## intel

## Beams, Streams & Finance

Neil Stevenson, Principal Architect, Hazelcast

#### Agenda

What Is Beam?

Beam in Finance

Choosing an Implementation Provider

Quick Demo

Hazelcast & Intel





## What Is Beam ?

#### https://beam.apache.org/

A standard for <u>b</u>atch and str<u>eam</u> processing

Origins from Google, became an Apache project in 2016

Based mainly around the dataflow paradigm

SDK available for Python, Java and GoLang

Multiple implementation providers



Declarative processing, described as a graph

Data flows from stage to stage

Each stage is self-contained, has no knowledge of what is before or after, only sees the data

Stages are stateful or stateless

One input to a stage need not result in one output

No loops









intel



(c) Sink → 7 Write to database, file, etc

(b) Intermediate  $\rightarrow$  2  $\rightarrow$  Filter, compute, enrich, deplete, aggregate

(a) Source Read from database, file, etc

3 building blocks



### What is Beam ?

Parallelized!

One processing job, two instances

Data can flow from host to host



HOST 2













#### Tip:

Some sources are infinite – eg. temperature sensor Some sources are finite – read a database table

But \*think\* of a bounded source as unbounded

Then a trigger or CDC on a database table can give a continuous feed of changes





## **Beam In Finance**

#### Reporting

Any type of "*evaluate all X against Y scenario*" type calculation. Partition the *X* across hosts to achieve the SLA. Add more hosts when Y goes up, such as month end.

#### Alerting

Watch the markets to detect trends in prices, eg. Bollinger Bands => issue alerts when trending changes from up->down or down->up

#### Management

Dynamically re-value portfolios as prices & forex rates fluctuate => force sell if over-exposed to to one market





Credit Value Adjustment

(a) Interest Rate Swap trades(b) Interest Rate curves

Simplistically:

Calculate the Mark-To-Market for each trade with each curve Convert to the MTM to a CVA exposure Discard out-of-the-money Calculate the average CVA exposure per trade Sum these per counterparty



Credit Value Adjustment

#### In Beam terms:

Combine every (a) and every (b) Process these combinations to produce (c) Transform (c) to (d) Filter some (d) Average (d) by primary key giving (e) Sum (e) by foreign key



Credit Value Adjustment

In Beam terms:

PCollection<D> d
 = c.apply(ParDo.of(new TransformCintoD()));
 .apply(ParDo.of(new TransformAintoKV()));
PCollectionView<Map<String, B>> bView
= pipeline
 .apply(TextIO.read().from("./curves.csv"))
 .apply(ParDo.of(new TransformBintoKV()))
 .apply(View.asMap());
PCollection<C>.c

= a.apply(ParDo.of(new BusinessLogic(bView)).withSideInputs(bView))

Combine every (a) and every (b) Process these combinations to produce (c) Transform (c) to (d) Filter some (d) Average (d) by primary key giving (e) Sum (e) by foreign key



hazelcast

Credit Value Adjustment

#### In Beam terms:

Combine every (a) and every (b) Process these combinations to produce (c) Transform (c) to (d) Filter some (d) Average (d) by primary key giving (e) Sum (e) by foreign key



Credit Value Adjustment



In Beam terms:

Combine every (a) and every (b) Process these combinations to produce (c) Transform (c) to (d) Filter some (d) Average (d) by primary key giving (e) Sum (e) by foreign key



Credit Value Adjustment

```
PCollection<D> d2
= d.apply(Filter.by(new SerializableFunction<D, Boolean>() {
        public Boolean apply(D d) {
            return d.value > 0;
        }
    }));
```

In Beam terms:

Combine every (a) and every (b) Process these combinations to produce (c) Transform (c) to (d) Filter some (d) Average (d) by primary key giving (e) Sum (e) by foreign key

© HAZELCAST | 18



hazelcast

Credit Value Adjustment

PCollection<KV<String, E>> e =
 d2
 .apply(ParDo.of(new TransformDintoKV()))
 .apply(GroupByKey.create())
 .apply(ParDo.of(new TransformDintoE()));

In Beam terms:

Combine every (a) and every (b) Process these combinations to produce (c) Transform (c) to (d) Filter some (d) Average (d) by primary key giving (e) Sum (e) by foreign key



Credit Value Adjustment

e
.apply(ParDo.of(new TransformEintoF()))
.apply(GroupByKey.create())
.apply(ParDo.of(new TransformFintoString()))
.apply(TextIO.write().to("./results.txt"));

In Beam terms:

Combine every (a) and every (b) Process these combinations to produce (c) Transform (c) to (d) Filter some (d) Average (d) by primary key giving (e) Sum (e) by foreign key



Credit Value Adjustment

#### BUT

When we did a CVA example at Hazelcast, we discarded Beam in favor of directly coding in Jet

