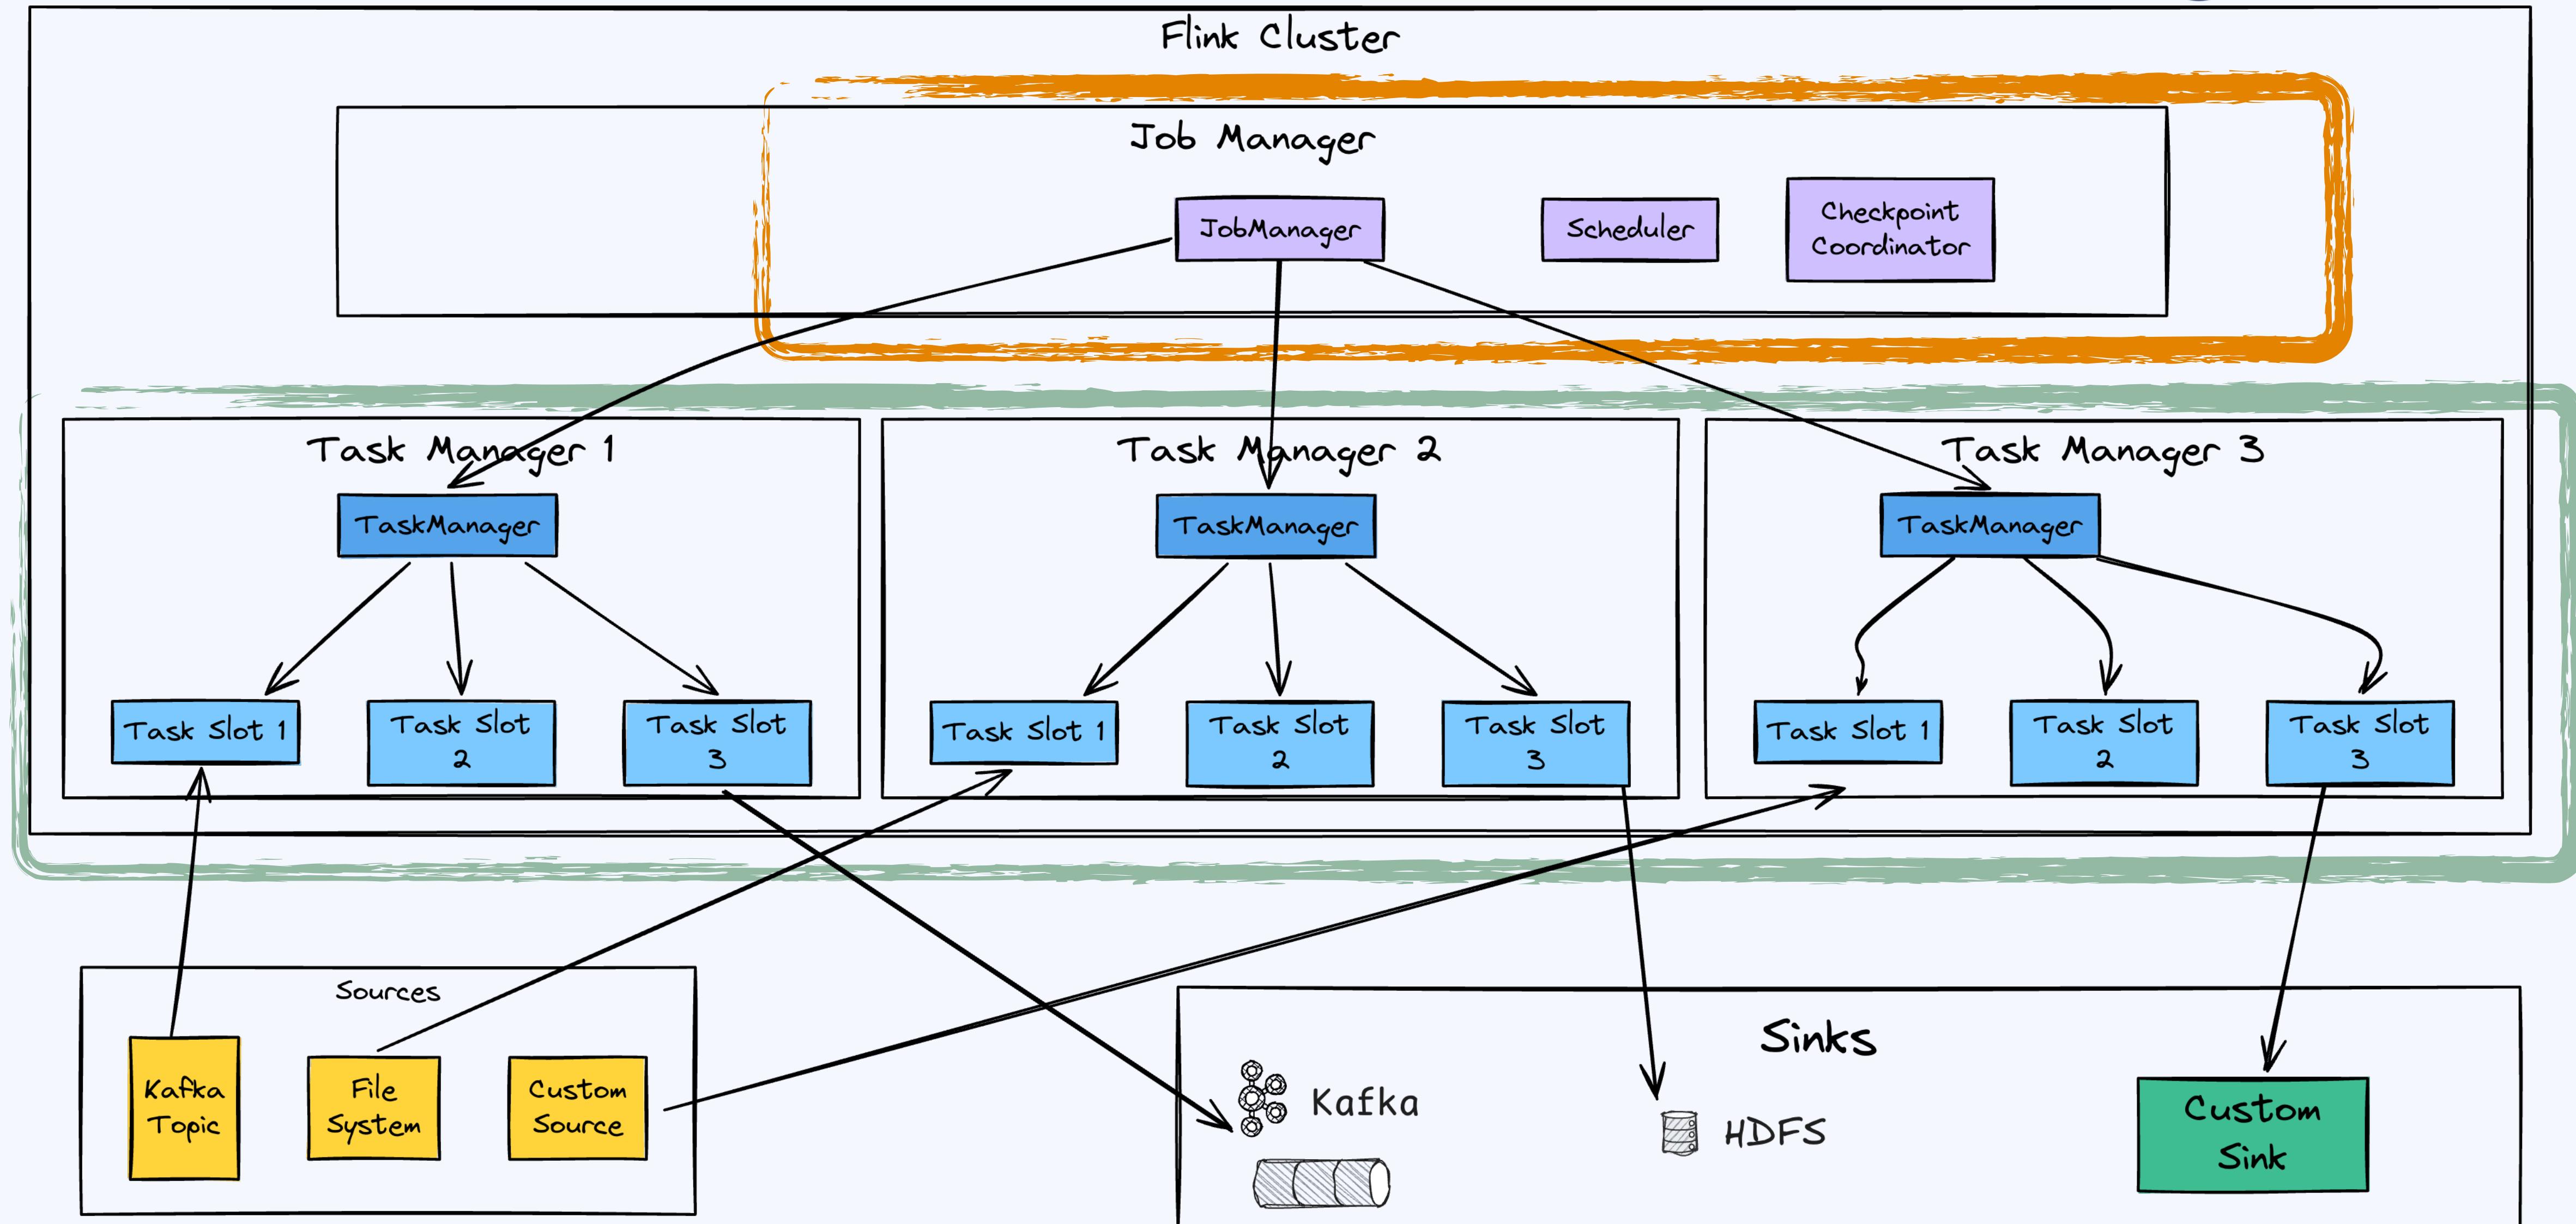


# Flink's APIs: mix & match





