

Google Docs version of slides (including animations): <https://goo.gl/yzvLXe>

Fundamentals of Stream Processing with Apache Beam (incubating)



Frances Perry & Tyler Akidau
@francesjperry, @takidau
Apache Beam Committers & Google Engineers

Agenda

1

Infinite, Out-of-Order Data Sets

2

What, Where, When, How

3

Reasons This is Awesome

4

Apache Beam (incubating)

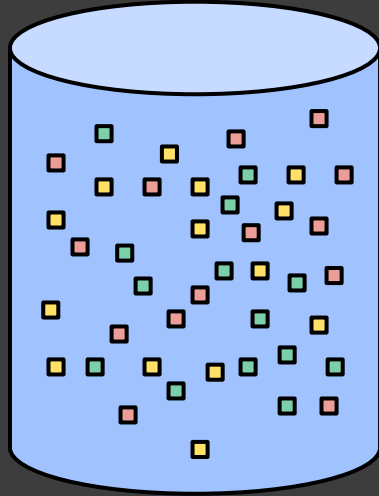





Infinite, Out-of-Order Data Sets



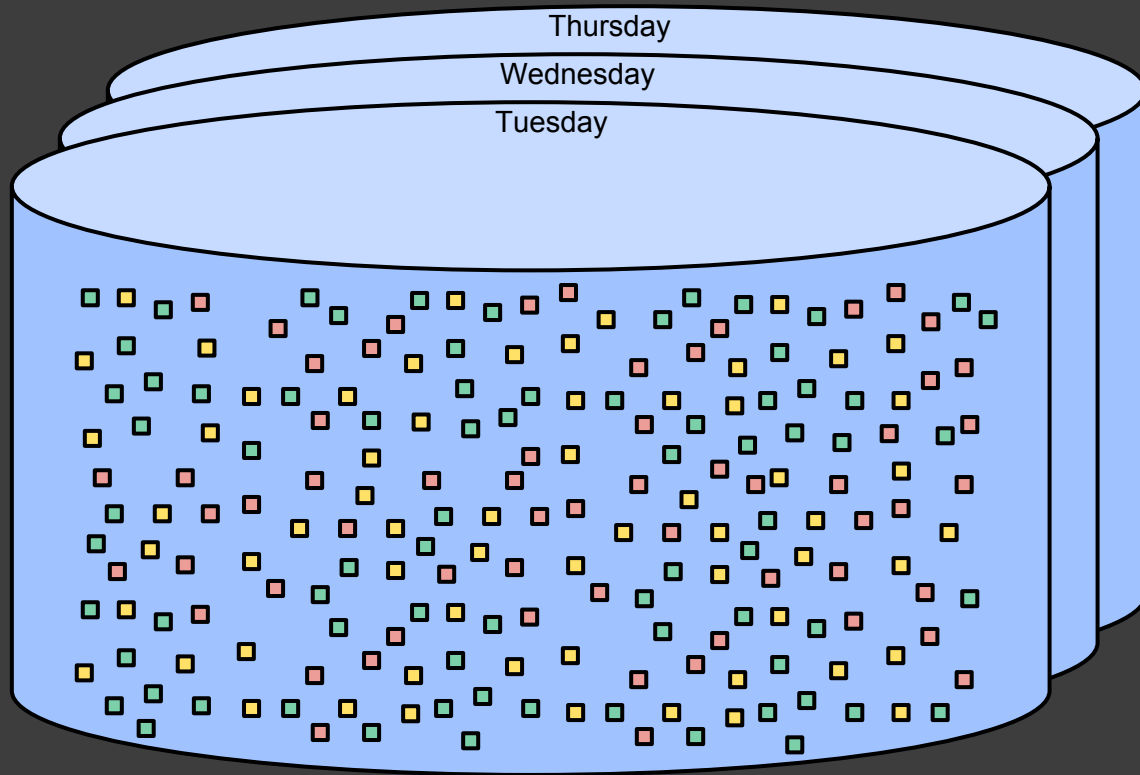
Data...




...can be big...



...really, really big...



... maybe infinitely big...



8:00

9:00


10:00

11:00


12:00

13:00


14:00




... with unknown delays.




Element-wise transformations




Aggregating via Processing-Time Windows




Aggregating via Event-Time Windows



Formalizing Event-Time Skew



Formalizing Event-Time Skew



Watermarks describe event time progress.

"No timestamp earlier than the watermark will be seen"

Often heuristic-based.

Too Slow? Results are *delayed*.

Too Fast? Some data is *late*.



What, Where, When, How



What are you computing?


Where in event time?

When in processing time?


How do refinements relate?




What are you computing?