Because

We could optimize it more, for intraday CVA





## **Choosing An Implementation Provider**

Beam is an SDK or API to define a job in a implementation neutral way

A "Beam runner" is software that is an execution environment for such a job, dealing with parallelism, data marshalling and general runtime

Using Beam makes it easier

- developers are easier to find, a standard skillset
- no vendor lock-in

- using Beam now allows projects to be developed in parallel to strategic projects

- (change from Beam runner A to Beam runner B)



#### Various exist (12):

- Hazelcast Jet
- Google Cloud Dataflow
- Apache Spark
- Apache Flink
- etc

#### All the usual considerations

- Operational familiarity, performance, licensing model, cost, export restrictions, cloud enabled, developer preference

PLUS one major and one minor caveat.....



The minor caveat - Input & Output

There are standard connectors for reading files, writing files - Baked in assumptions on directory paths, cross-mounts may need a little adjustment when changing Beam runner

Some standard connectors for specific set-ups - BigQueryIO, SpannerIO for Google source/sink



The major caveat

- all Beam runners do not implement all features

#### https://beam.apache.org/documentation/runners/capability-matrix

	Beam Model	Google Cloud Dataflow	Apache Flink	Apache Spark (RDD/DStream based)	Apache Spark Structured Streaming (Dataset based)	Apache Hadoop MapReduce	JStorm	IBM Streams	Apache Samza	Apache Nemo	Hazelcast Jet	Twister2
ParDo	1	×	<	<	~	<	<ul> <li>✓</li> </ul>	1	<ul> <li>✓</li> </ul>	<ul> <li>✓</li> </ul>	1	×
GroupByKey	<	×	<ul> <li>✓</li> </ul>	~	~	<	<ul><li>✓</li></ul>	<ul> <li>✓</li> </ul>	<	<ul><li>✓</li></ul>	1	×
Flatten	1	×	<ul> <li>✓</li> </ul>	<	~	<	<ul> <li>✓</li> </ul>	<ul> <li>✓</li> </ul>	<	<ul> <li>✓</li> </ul>	1	×
Combine	1	1	1	<	~	<	1	1	<ul> <li>✓</li> </ul>	<ul> <li>✓</li> </ul>	1	<
Composite Transforms	<	~	~	~	~	<	1	~	~	<	~	~
Side Inputs	<	1	✓	✓	~	<	1	<ul> <li>✓</li> </ul>	<	<ul> <li>✓</li> </ul>	~	<
Source API	<	1	<ul><li>✓</li></ul>	✓	~	~	1	<ul> <li>✓</li> </ul>	<ul> <li>✓</li> </ul>	<ul><li>✓</li></ul>	1	<
Splittable DoFn (SDF)	~	×	1	~	×	×	×	×	~	×	×	~
Metrics	~	~	~	~	~	~	~	~	~	×	~	×
Stateful Processing	1	~	~	~	×	~	~	~	~	×	~	×



The major caveat

- all Beam runners do not implement all features

https://beam.apache.org/documentation/runners/capability-matrix

You need to know what features you need, now and in the future



Pragmatic points:

No one ever changes implementation provider

**Blurred distinction** 

- Be careful:

Developers will put "optimizations" into Beam SDK code that rely on the execution environment

Performance

- If you can:

POC without business logic as Beam and without Beam on your runner





## Demo – "VaR"

An outline of VaR

- (1) Take all portfolios
- (2) Take all prices
- (3) For all combinations
- (4) Simulate a range of possible movements
- (5) Select 5% worst, 1% worst, whatever



intel

An outline of VaR

- (1) Take all portfolios
- (2) Take all prices
- (3) For all combinations
- (4) Simulate a range of possible movements
- (5) Select 5% worst, 1% worst, whatever

```
for (int i = 0 ; i < portfolios.length ; i = i+1) {
    double sum = 0;
    for (int j = 0 ; j < prices.length; j = j+1) {
        sum = sum + businessLogic(portfolios[i], prices[j]);
    }
}</pre>
```





An outline of VaR

(1) Take all portfolios
(2) Take all prices
(3) For all combinations



An outline of VaR

(1) Take all portfolios
(2) Take all prices
(3) For all combinations



PCollection<KV<String, Account>> account

= pipeline

.apply("bounded-account-source", Read.from(new MyAccountSource()))
.apply("account-string-to-entry", ParDo.of(new MyAccountTransform()));



An outline of VaR

(1) Take all portfolios
(2) Take all prices
(3) For all combinations



PCollection<KV<String, Price>> prices

= pipeline

.apply("bounded-price-source", Read.from(new MyPriceSource()))
.apply("price-string-to-entry", ParDo.of(new MyPriceTransform()));



An outline of VaR

(1) Take all portfolios (2) Take all prices (3) For all combinations



intel

hazelcast

PCollectionView<Map<String,Price>> pricesView = prices .apply("price-entry-viewmap", View.asMap());

PCollection<Valuation> joined

= account.apply("join", ParDo.of(new MyJoinTransform(pricesView)).withSideInputs(pricesView))

#### Like an outer join

An outline of VaR

- (1) Take all portfolios
- (2) Take all prices
- (3) For all combinations
- (4) Simulate a range of possible movements .. Historical returns, Monte Carlo
- (5) Select 5% worst, 1% worst, whatever

This is not a "for loop", it is a cartesian product == > RUNNING IN PARALLEL



Two hosts....

