

Streaming SQL with

# PipelineDB

# What is PipelineDB?

# What is PipelineDB?

- Continuous SQL on streams (continuous views)

# What is PipelineDB?

- Continuous SQL on streams (continuous views)
- High-throughput, incremental materialized views

# What is PipelineDB?

- Continuous SQL on streams (continuous views)
- High-throughput, incremental materialized views
- Based on PostgreSQL 9.5



# What is PipelineDB?

- Continuous SQL on streams (continuous views)
- High-throughput, incremental materialized views
- Based on PostgreSQL 9.5
- No special client libraries



# What is PipelineDB?

- Continuous SQL on streams (continuous views)
- High-throughput, incremental materialized views
- Based on PostgreSQL 9.5
- No special client libraries
- Free and open-source (GPLv3)



# What is PipelineDB?

- Continuous SQL on streams (continuous views)
- High-throughput, incremental materialized views
- Based on PostgreSQL 9.5
- No special client libraries
- Free and open-source (GPLv3)
- *(30-second demo)*



When is PipelineDB **not useful**?

# When is PipelineDB **not useful**?

- SQL isn't a fit
- Ad-hoc on granular data

When is PipelineDB **useful**?

# When is PipelineDB **useful**?

- High throughput aggregations (**realtime reporting/analytics**)
- Computations over sliding windows (**continuous monitoring/ops**)
- Queries are known in advance

# 100,000 feet

Produce

```
2010-02-10 00:01:07 w3svc1446 WEB100 216.167.204.29 GET /tips-tricks/f
downloaded-error-in-ie/ - 80 - 12.178.189.252 HTTP/1.1 Mozilla/5.0(Ma
(KHTML,like+Gecko)+version/4.0.3+Safari/531.9 - http://www.google.com
3A+your+current+security+settings+do+not+allow+this+program+to+be+down
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 811 475 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-includes/f
Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+AppleWebKit/5
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1537 446 124
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 5941 478 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1365 474 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
- 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1414 495 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1443 483 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
```

Process



Consume





# 100,000 feet

Produce

```
2010-02-10 00:01:07 w3svc1446 WEB100 216.167.204.29 GET /tips-tricks/f
downloaded-error-in-ie/ - 80 - 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Ma
(KHTML+like+Gecko)+version/4.0.3+Safari/531.9 - http://www.google.com
3A+your+current+security+settings+do+not+allow+this+program+to+be+down
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 811 475 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-includes/f
Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+AppleWebKit/5
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1537 446 124
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 5941 478 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1365 474 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
- 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1414 495 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1443 483 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
```

Process

# SQL

Aggregation

Filtering

Sliding windows

Consume



# 100,000 feet

Produce

```
2010-02-10 00:01:07 w3svc1446 WEB100 216.167.204.29 GET /tips-tricks/f
downloaded-error-in-ie/ - 80 - 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Ma
(KHTML+like+Gecko)+version/4.0.3+safari/531.9 - http://www.google.com
3A+your+current+security+settings+do+not+allow+this+program+to+be+down
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 811 475 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-includes/f
Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+AppleWebKit/5
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1537 446 124
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 5941 478 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1365 474 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
- 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1414 495 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1443 483 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
```

Process

# SQL

Aggregation

Filtering

Sliding windows

# = Reduction

Consume



**Why did we build it?**

# Produce

```
2010-02-10 00:01:07 w3svc1446 WEB100 216.167.204.29 GET /tips-tricks/f
downloaded-error-in-ie/ - 80 - 12.178.189.252 HTTP/1.1 Mozilla/5.0(Ma
(KHTML,like+Gecko)+version/4.0.3+Safari/531.9 - http://www.google.com
3A+your+current+security+settings+do+not+allow+this+program+to+be+down
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 811 475 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-includes/f
Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+xx+10_6_2;+en-us)+AppleWebKit/5
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1537 446 124
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+xx+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 5941 478 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1365 474 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
- 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+xx+10
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1414 495 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+xx+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1443 483 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+xx+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
```

# Process



# Consume





Produce

```
2010-02-10 00:01:07 w3SVC1446 WEB100 216.167.204.29 GET /tips-tricks/f
downloaded-error-in-ie/ - 80 - 12.178.189.252 HTTP/1.1 Mozilla/5.0(Ma
(KHTML+like+Gecko)+version/4.0.3+Safari/531.9 - http://www.google.com
3A+your+current+security+settings+do+not+allow+this+program+to+be+down
2010-02-10 00:01:08 w3SVC1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 811 475 93
2010-02-10 00:01:08 w3SVC1446 WEB100 216.167.204.29 GET /wp-includes/i
Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+AppleWebKit/5
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1537 446 124
2010-02-10 00:01:08 w3SVC1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 5941 478 93
2010-02-10 00:01:08 w3SVC1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1365 474 93
2010-02-10 00:01:08 w3SVC1446 WEB100 216.167.204.29 GET /wp-content/pl
- 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1414 495 109
2010-02-10 00:01:08 w3SVC1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1443 483 109
2010-02-10 00:01:08 w3SVC1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
```