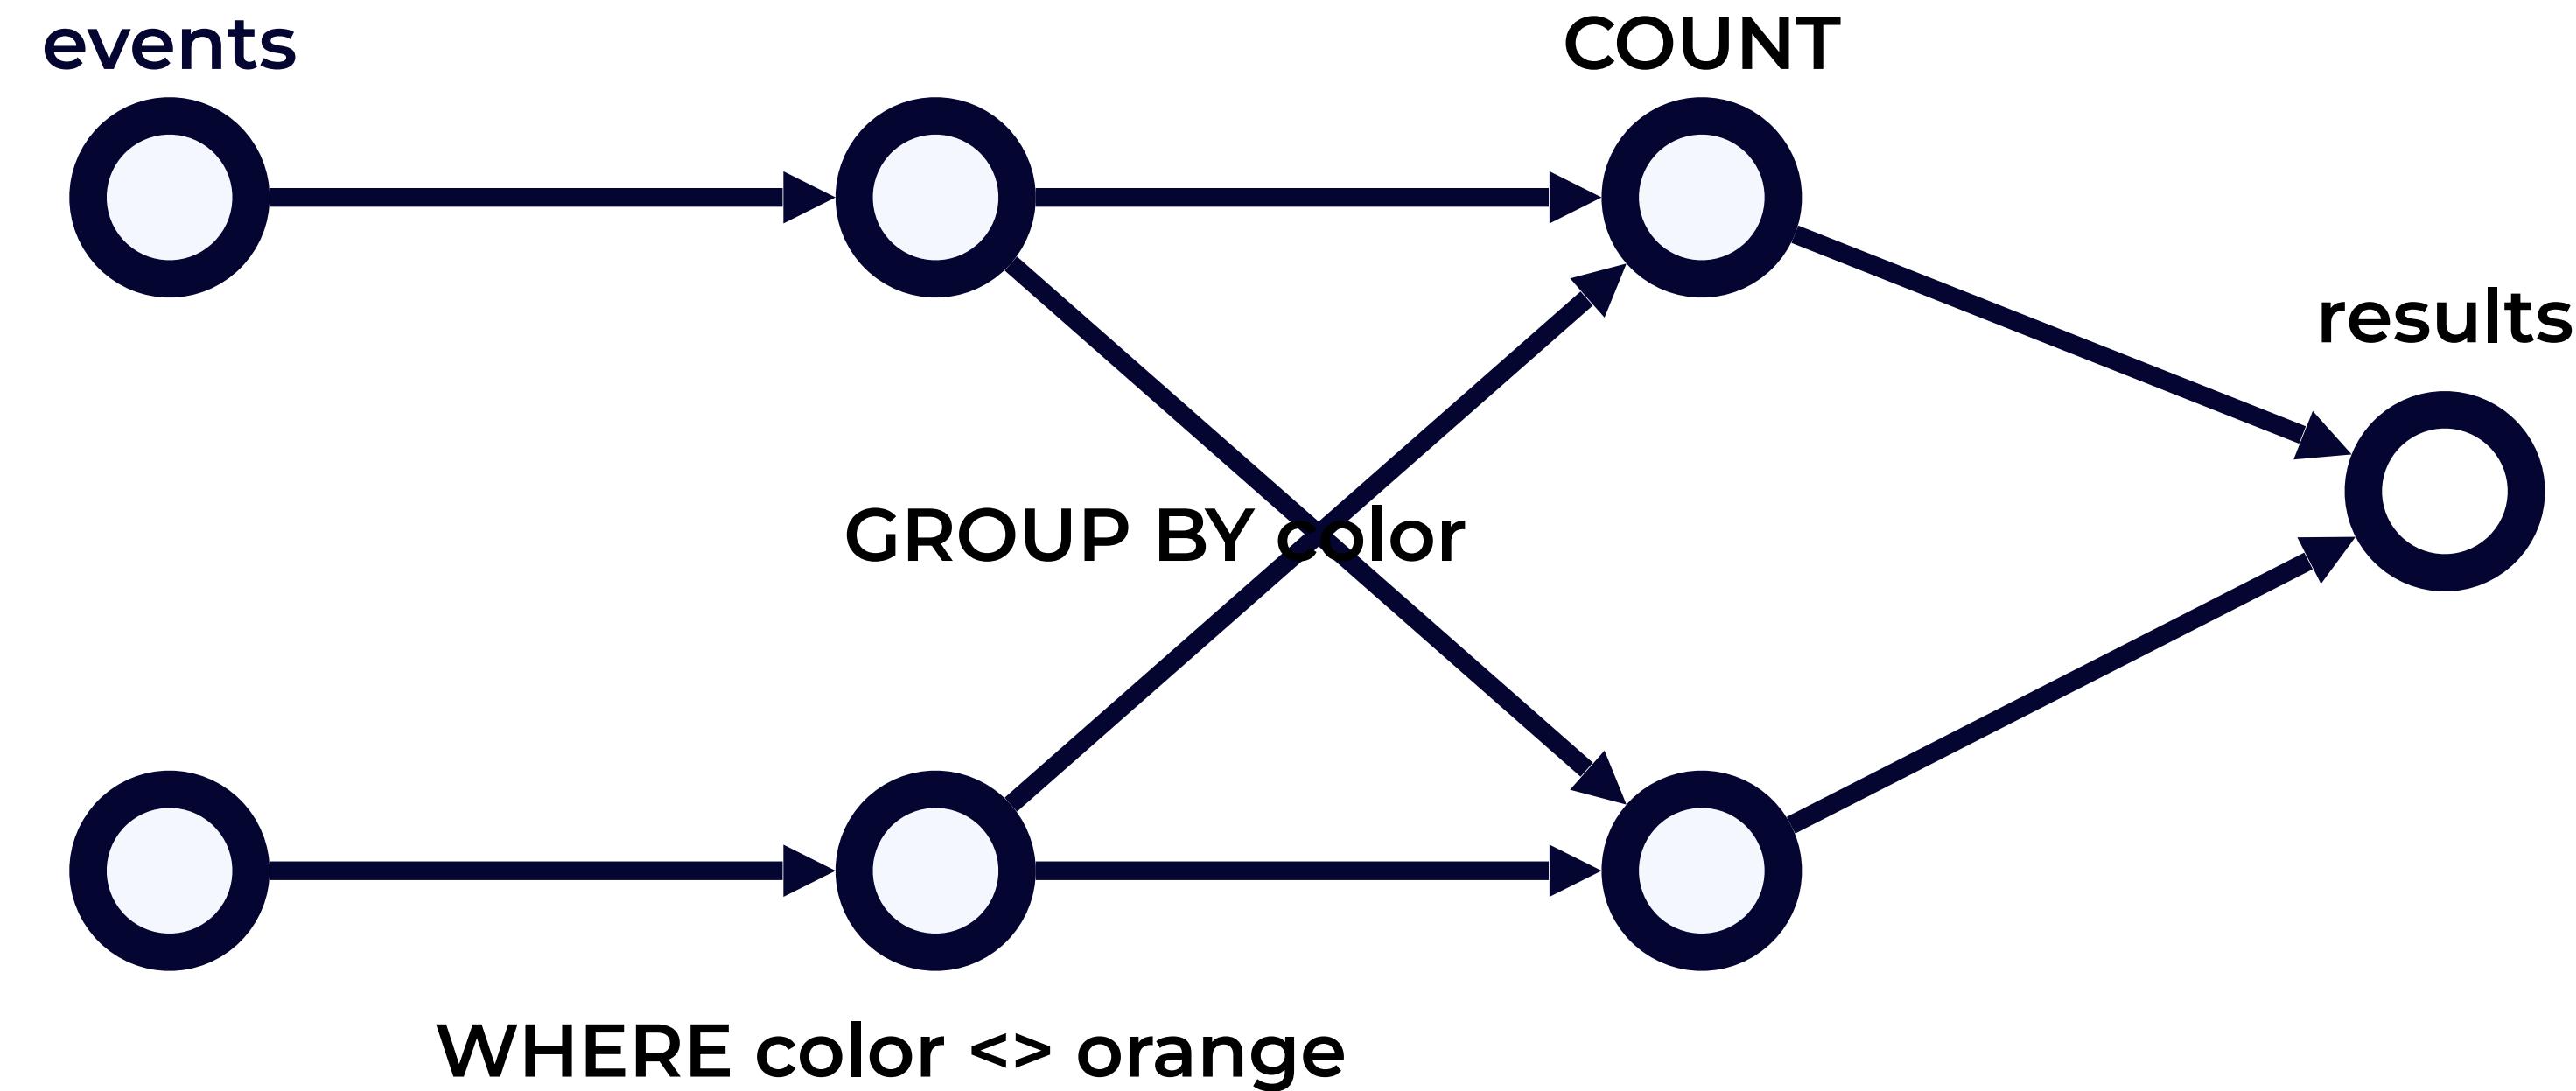
# State Management



# Stateful stream processing with Flink SQL



```
INSERT INTO results
SELECT color, COUNT(*)
FROM events
WHERE color <> orange
GROUP BY color;
```