Element-Wise



Aggregating



Composite

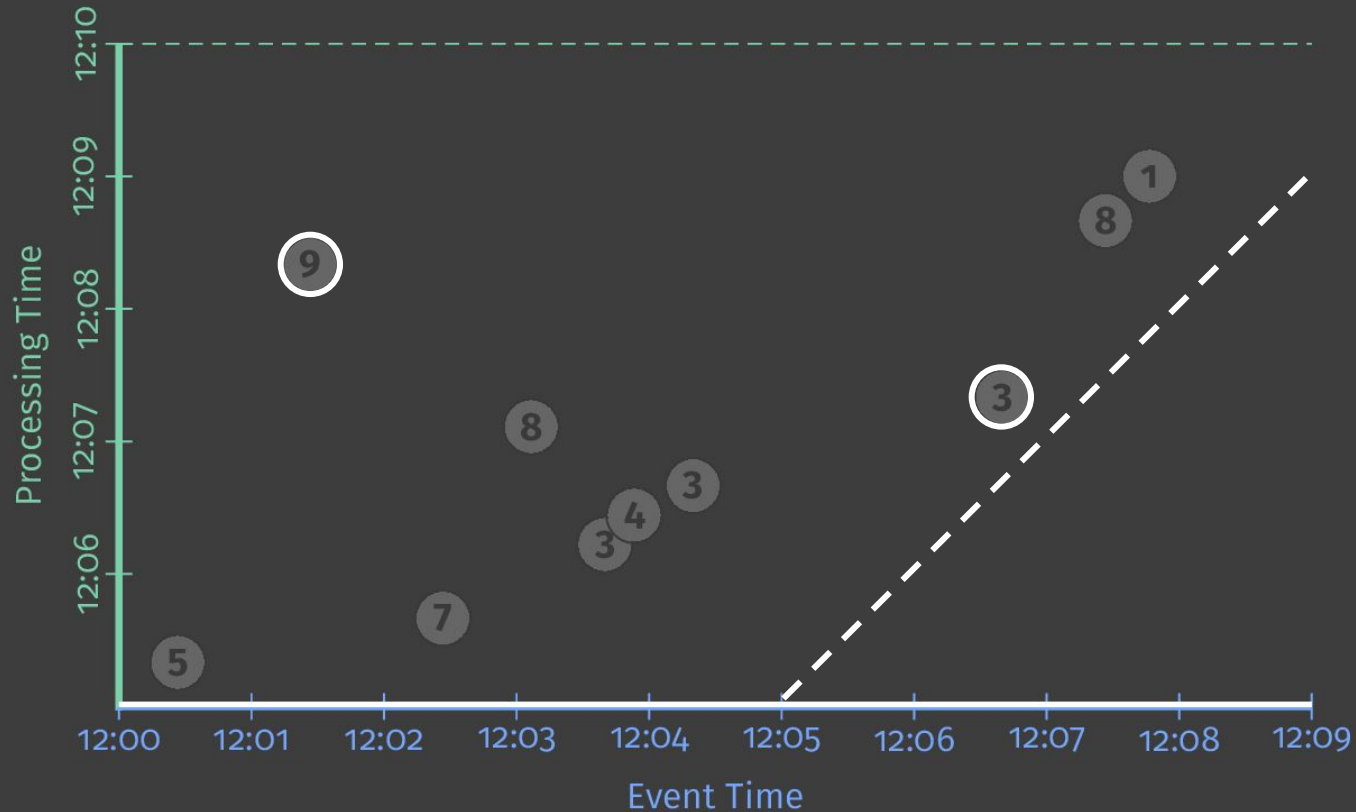
What: Computing Integer Sums

```
// Collection of raw log lines
PCollection<String> raw = IO.read(...);

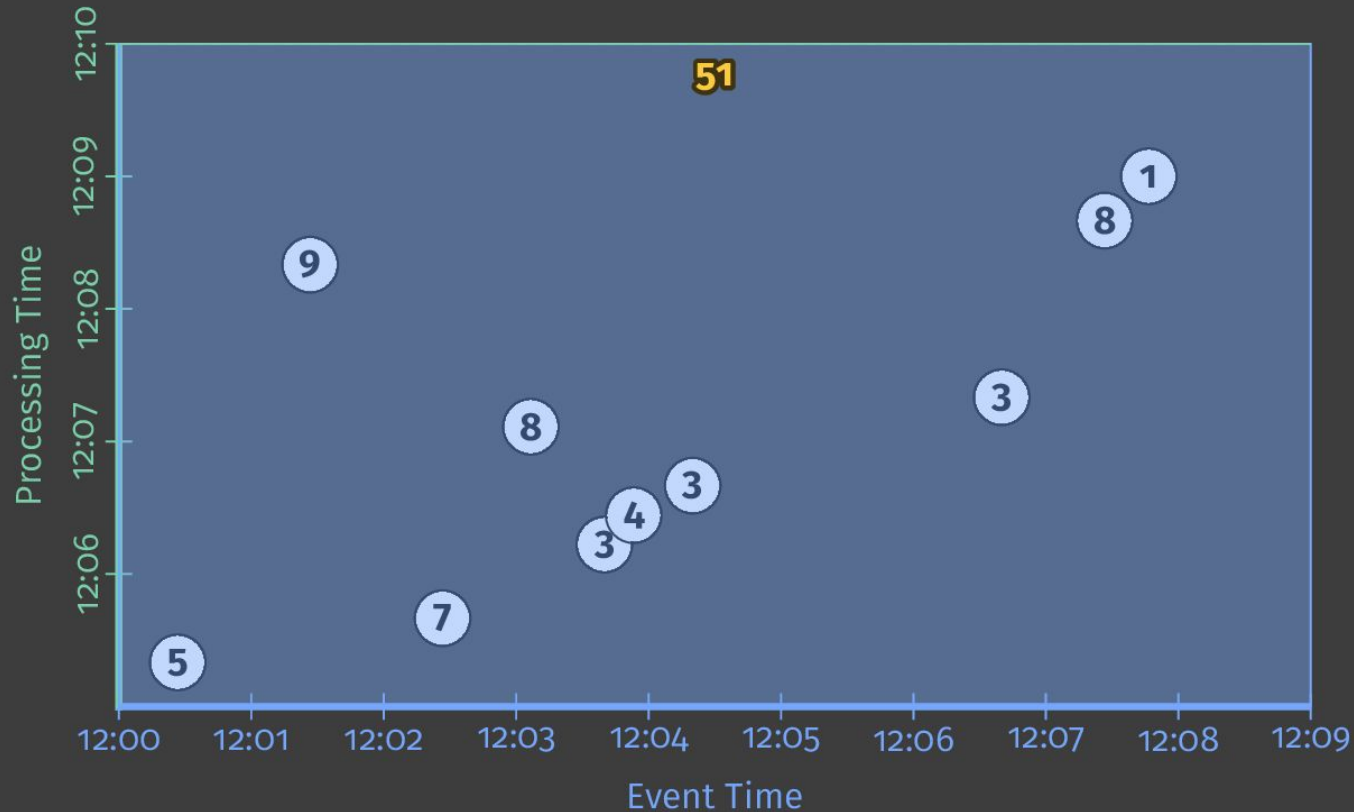
// Element-wise transformation into team/score pairs
PCollection<KV<String, Integer>> input =
    raw.apply(ParDo.of(new ParseFn()));

// Composite transformation containing an aggregation
PCollection<KV<String, Integer>> scores =
    input.apply(Sum.integersPerKey());
```