Consume



Produce

```
2010-02-10 00:01:07 w3svc1446 WEB100 216.167.204.29 GET /tips-tricks/f
downloaded-error-in-ie/ - 80 - 12.178.189.252 HTTP/1.1 Mozilla/5.0(Ma
(KHTML,like+Gecko)+version/4.0.3+Safari/531.9 - http://www.google.com
3A+your+current+security+settings+do+not+allow+this+program+to+be+down
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 811 475 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-includes/f
Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+AppleWebKit/5
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1537 446 124
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 5941 478 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1365 474 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
- 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1414 495 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1443 483 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+U;+Intel+Mac+OS+X+10_6
+Version/4.0.3+Safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
```

PipelineDB  
continuous view

Consume



# Simplicity is nice, but what else?

Produce

```
2010-02-10 00:01:07 w3svc1446 WEB100 216.167.204.29 GET /tips-tricks/f
downloaded-error-in-ie/ - 80 - 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Ma
(KHTML+like+Gecko)+version/4.0.3+safari/531.9 - http://www.google.com
3A+your+current+security+settings+do+not+allow+this+program+to+be+down
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 811 475 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-includes/i
Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+AppleWebKit/5
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1537 446 124
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 5941 478 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6_2;+en-us)+Appl
http://blog.caneja.com/tips-tricks/fix-your-current-security-settings-
200 0 0 1365 474 93
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
- 12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1414 495 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
error-in-ie/ blog.caneja.com 200 0 0 1443 483 109
2010-02-10 00:01:08 w3svc1446 WEB100 216.167.204.29 GET /wp-content/pl
12.178.189.252 HTTP/1.1 Mozilla/5.0+(Macintosh;+u;+Intel+Mac+OS+xx+10_6
+version/4.0.3+safari/531.9 - http://blog.caneja.com/tips-tricks/fix-y
```

PipelineDB

continuous view

Consume



# Benefits of continuous SQL on streams

- Aggregate before writing to disk

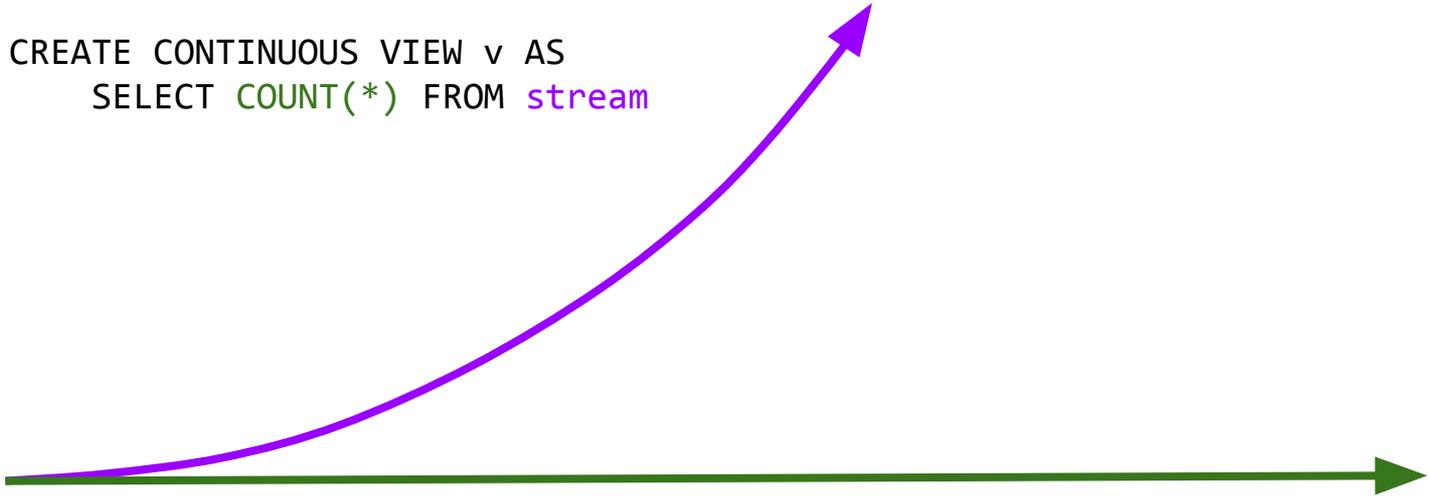
# Benefits of continuous SQL on streams

- Aggregate before writing to disk

```
CREATE CONTINUOUS VIEW v AS  
  SELECT COUNT(*) FROM stream
```