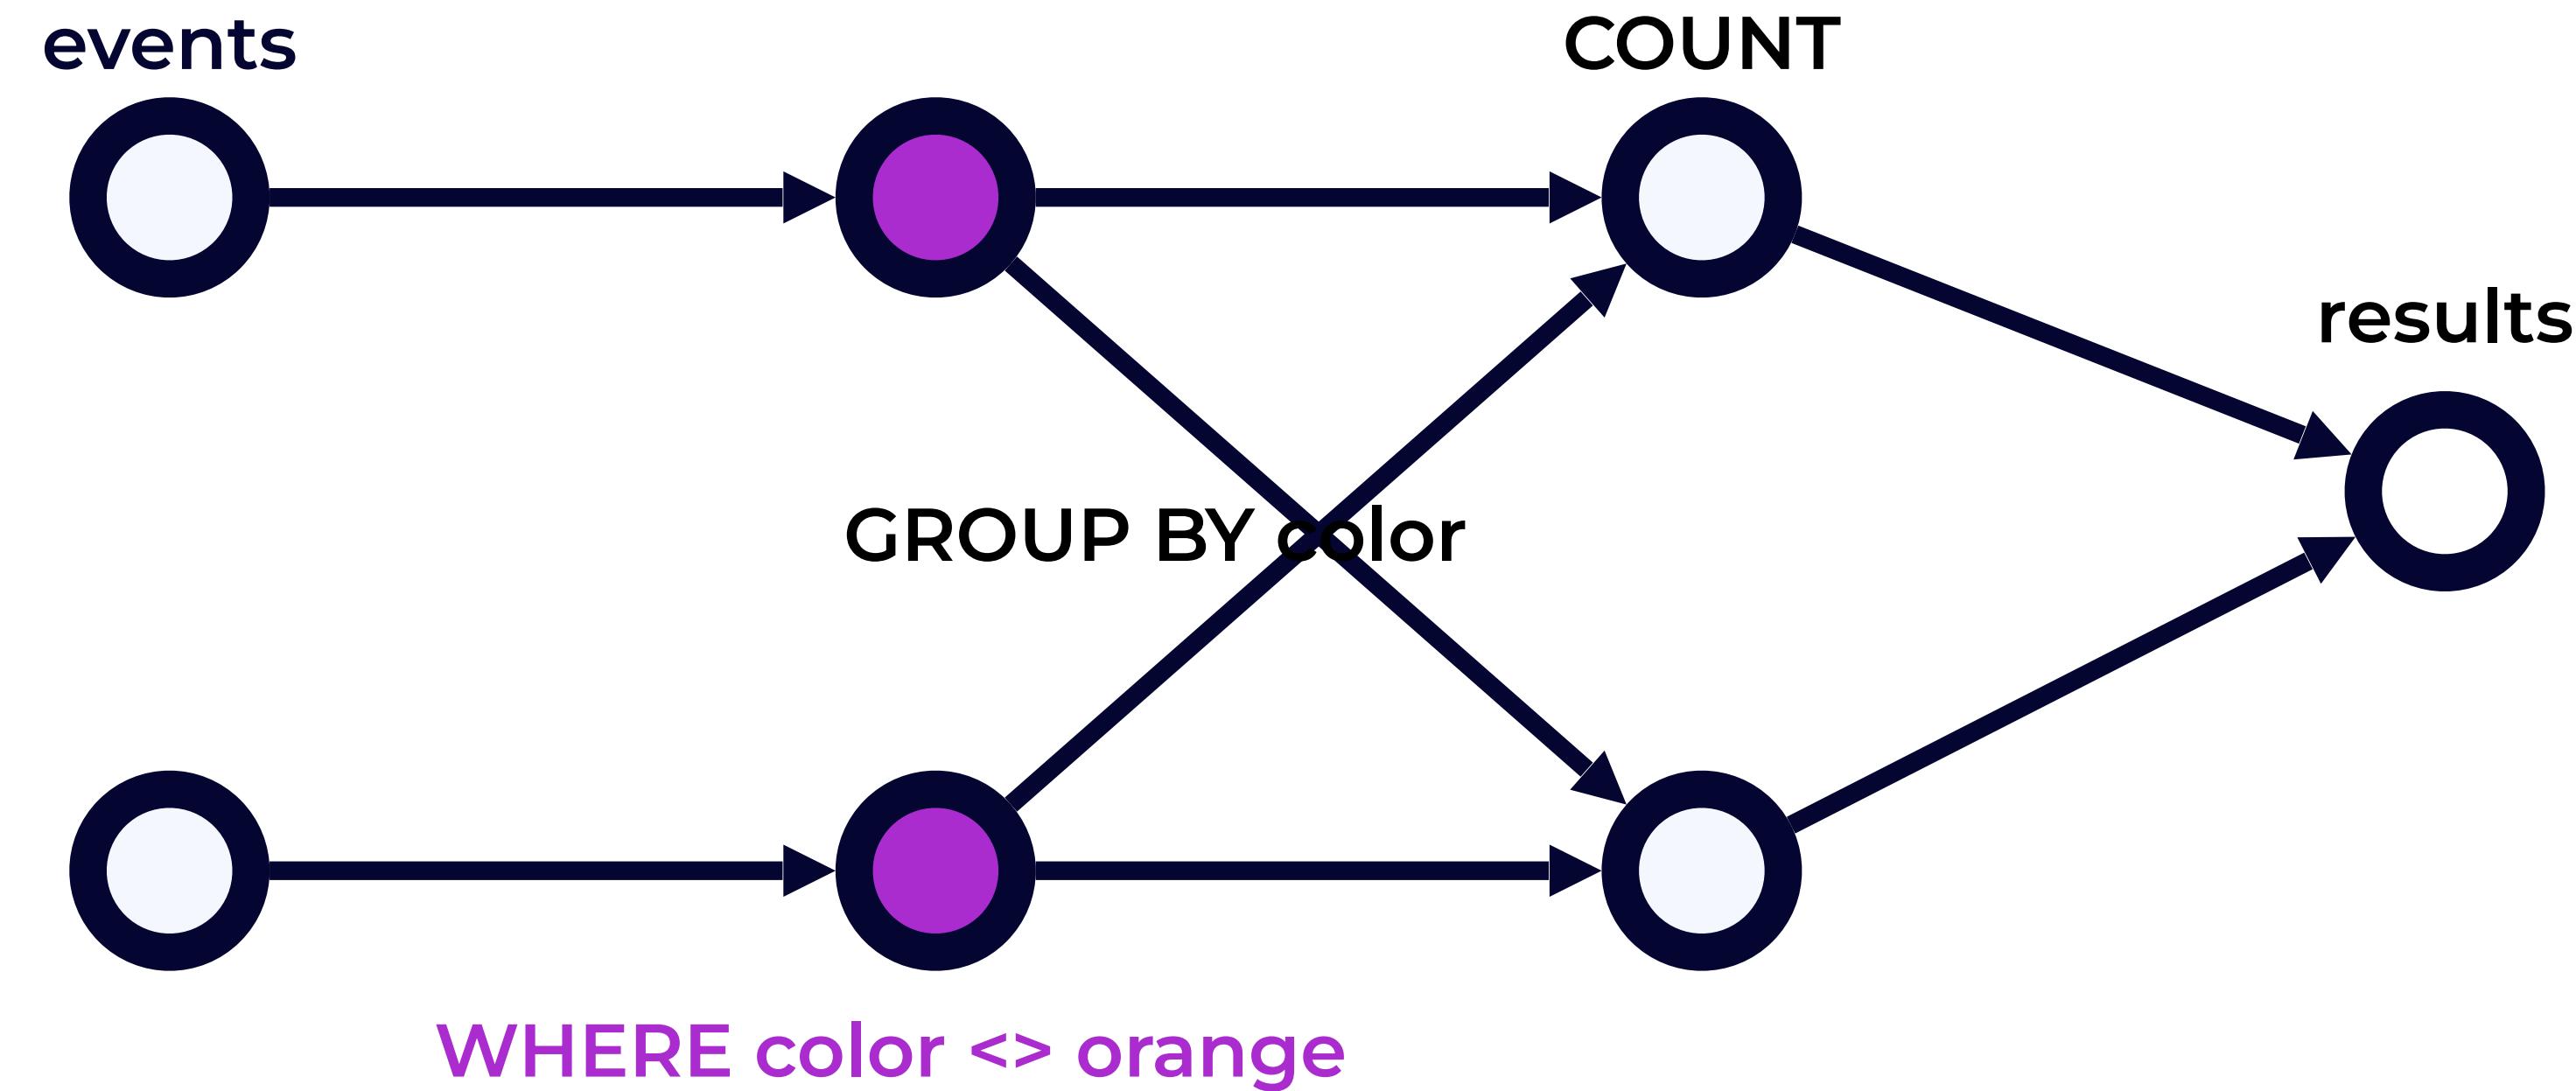


# Stateful stream processing with Flink SQL



```
INSERT INTO results
SELECT color, COUNT(*)
FROM events
WHERE color != orange
GROUP BY color;
```