What: Computing Integer Sums




What: Computing Integer Sums



Where in event time?

Windowing divides data into event-time-based finite chunks.




Often required when doing aggregations over unbounded data.


Where: Fixed 2-minute Windows

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2))))
    .apply(Sum.integersPerKey());
```

Where: Fixed 2-minute Windows



When in processing time?




- Triggers control when results are emitted.
- Triggers are often relative to the watermark.



When: Triggering at the Watermark

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2))
        .triggering(AtWatermark())))
    .apply(Sum.integersPerKey());
```


When: Triggering at the Watermark

Perfect Watermark

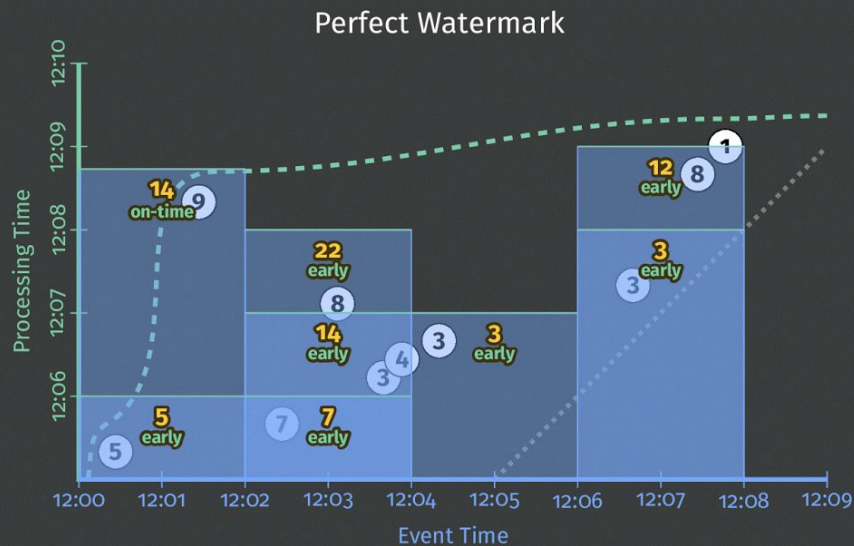


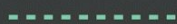

Perfect watermark: 
Ideal watermark: 