total data  
ingested

database size



# Benefits of continuous SQL on streams

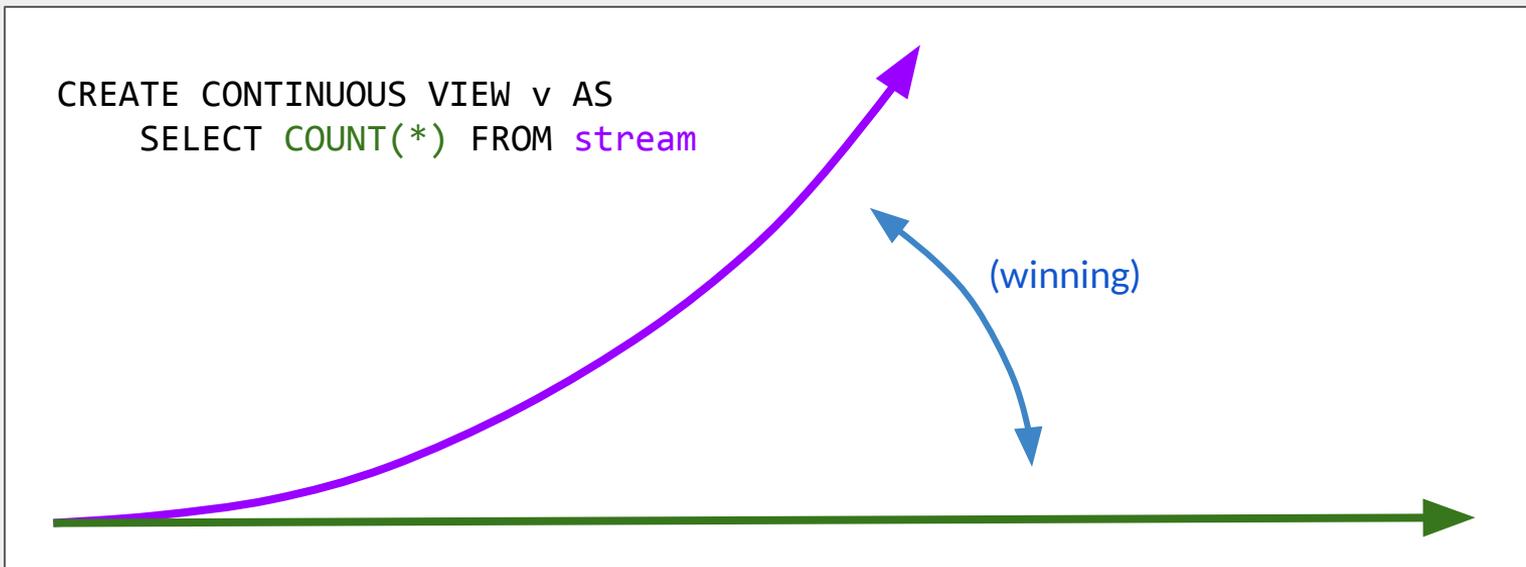
- Aggregate before writing to disk

total data  
ingested

```
CREATE CONTINUOUS VIEW v AS  
  SELECT COUNT(*) FROM stream
```

(winning)

database size



# Benefits of continuous SQL on streams

- Sliding window queries

```
CREATE CONTINUOUS VIEW v WITH (max_age = '1 hour') AS  
  SELECT COUNT(*) FROM stream
```

- Any information outside of the window is **excluded** from results and **deleted** from disk
- Essentially automatic TTL

# Benefits of continuous SQL on streams

- Probabilistic computations on infinite inputs

```
CREATE CONTINUOUS VIEW v AS  
  SELECT COUNT(DISTINCT x) FROM never_ending_stream
```

- Streaming **Top-K, Percentiles, distincts**, large set **cardinalities**
- Constant space
- No sorting
- Small margin of error

Internals (part 1/2)

---

# Streams and Workers



# Streams

- Internally, a stream is **Foreign Table**

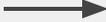
```
CREATE STREAM stream (x int, y int, z int);  
INSERT INTO stream (x, y, z) VALUES (0, 1, 2);
```

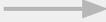
# Streams

- Internally, a stream is **Foreign Table**

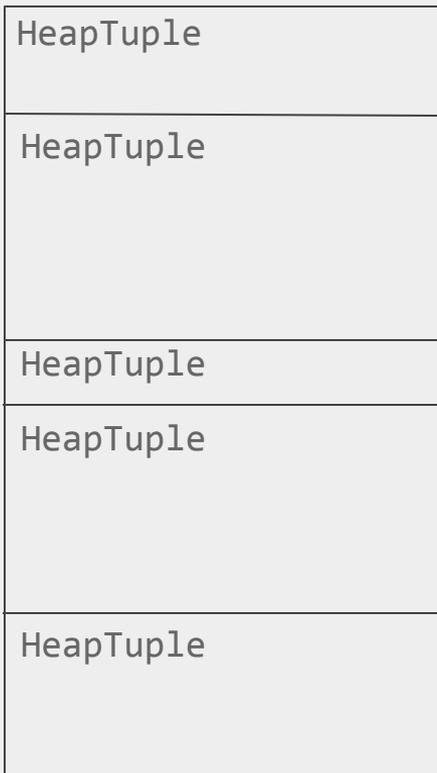
```
CREATE STREAM stream (x int, y int, z int);  
INSERT INTO stream (x, y, z) VALUES (0, 1, 2);
```