(1) Take all portfolios(2) Take all prices(3) For all combinations



This is not a "for loop", it is a cartesian product == > RUNNING IN PARALLEL



Two hosts....

(1) Take all portfolios(2) Take all prices(3) For all combinations



This is not a "for loop", it is a cartesian product == > RUNNING IN PARALLEL





## Demo



## Demo



## Hazelcast & Intel

#### Hazelcast



An In-Memory Data Grid with a stream processing engine (Jet)

The data grid is collection of Java processes caching data in memory

Hazelcast provides reliability, securability, scalability, performance Widely used by most of the world's large organisations on cloud or premise Multiple client connectors – Java, Python, C++, C#, GoLang, Node.js..

(https://upload.wikimedia.org/wikipedia/commons/9/94/Filing cabinet icon.svg)







Hybrid memory

Packaged as a DIMM, can use as memory

- up to 512GB, larger than typical DRAM cards
- a little slower than DRAM but wins on cost/GB



- 6x faster than SSD

- why is this relevant to Hazelcast, an in-memory data grid?



intel

hazelcast

+ Hardware encryption!



Using Optane for save/restore file storage

Reload time for data:

87GB in 19 seconds



(5.5 million records, 16Kb each + metadata)





#### Using Optane for save/restore file storage

Reload time for data:

87GB in 19 seconds

12 node cluster, 1TB of data

Still 19 seconds!

(+ integrity cross-checks)







Using Optane for save/restore file storage

In the real world...

A few minutes for off, on, verify, ready an <u>empty</u> cache + however long to reload from legacy before useable

Now...

Same few minutes, but now can have 1TB of data!







Using Optane for save/restore file storage

In the real world...

A few minutes for off, on, verify, ready an <u>empty</u> cache + however long to reload from legacy before useable



Now...

Same few minutes, but now can have 1TB of data! Or 2TB!





Using Optane for save/restore file storage

In the real world...

A few minutes for off, on, verify, ready an <u>empty</u> cache + however long to reload from legacy before useable

	hazelcast imdg	hazelcast imdg			hazelcasť imdg	📘 hazelcast imdg 📃	
	hazelcast imdg	hazelcast IMDG	◄		hazelcast' IMDG	hazelcasť imdg 🚽	
	hazelcast imdg	hazelcast' IMDG	◄		hazelcast IMDG	📙 hazelcast imdg 📃	
<b>=</b>	hazelcast imdg	hazelcast <sup>®</sup> IMDG		┢┣┫╛	hazelcast imdg	📑 hazelcast imdg 📃	
	hazelcast imdg	hazelcast' IMDG			hazelcast imdg	📑 hazelcast imdg 📃	
▤ ≠	hazelcast' IMDG	hazelcast' IMDG			hazelcast imdg	📙 hazelcast' IMDG 🛛 🖊	
	hazelcast imdg	hazelcast' IMDG			hazelcasť imdg	📑 hazelcast imdg 🚽	
	hazelcast imdg	hazelcast' IMDG			hazelcast IMDG	📑 hazelcast imdg 📃	
	hazelcast imdg	hazelcast' IMDG			hazelcast' IMDG	📑 hazelcast imdg 📃	
	hazelcast' IMDG	hazelcast imdg			hazelcast imdg	📑 hazelcast imdg 📃	
	hazelcast imdg	hazelcast' IMDG	◄		hazelcast imdg	📑 hazelcast imdg 📃	
	hazelcast imdg	hazelcast' IMDG	1		hazelcast imdg	🖡 hazelcast' IMDG 📃	

Now...

Same few minutes, but now can have 1TB of data! Or 2TB! Or 4TB!...etc





#### Hazelcast

Want to try it ?

A Hazelcast & Beam example, tracking the movement of a train from a stream of GPS points



Look for "Train Track" on <a href="https://github.com/hazelcast/hazelcast-jet-demos">https://github.com/hazelcast/hazelcast-jet-demos</a>





# Summary



Imagine your Batch as a Stream, one day it may be

Performance of Beam runners and Beam overhead may matter

Not all Beam runners have complete implementations, may matter

Monitoring/metrics are vital, debugging is difficult

Senior developers only: deplete > enrich

Hazelcast & Intel, other runners are available





# Thank you