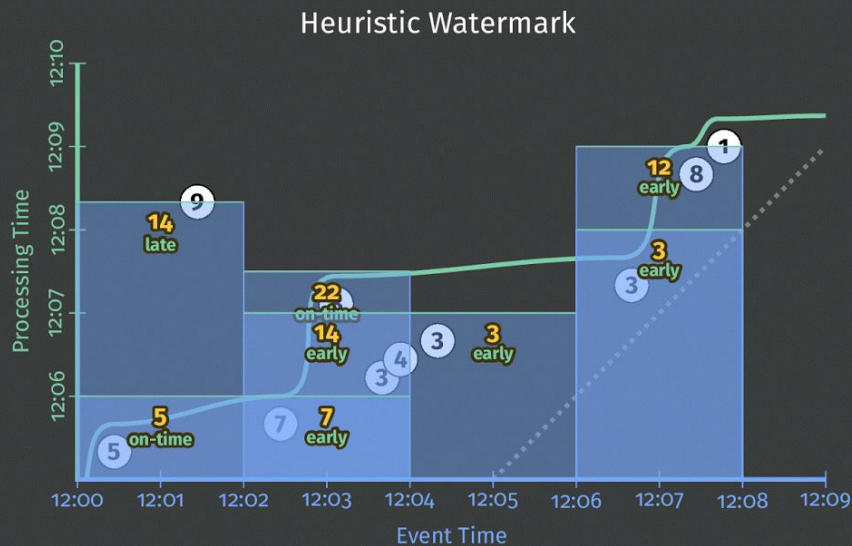
When: Early and Late Firings

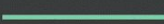

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2))
        .triggering(AtWatermark()
            .withEarlyFirings(AtPeriod(Minutes(1)))
            .withLateFirings(AtCount(1))))))
    .apply(Sum.integersPerKey());
```

When: Early and Late Firings



Perfect watermark: 
Ideal watermark: 



Heuristic watermark: 
Ideal watermark: 

How do refinements relate?

- How should multiple outputs per window accumulate?
- Appropriate choice depends on consumer.


Firing	Elements	Discarding	Accumulating	Acc. & Retracting
Speculative	[3]	3	3	3
Watermark	[5, 1]	6	9	9, -3
Late	[2]	2	11	11, -9
<i>Last Observed</i>		2	11	11
<i>Total Observed</i>		11	23	11

(Accumulating & Retracting not yet implemented.)

How: Add Newest, Remove Previous

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2))
        .triggering(AtWatermark()
            .withEarlyFirings(AtPeriod(Minutes(1)))
            .withLateFirings(AtCount(1)))
        .accumulatingAndRetractingFiredPanels()))
    .apply(Sum.integersPerKey());
```

How: Add Newest, Remove Previous



Heuristic watermark: —————

Ideal watermark:
—————



Reasons This is Awesome



What / Where / When / How

Correctness

Power

Composability

Flexibility

Modularity



What / **Where** / **When** / **How**

Correctness

Power


Composability

Flexibility

Modularity





Distributed Systems are Distributed




Ideal watermark:




Processing Time Results Differ



Event Time Results are Stable



Heuristic watermark: —————
Ideal watermark: ··········



Heuristic watermark: —————
Ideal watermark: ··········

What / Where / When / How

Correctness

Power

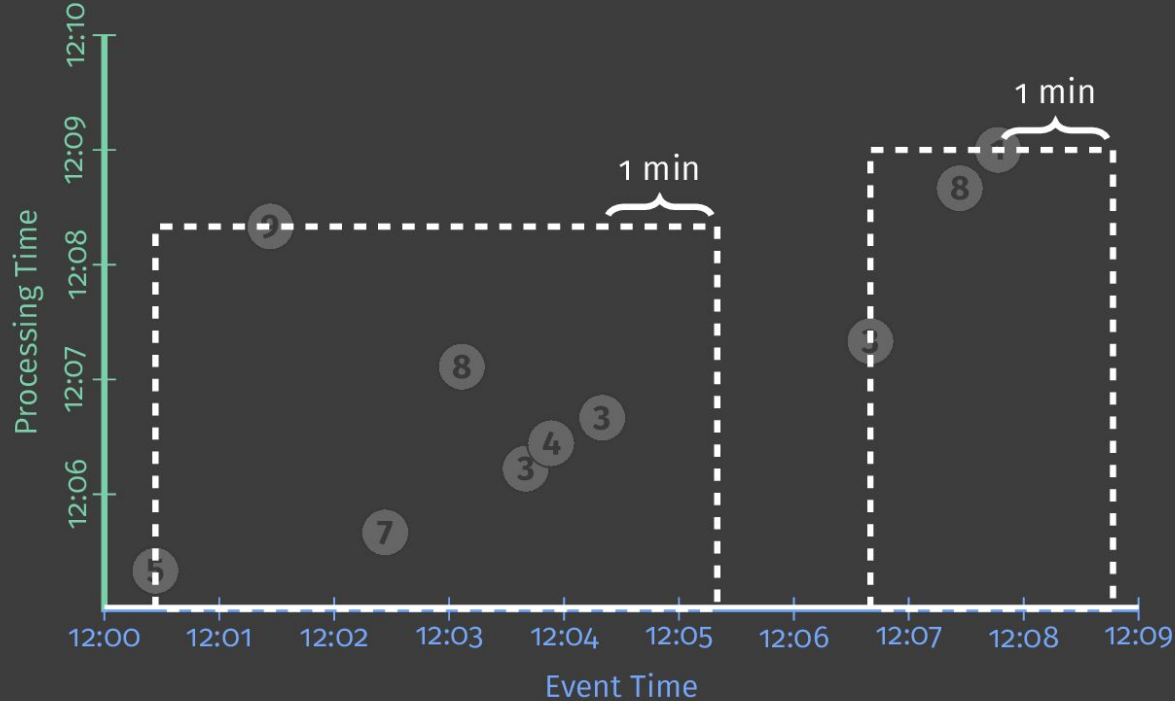
Composability

Flexibility

Modularity



Identifying Bursts of User Activity



Heuristic watermark: 