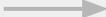
- System-wide **Foreign Server** called **pipeline\_streams**
- **stream\_fdw** reads from/writes to the **Stream Buffer**
- No permanent storage
- Stream rows only exist until they've been fully read

stream buffer  query on microbatch  incremental table update

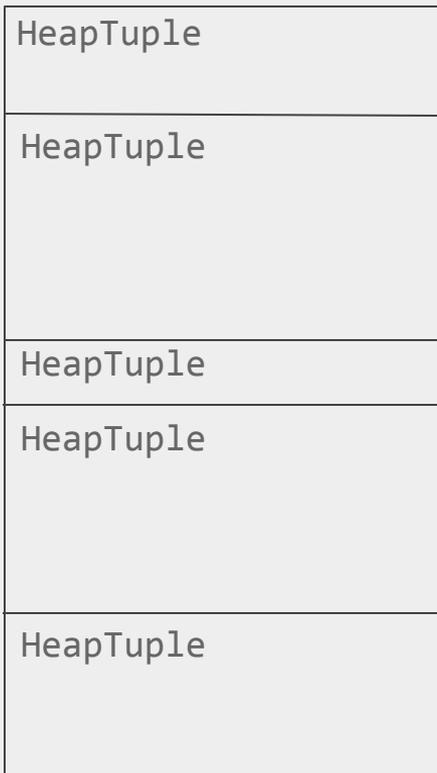
stream buffer  query on microbatch  incremental table update

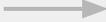
- INSERT INTO ...



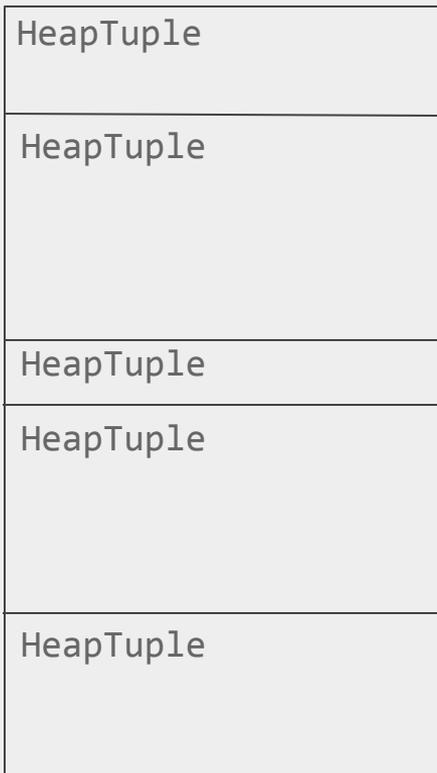
**stream buffer**  `query on microbatch`  `incremental table update`

- **INSERT INTO ...**
- **Concurrent circular buffer**



stream buffer  query on microbatch  incremental table update

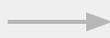
- **INSERT INTO ...**
- **Concurrent circular buffer**
- **Preallocated block of shared memory**



stream buffer

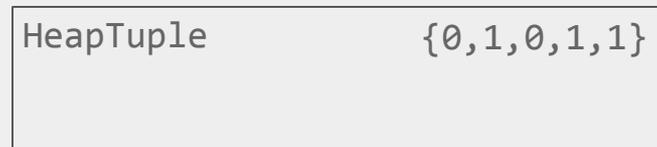
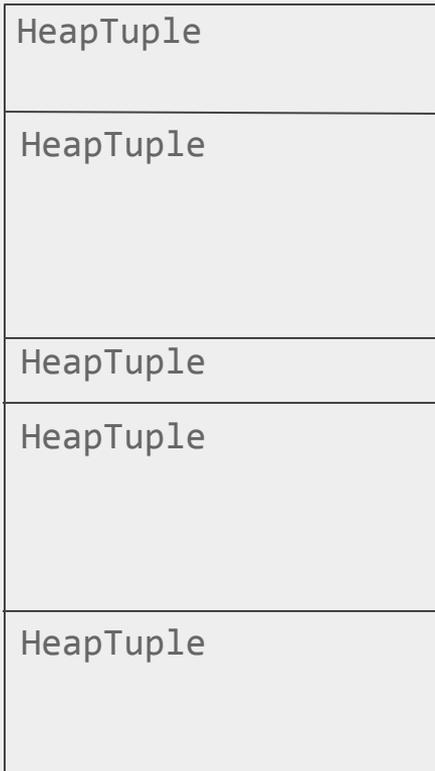


query on microbatch



incremental table update

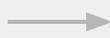
- **INSERT INTO ...**
- **Concurrent circular buffer**
- **Preallocated block of shared memory**