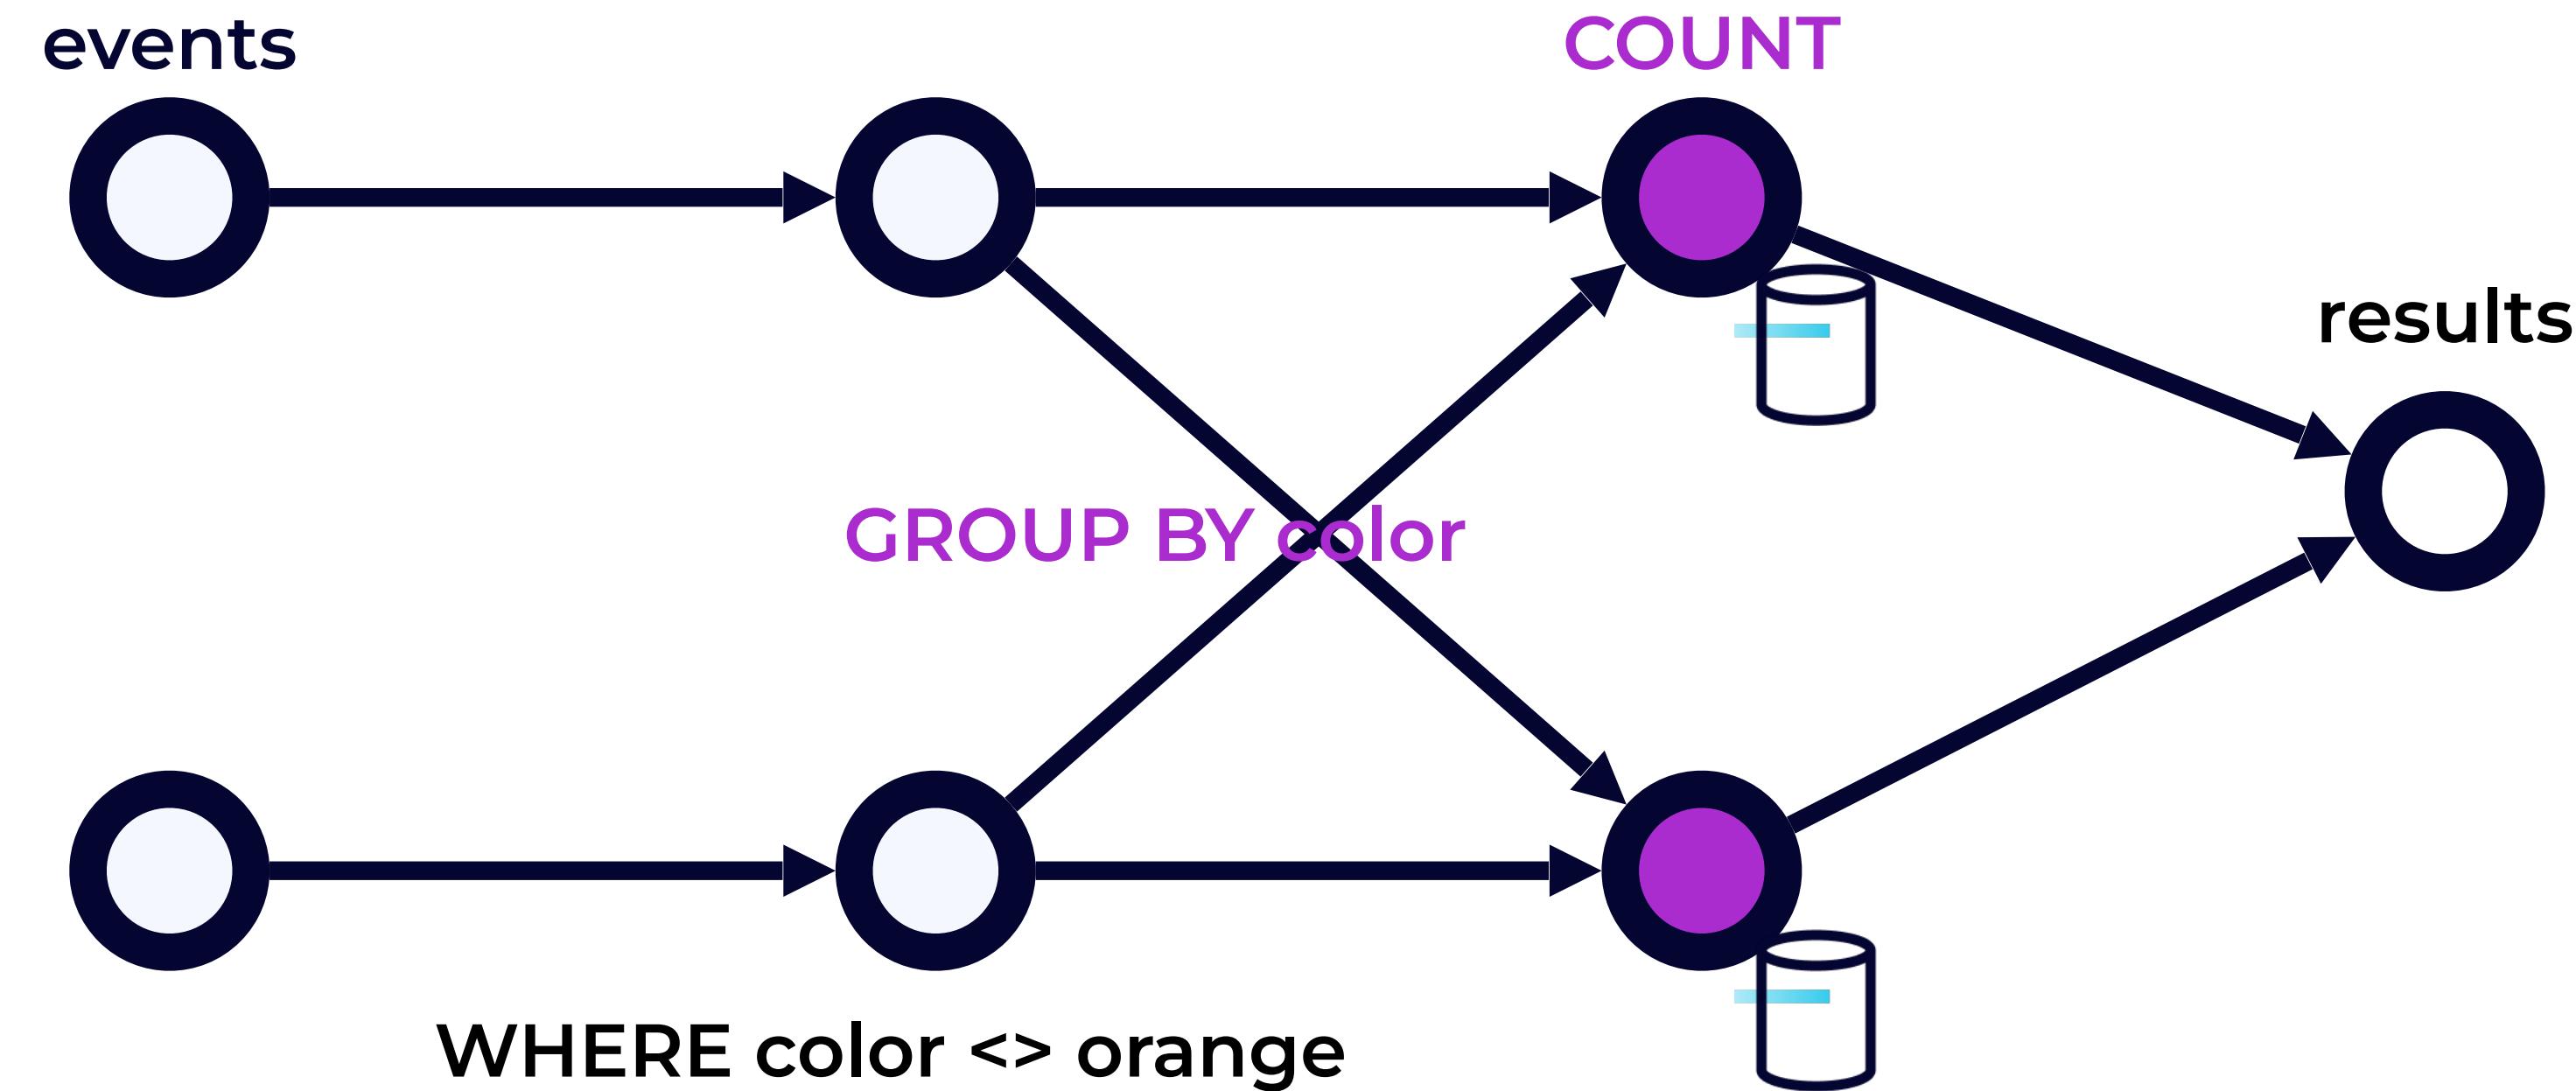
- Filtering is stateless



# Stateful stream processing with FlinkSQL



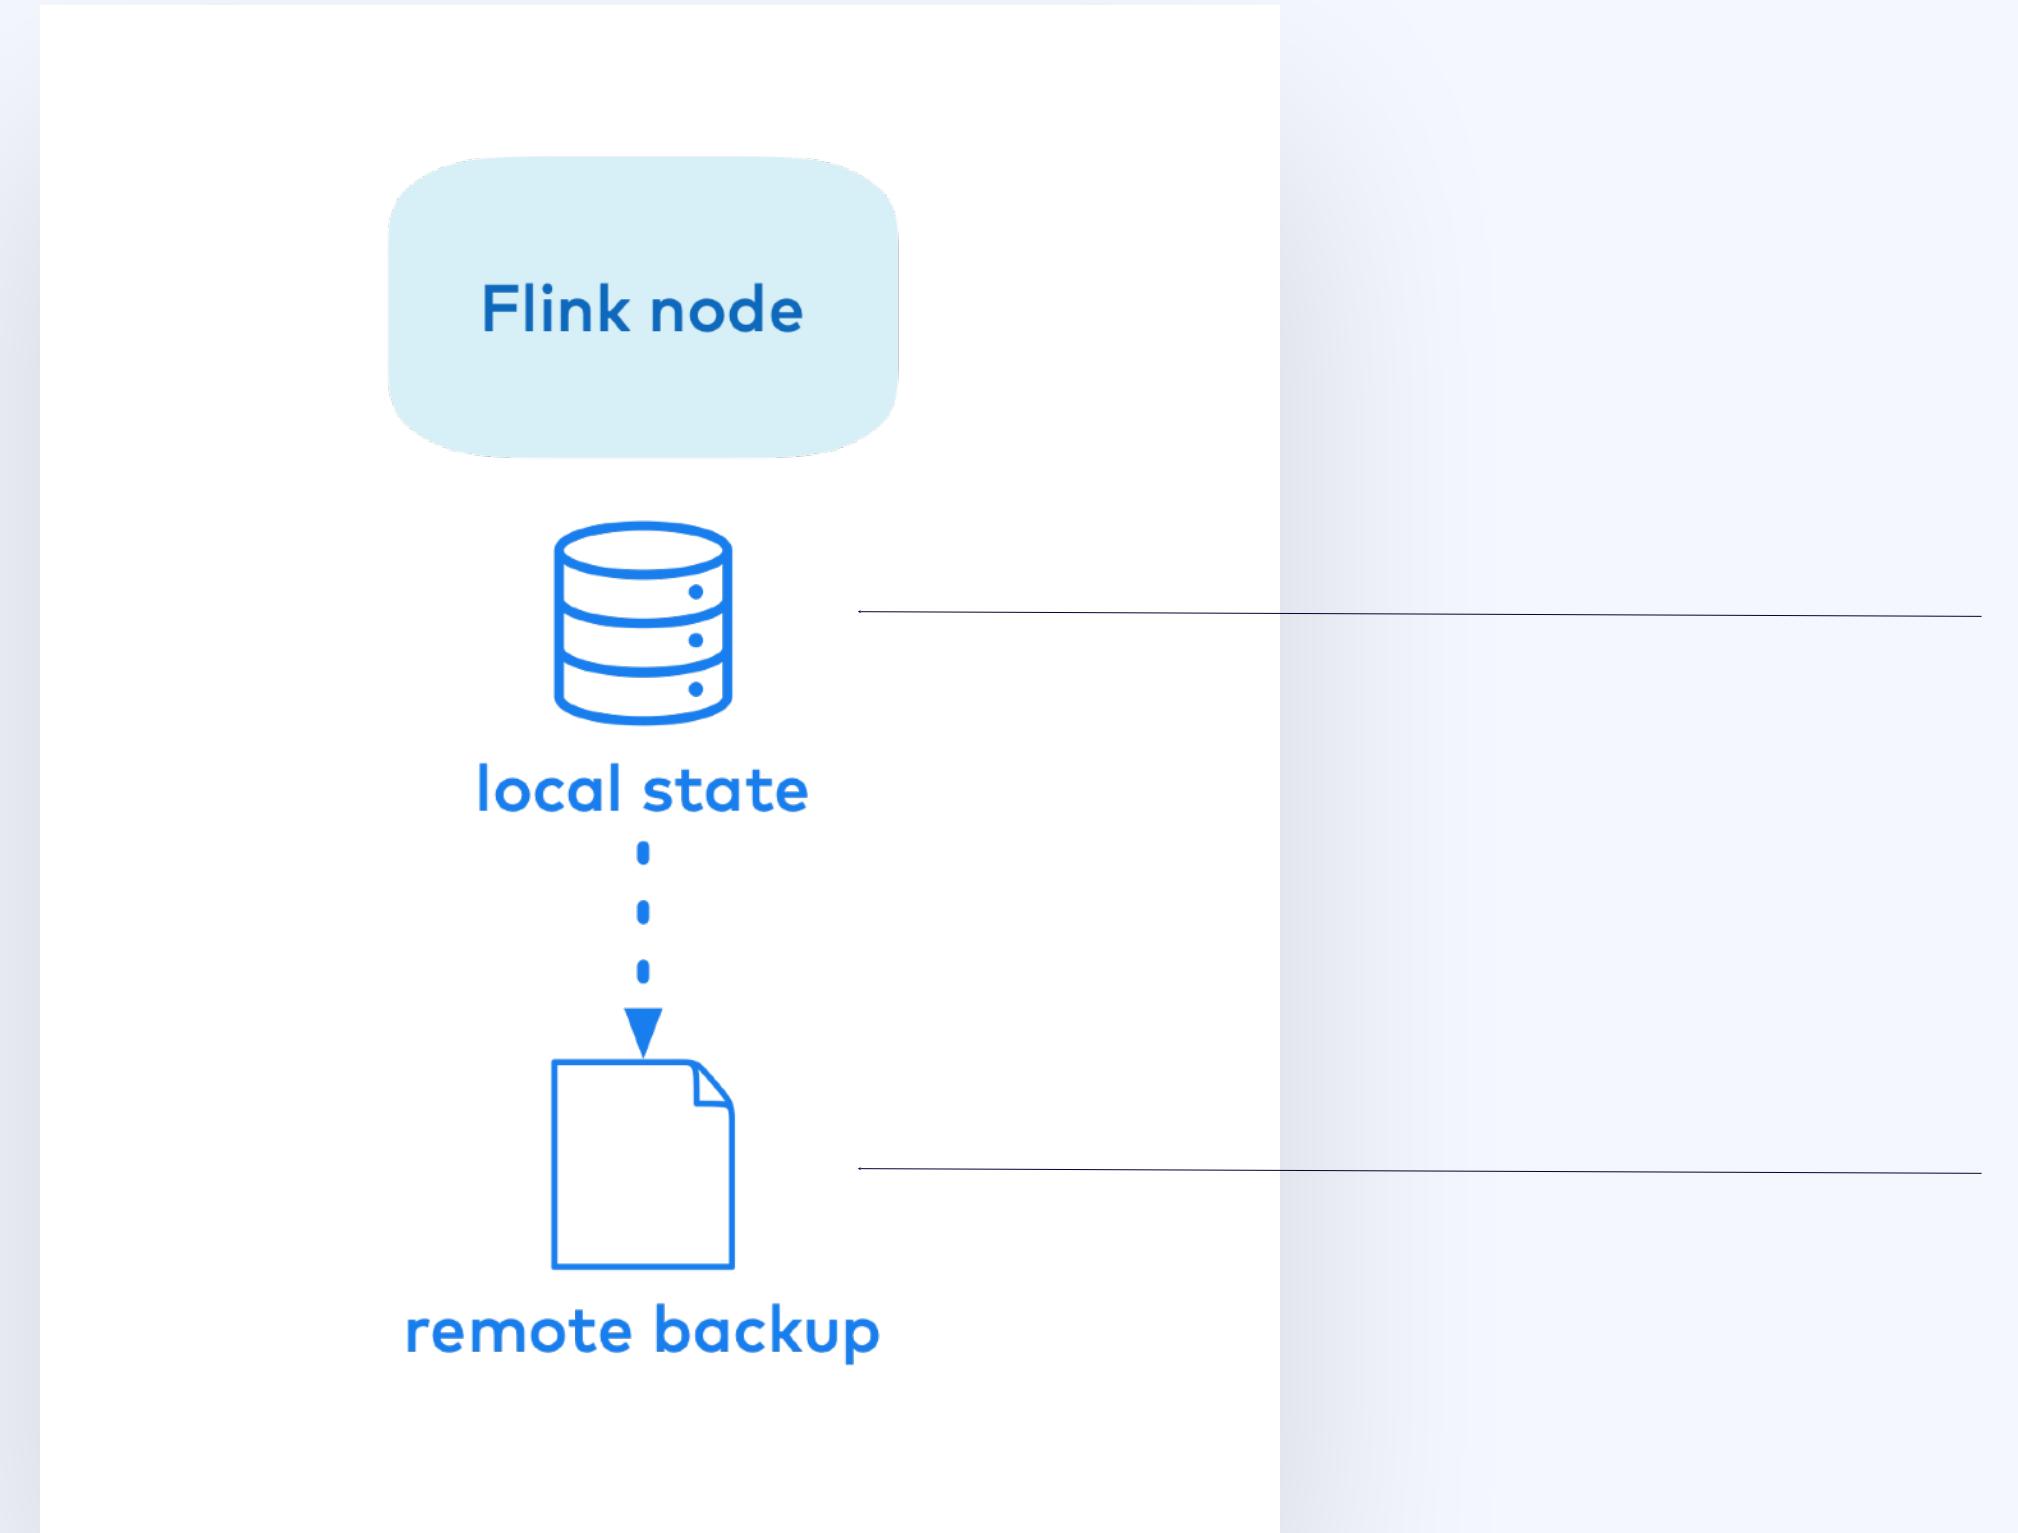
- Counting requires state



# State



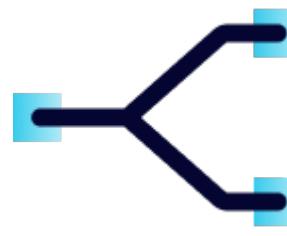
- **Local**
- **Fast**
- **Fault tolerant**



Stored on the heap or  
on disk using RocksDB  
(a KV store)

Distributed, reliable  
storage such as S3 or  
HDFS

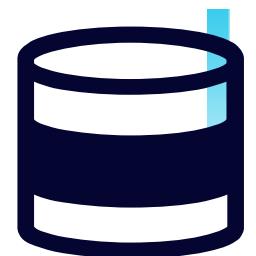
# Summary



## Streaming

A sequence of events.

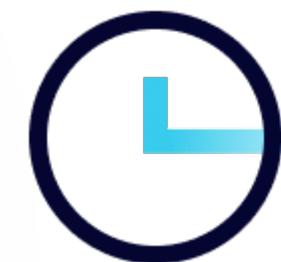
Unfamiliar to many developers, but ultimately straightforward.



## State

Delightfully simple

- local
- key/value
- single-threaded



## Event time and watermarks

Watermarks indicate how much progress the time in the stream has made.



## State snapshots for recovery

Transparent to application developers, enables correctness and operations.

*As Always  
Have a Nice Day...*