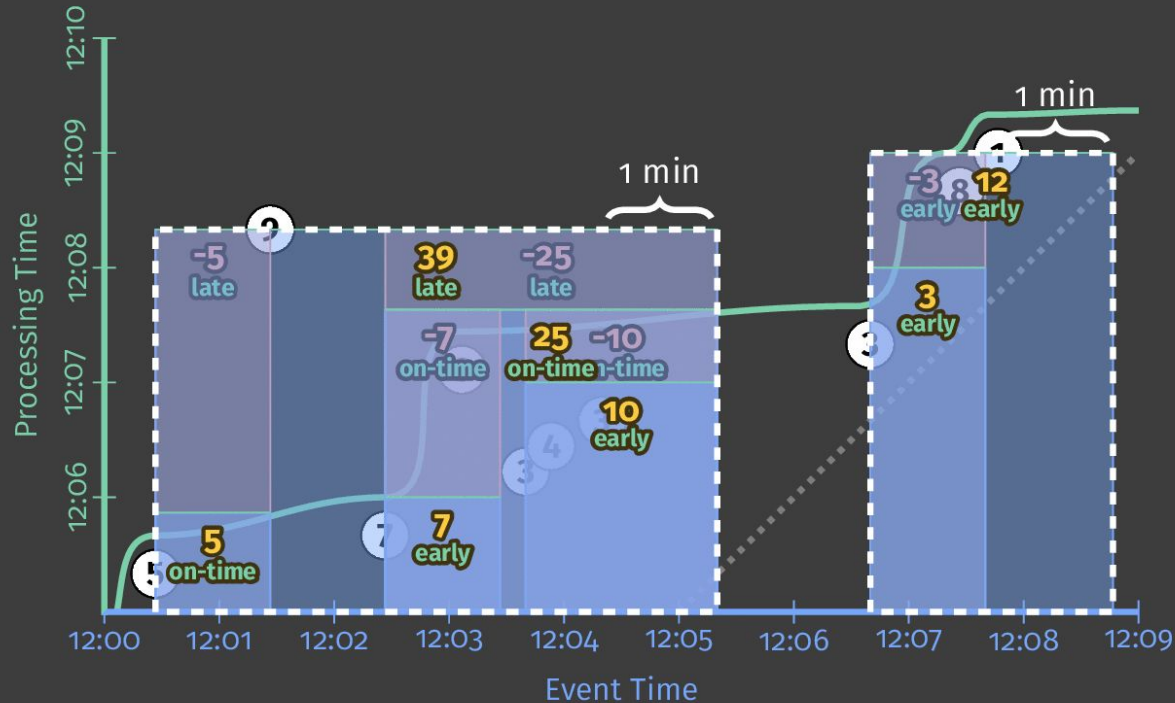
Ideal watermark: 



Sessions

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(Sessions.withGapDuration(Minutes(1)))
        .triggering(AtWatermark()
            .withEarlyFirings(AtPeriod(Minutes(1)))
            .withLateFirings(AtCount(1)))
        .accumulatingAndRetractingFiredPanels())
    .apply(Sum.integersPerKey());
```

Identifying Bursts of User Activity



Heuristic watermark: —————

Ideal watermark:
—————

What / Where / When / How

Correctness

Power


Composability

Flexibility


Modularity

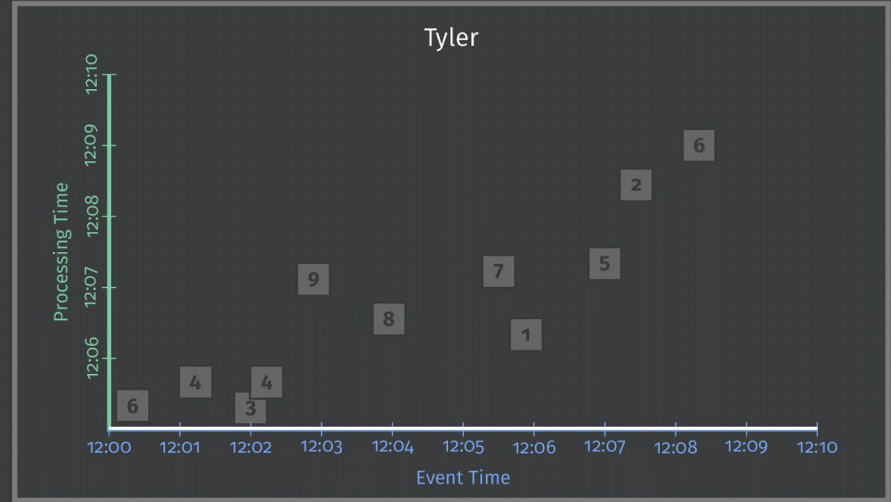
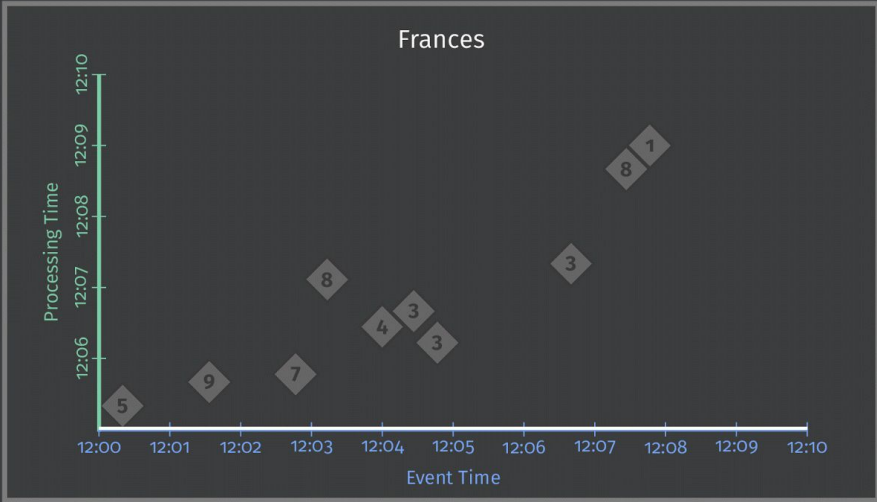


Calculating Session Lengths