stream buffer

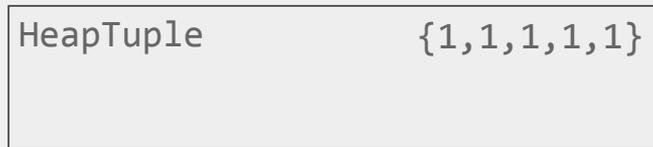
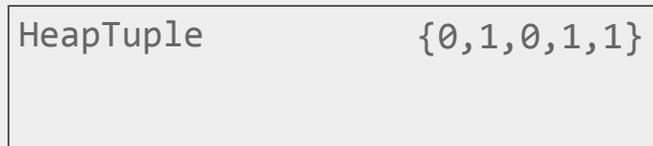
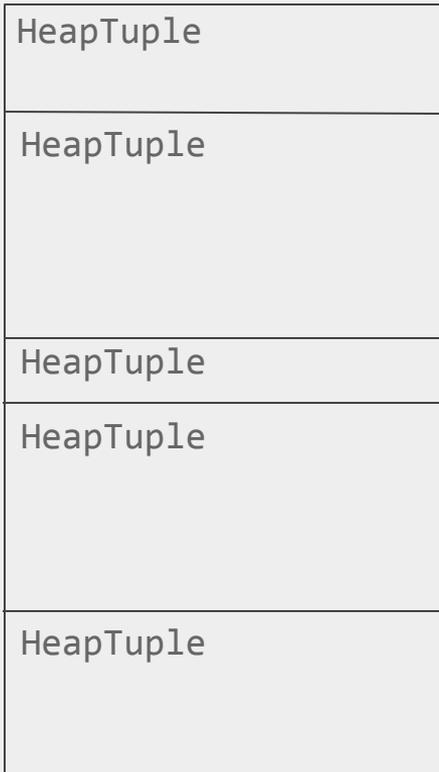


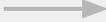
query on microbatch



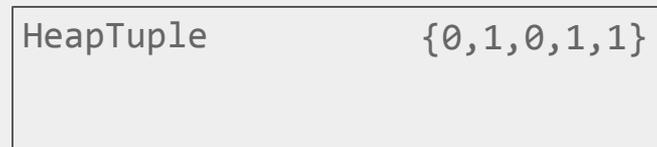
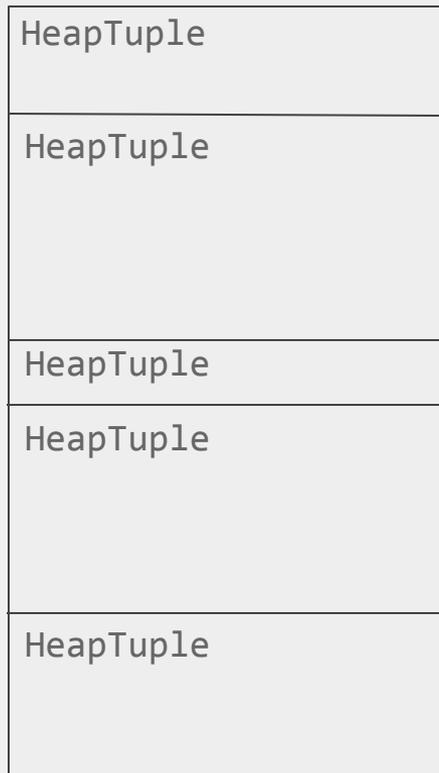
incremental table update

- INSERT INTO ...
- Concurrent circular buffer
- Preallocated block of shared memory



stream buffer  query on microbatch  incremental table update

- INSERT INTO ...
- Concurrent circular buffer
- Preallocated block of shared memory



stream buffer



**query on microbatch**



incremental table update

```
/* At Postmaster startup time ... */  
worker.bgw_main = any_function;  
worker.bgw_main_arg = (Datum) arg;  
  
RegisterDynamicBackgroundWorker(&worker, &handle);
```