```
input
  .apply(Window.into(Sessions.withGapDuration(Minutes(1)))
    .trigger(AtWatermark())
    .discardingFiredPanes())
  .apply(CalculateWindowLength()));
```






Calculating the Average Session Length







input

```
.apply(Window.into(Sessions.withGapDuration(Minutes(1)))  
      .trigger(AtWatermark())  
      .discardingFiredPanels())  
.apply(CalculateWindowLength()));
```



```
.apply(Window.into(FixedWindows.of(Minutes(2)))  
      .trigger(AtWatermark())  
      .withEarlyFirings(AtPeriod(Minutes(1))))  
      .accumulatingFiredPanels())  
.apply(Mean.globally()));
```





What / Where / When / How

Correctness


Power

Composability


Flexibility

Modularity







1. Classic Batch




2. Batch with Fixed Windows




3. Streaming



4. Streaming with Speculative + Late Data



5. Streaming With Retractions



6. Sessions

What / Where / When / How

Correctness

Power

Composability

Flexibility

Modularity




```
PCollection<KV<String, Integer>> scores = input
    .apply(Sum.integersPerKey());
```

1. Classic Batch

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2))))
    .apply(Sum.integersPerKey());
```

2. Batch with Fixed Windows

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2)))
        .triggering(AtWatermark()))
    .apply(Sum.integersPerKey());
```

3. Streaming

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2)))
        .triggering(AtWatermark())
        .withEarlyFirings(AtPeriod(Minutes(1)))
        .withLateFirings(AtCount(1)))
    .apply(Sum.integersPerKey());
```

4. Streaming with Speculative + Late Data

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(FixedWindows.of(Minutes(2)))
        .triggering(AtWatermark())
        .withEarlyFirings(AtPeriod(Minutes(1)))
        .withLateFirings(AtCount(1)))
        .accumulatingAndRetractingFiredPanels())
    .apply(Sum.integersPerKey());
```

5. Streaming With Retractions