stream buffer



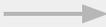
**query on microbatch**



incremental table update

```
SELECT count(*), avg(x) FROM stream
```

HeapTuple
HeapTuple
HeapTuple
...
HeapTuple

stream buffer  query on microbatch  incremental table update

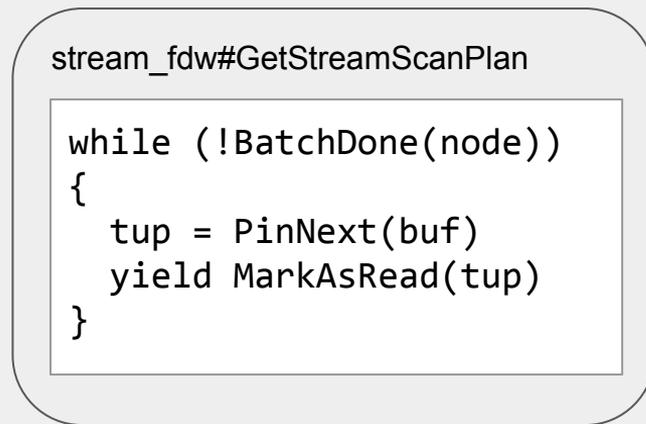
SELECT count(\*), avg(x) FROM stream

microbatch\_result

HeapTuple
HeapTuple
HeapTuple
...
HeapTuple



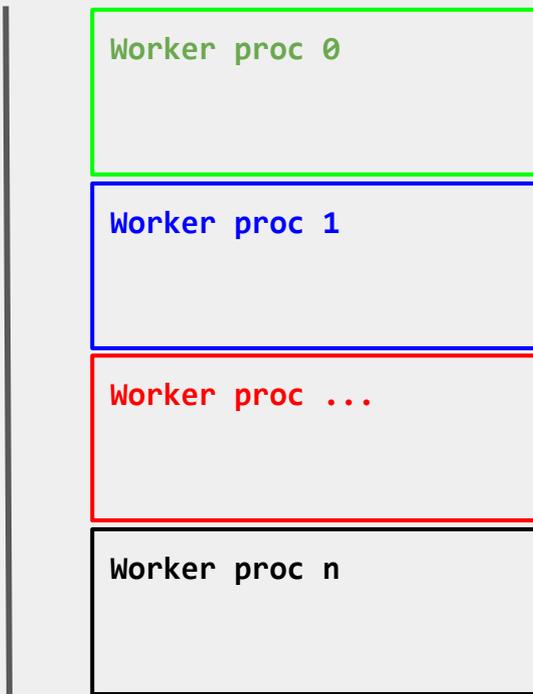
count	avg
<b>1000</b>	<b>{1000, 4000}</b>