```
PCollection<KV<String, Integer>> scores = input
    .apply(Window.into(Sessions.withGapDuration(Minutes(2)))
        .triggering(AtWatermark())
        .withEarlyFirings(AtPeriod(Minutes(1)))
        .withLateFirings(AtCount(1)))
        .accumulatingAndRetractingFiredPanels())
    .apply(Sum.integersPerKey());
```

6. Sessions

What / Where / When / How

Correctness

Power

Composability

Flexibility

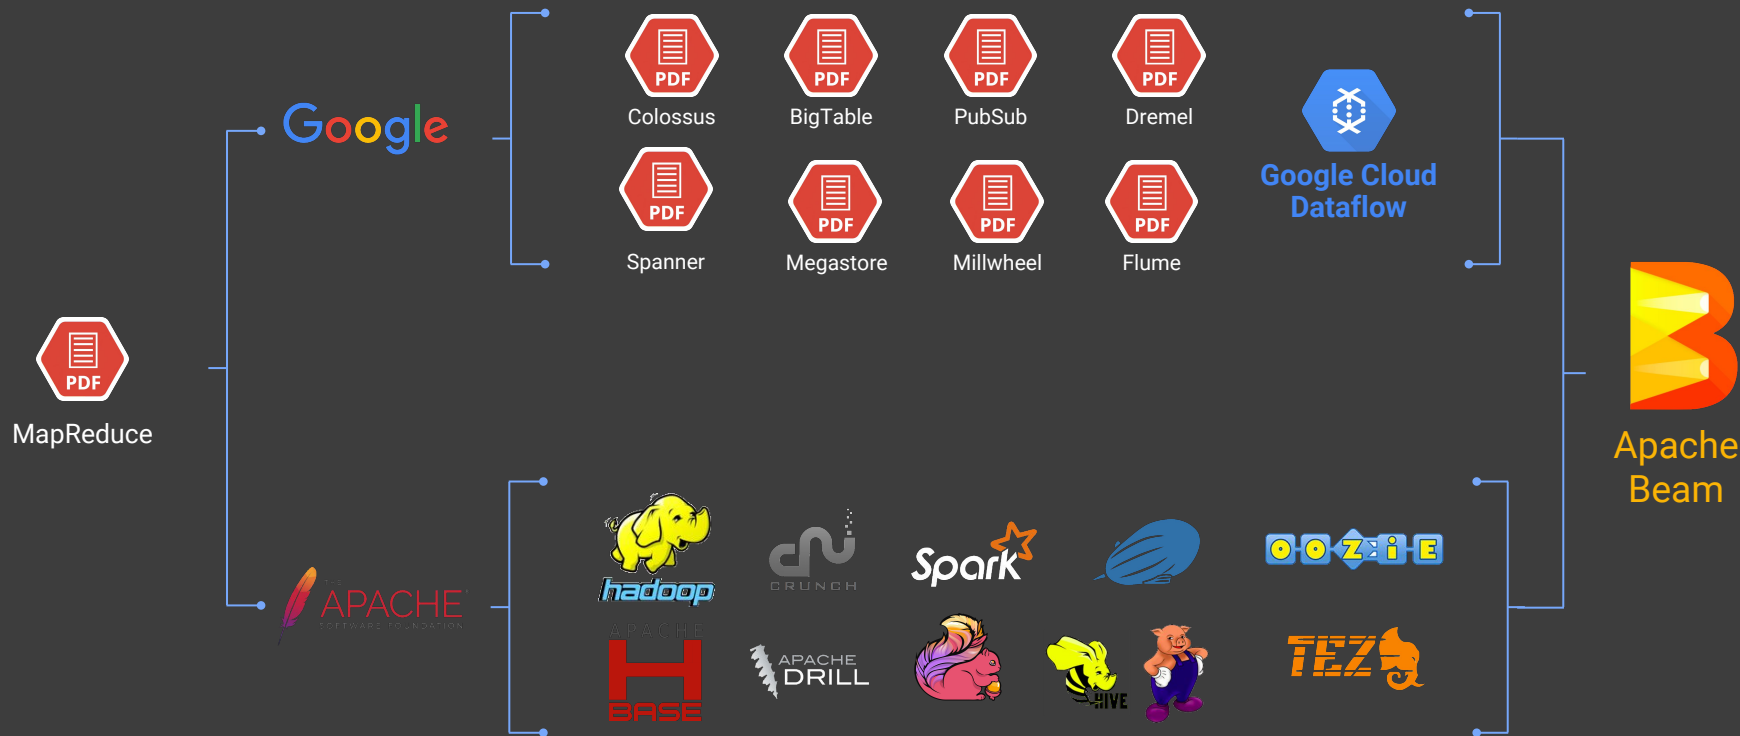
Modularity





Apache Beam (incubating)

The Evolution of Beam



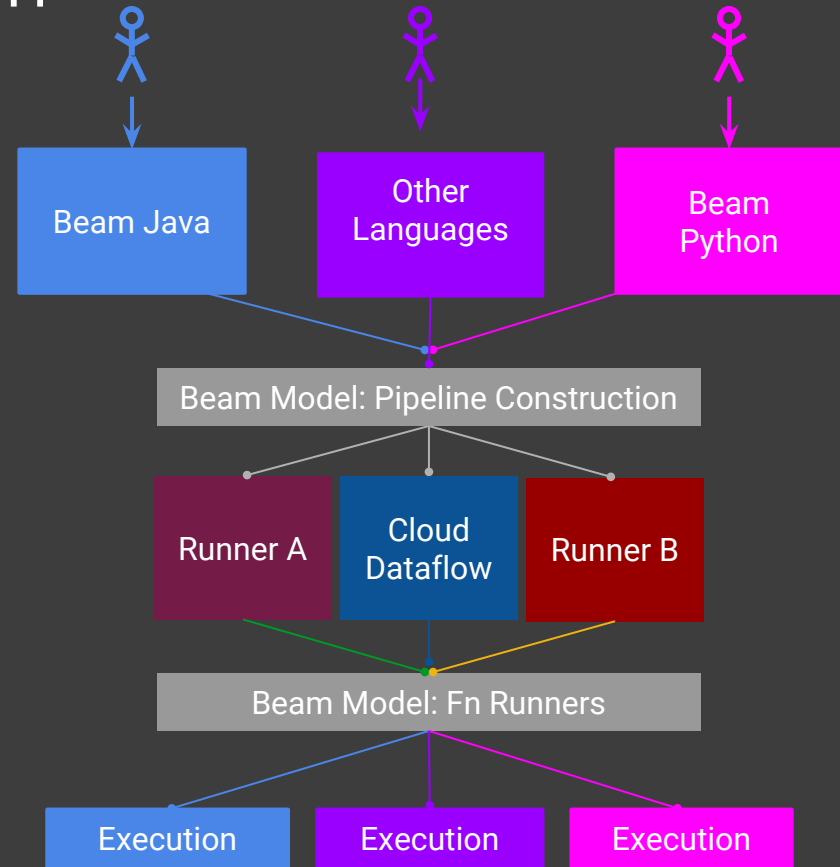
What is Part of Apache Beam?

1. The Beam Model: **What** / **Where** / **When** / **How**
2. SDKs for writing Beam pipelines -- Java and Python
3. Runners for Existing Distributed Processing Backends
 - Apache Flink
 - Apache Spark
 - Google Cloud Dataflow
 - Direct runner for local development and testing
 - In development: Apache Gearpump and Apache Apex

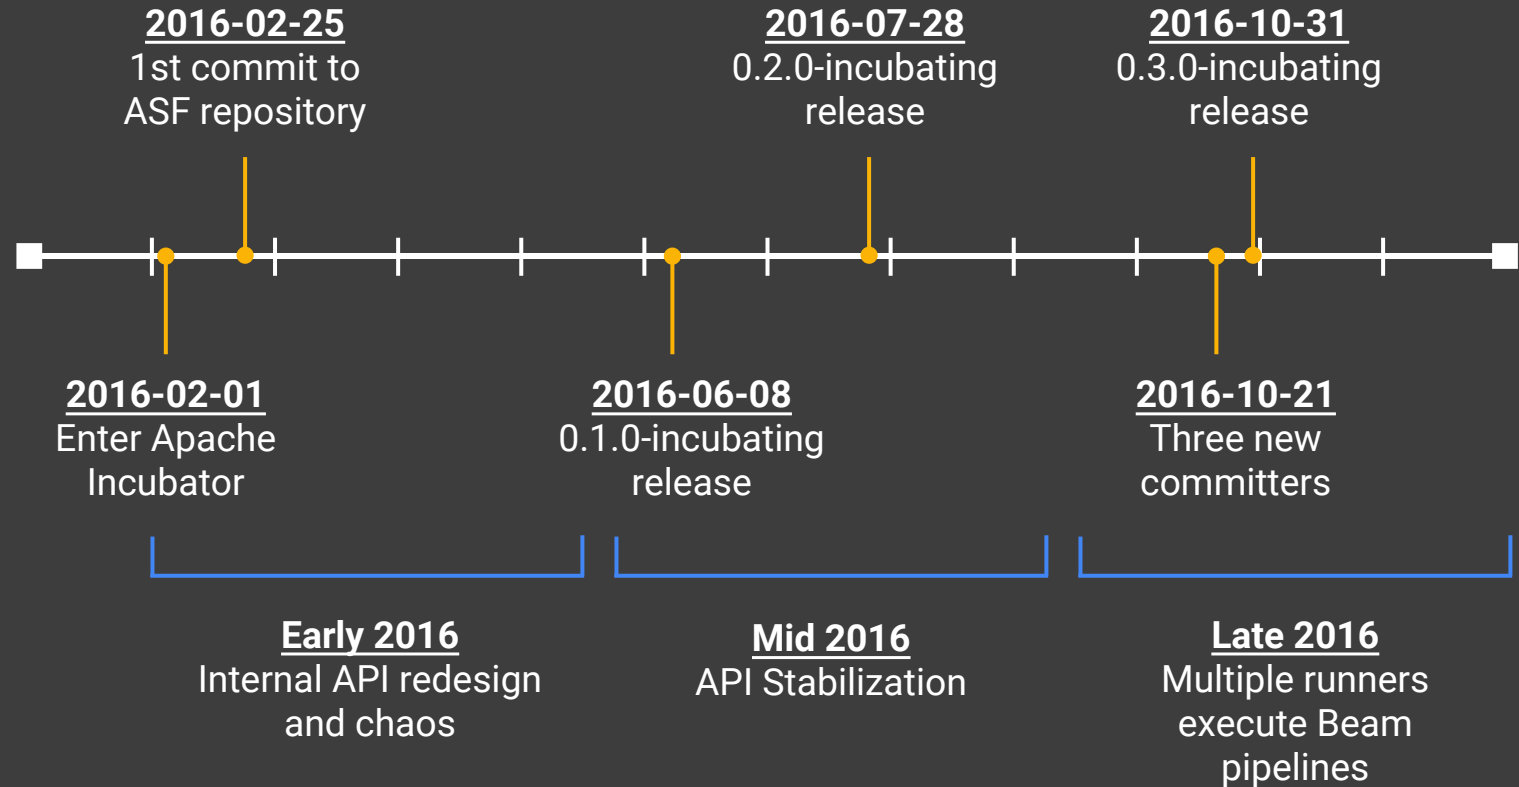


Apache Beam Technical Vision

1. **End users:** who want to write pipelines or transform libraries in a language that's familiar.
2. **SDK writers:** who want to make Beam concepts available in new languages.
3. **Runner writers:** who have a distributed processing environment and want to support Beam pipelines



Visions are a Journey



Categorizing Runner Capabilities

	Beam Model	Dataflow	Flink	Spark	Apex