# Worker process parallelism

tuples round-robin'd across  
n worker procs

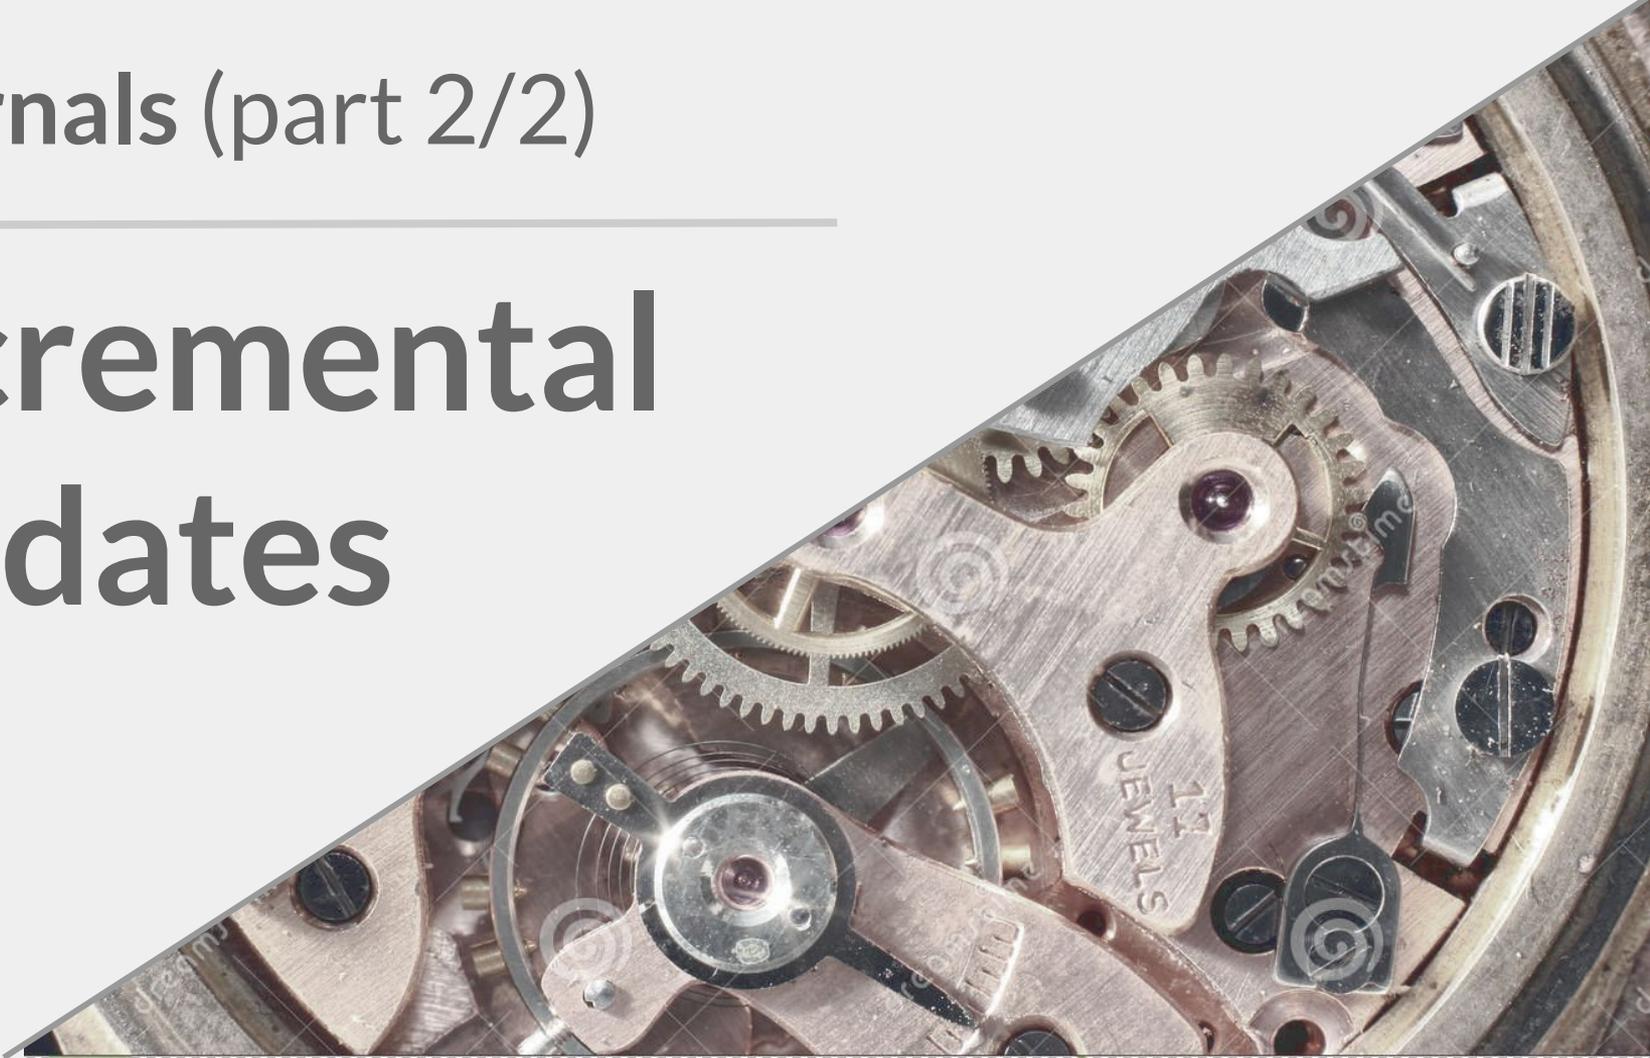
Stream buffer

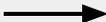


Internals (part 2/2)

---

# Incremental Updates



stream buffer  query on microbatch  **incremental table update**

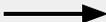
- `transition_state = combine(microbatch_tstate, existing_tstate)`

stream buffer  query on microbatch  incremental table update

- `transition_state = combine(microbatch_tstate, existing_tstate)`
- `pipeline_combine` catalog table maps combine functions to aggregates

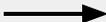
stream buffer  query on microbatch  incremental table update

- `transition_state = combine(microbatch_tstate, existing_tstate)`
- `pipeline_combine` catalog table maps combine functions to aggregates
- No changes to `pg_aggregate` catalog table or existing aggregate functions

stream buffer  query on microbatch  incremental table update

- `transition_state = combine(microbatch_tstate, existing_tstate)`
- `pipeline_combine` catalog table maps combine functions to aggregates
- No changes to `pg_aggregate` catalog table or existing aggregate functions
- User-defined aggregates just need a `combinefunc` to be combinable

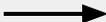
```
CREATE AGGREGATE combinable_agg(x)
(
    sfunc=sfunc,
    finalfunc=finalfunc,
    combinefunc=combinefunc,
);
```

stream buffer  query on microbatch  incremental table update

- `transition_state = combine(microbatch_tstate, existing_tstate)`
- `pipeline_combine` catalog table maps combine functions to aggregates

microbatch\_result

count	avg
<b>1000</b>	<b>{1000, 4000}</b>

stream buffer  query on microbatch  incremental table update

- `transition_state = combine(microbatch_tstate, existing_tstate)`
- `pipeline_combine` catalog table maps combine functions to aggregates

microbatch\_result

count	avg
<b>1000</b>	<b>{1000, 4000}</b>



combine()

count	avg
<b>1000</b>	<b>{1000, 4000}</b>
<b>5000</b>	<b>{5000, 10000}</b>

stream buffer  query on microbatch  incremental table update

- `transition_state = combine(microbatch_tstate, existing_tstate)`
- `pipeline_combine` catalog table maps combine functions to aggregates

microbatch\_result

count	avg
1000	{1000, 4000}



combine()

count	avg
1000	{1000, 4000}
5000	{5000, 10000}

existing on-disk row 

stream buffer  query on microbatch  incremental table update

- `transition_state = combine(microbatch_tstate, existing_tstate)`
- `pipeline_combine` catalog table maps combine functions to aggregates

microbatch\_result

count	avg
1000	{1000, 4000}



combine()

count	avg
1000	{1000, 4000}
5000	{5000, 10000}



updated\_result

count	avg
6000	{6000, 14000}

existing on-disk row



stream buffer



query on microbatch



**incremental table update**

```
lookup_plan = get_plan(SELECT * FROM matrel WHERE hash_group(x, y, z) IN (...))
```

stream buffer



query on microbatch



**incremental table update**

```
lookup_plan = get_plan(SELECT * FROM matrel WHERE hash_group(x, y, z) IN (...))
```

```
/* dynamically generate a VALUES node */
```

```
foreach(row, microbatch)
```

```
  values = lappend(values, hash_group(row));
```

stream buffer



query on microbatch



**incremental table update**

```
lookup_plan = get_plan(SELECT * FROM matrel WHERE hash_group(x, y, z) IN (...))
```

```
/* dynamically generate a VALUES node */
```

```
foreach(row, microbatch)
```

```
  values = lappend(values, hash_group(row));
```

```
set_values(lookup_plan, values)
```

stream buffer



query on microbatch



**incremental table update**

```
lookup_plan = get_plan(SELECT * FROM matrel WHERE hash_group(x, y, z) IN (...))
```

```
/* dynamically generate a VALUES node */
```

```
foreach(row, microbatch)
```

```
  values = lappend(values, hash_group(row));
```

```
set_values(lookup_plan, values)
```

```
existing = PortalRun(lookup_plan, ...)
```

```
/* now we're reading to combine these on-disk tuples with the incoming batch result */
```

stream buffer



query on microbatch



**incremental table update**

```
SELECT * FROM matrel WHERE hash_group(x, y, z) IN (hash(microbatch group), ...)
```

- This query needs to be as fast as possible
- Continuous views indexed on a 32-bit hash of grouping
- **Pro:** maximize cardinality of the index keyspace, great for random perf
- **Con:** must deal with collisions programmatically

stream buffer



query on microbatch



incremental table update

```
SELECT * FROM matrel WHERE hash_group(x, y, z) IN (hash(microbatch group), ...)
```

- If the grouping contains a **time-based column**, we can do better

```
CREATE ... AS SELECT day(timestamp), count(*) FROM stream GROUP BY day
```

stream buffer



query on microbatch



incremental table update

```
SELECT * FROM matrel WHERE hash_group(x, y, z) IN (hash(microbatch group), ...)
```

- If the grouping contains a **time-based column**, we can do better

```
CREATE ... AS SELECT day(timestamp), count(*) FROM stream GROUP BY day
```

- These continuous views are indexed with **64 bits**:hash of grouping

Timestamp from group (32 bits)

Regular 32-bit grouping hash

stream buffer



query on microbatch



incremental table update

```
SELECT * FROM matrel WHERE hash_group(x, y, z) IN (hash(microbatch group), ...)
```

- If the grouping contains a **time-based column**, we can do better

```
CREATE ... AS SELECT day(timestamp), count(*) FROM stream GROUP BY day
```

- These continuous views are indexed with **64 bits**:hash of grouping

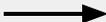
Timestamp from group (32 bits)

Regular 32-bit grouping hash

- **Pro**: most incoming groups will have a similar timestamp, so better index caching
- **Con**: larger index footprint

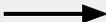
stream buffer  query on microbatch  **incremental table update**

 **microbatch result** generated from stream by worker

stream buffer  query on microbatch  **incremental table update**

✓ **microbatch result** generated from stream by worker

✓ **existing result** retrieved from disk

stream buffer  query on microbatch  incremental table update

- ✓ microbatch result generated from stream by worker
- ✓ existing result retrieved from disk

```
combine_plan = get_plan(SELECT group, combine(count), combine(avg)
  FROM microbatch_result UNION existing GROUP BY group);
```

```
combined = PortalRun(combine_plan, ...)
```

```
foreach(row, combined)
{
  if (new_tuple(row))
    heap_insert(row, ...);
  else
    heap_update(row, ...);
}
```

# Combiner process parallelism

On-disk groupings are **sharded over combiners by group**

Each row is guaranteed to only ever be updated by **one combiner process**

Continuous view

grouping (a, b, c)

grouping (d, e, f)

grouping (g, h, i)

grouping (j, k, l)

Just released! **Continuous transforms**

# Just released! **Continuous transforms**

- Worker-only continuous queries

# Just released! **Continuous transforms**

- Worker-only continuous queries
- Arbitrary procedure called on its output rows

# Just released! **Continuous transforms**

- Worker-only continuous queries
- Arbitrary procedure called on its output rows
- Enable work sharing between continuous views

# Just released! Continuous transforms

- Worker-only continuous queries
- Arbitrary procedure called on its output rows
- Enable work sharing between continuous views

```
CREATE CONTINUOUS TRANSFORM xform AS
  SELECT foo(col), bar(col) FROM raw_stream
  THEN EXECUTE PROCEDURE pipeline_stream_insert('normalized_stream')
```

# Just released! Continuous transforms

- Worker-only continuous queries
- Arbitrary procedure called on its output rows
- Enable work sharing between continuous views

```
CREATE CONTINUOUS TRANSFORM xform AS
  SELECT foo(col), bar(col) FROM raw_stream
  THEN EXECUTE PROCEDURE pipeline_stream_insert('normalized_stream')
```

```
CREATE CONTINUOUS VIEW v0 AS SELECT ... FROM normalized_stream;
CREATE CONTINUOUS VIEW v1 AS SELECT ... FROM normalized_stream;
```

# Thanks !

- [derek@pipelinedb.com](mailto:derek@pipelinedb.com)
- [pipelinedb.com](http://pipelinedb.com)
- [docs.pipelinedb.com](http://docs.pipelinedb.com)
- [github.com/pipelinedb](https://github.com/pipelinedb)