ParDo	✓	✓	✓	✓	✓
GroupByKey	✓	✓	✓	~	✓
Flatten	✓	✓	✓	✓	✓
Combine	✓	✓	✓	✓	✓
Composite Transforms	✓	~	~	~	~
Side Inputs	✓	✓	✓	✓	✓
Source API	✓	✓	✓	✓	✓
Aggregators	~	~	~	~	✗
Keyed State	✗	✗	✗	✗	✗

	Beam Model	Dataflow	Flink	Spark	Apex
Configurable triggering	✓	✓	✓	✗	✓
Event-time triggers	✓	✓	✓	✗	✓
Processing-time triggers	✓	✓	✓	✓	✓
Count triggers	✓	✓	✓	✗	✓
[Meta]data driven triggers	✗	✗	✗	✗	✗
Composite triggers	✓	✓	✓	✗	✓
Allowed lateness	✓	✓	✓	✗	✓
Timers	✗	✗	✗	✗	✗

	Beam Model	Dataflow	Flink	Spark	Apex
Global windows	✓	✓	✓	✓	✓
Fixed windows	✓	✓	✓	✓	✓
Sliding windows	✓	✓	✓	✓	✓
Session windows	✓	✓	✓	✓	✓
Custom windows	✓	✓	✓	✓	✓
Custom merging windows	✓	✓	✓	✓	✓
Timestamp control	✓	✓	✓	✓	✓

	Beam Model	Dataflow	Flink	Spark	Apex
Discarding	✓	✓	✓	✓	✓
Accumulating	✓	✓	✓	✗	✓
Accumulating & Retracting	✗	✗	✗	✗	✗

<http://beam.incubator.apache.org/documentation/runners/capability-matrix/>

Learn More !

Streaming Fundamentals: The World Beyond Batch 101 & 102

<http://www.oreilly.com/ideas/the-world-beyond-batch-streaming-101>

<http://www.oreilly.com/ideas/the-world-beyond-batch-streaming-102>

Apache Beam (incubating)

<http://beam.incubator.apache.org>

Join the Beam community

user-subscribe@beam.incubator.apache.org

dev-subscribe@beam.incubator.apache.org

Slides for this talk

<http://goo.gl/yzvLXe>

Follow @ApacheBeam on Twitter



Thank you!

