



SPARK+AI
SUMMIT EUROPE

Spline: Data Lineage for Spark Structured Streaming


Marek Novotny, ABSA
Vaclav Kosar, ABSA

#SAISExp18

About Us



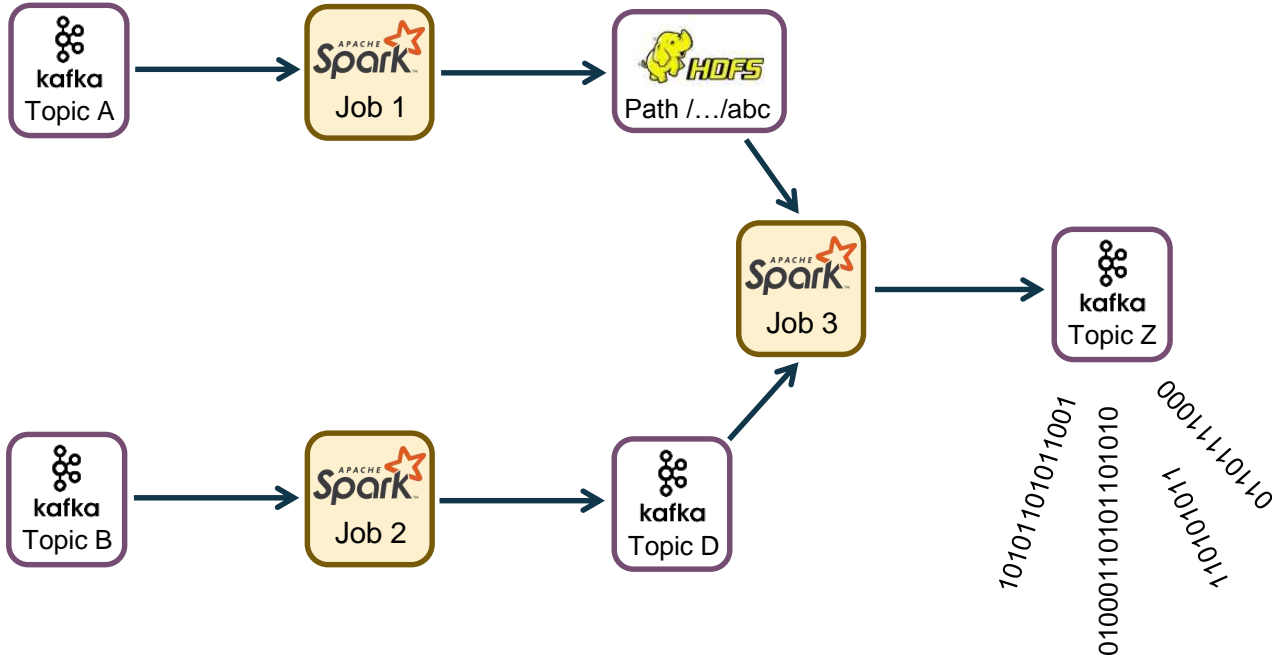
- ABSA is a Pan-African financial services provider
 - With Apache Spark at the core of its data engineering
- We try to fill gaps in the Hadoop eco-system
- Contributions to Apache Spark
- Spark-related open-source projects (github.com/AbsaOSS)
 - **ABRiS** – Avro SerDe for structured APIs (#SAISDev5)
 - **Cobrix** – Cobol data source
 - **Atum** – Completeness and accuracy library
 - **Spline** – Data lineage tracking and visualization tool (#EUent3)



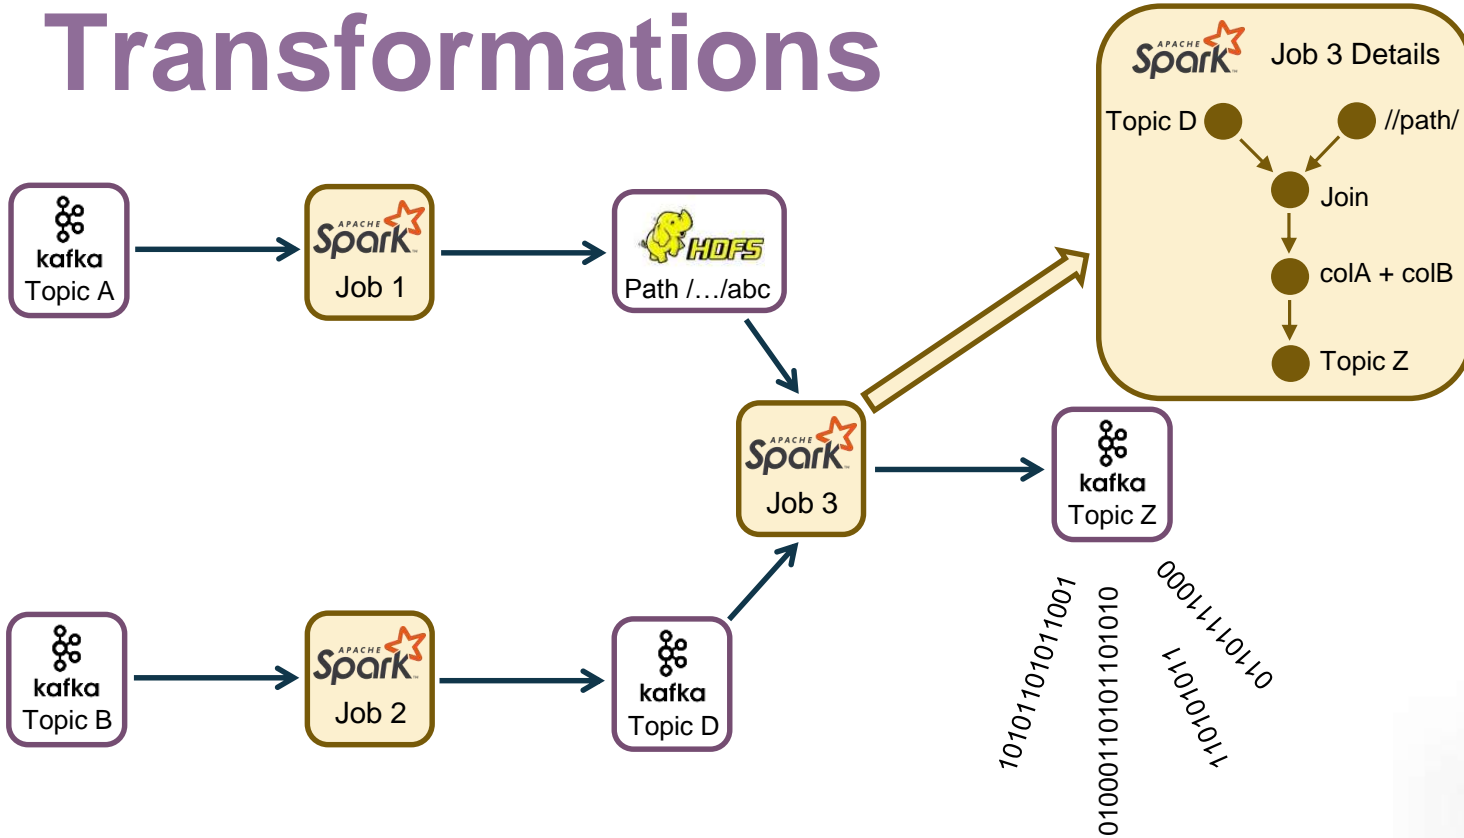
- How data is calculated?
- What is the schema and format of streamed data?



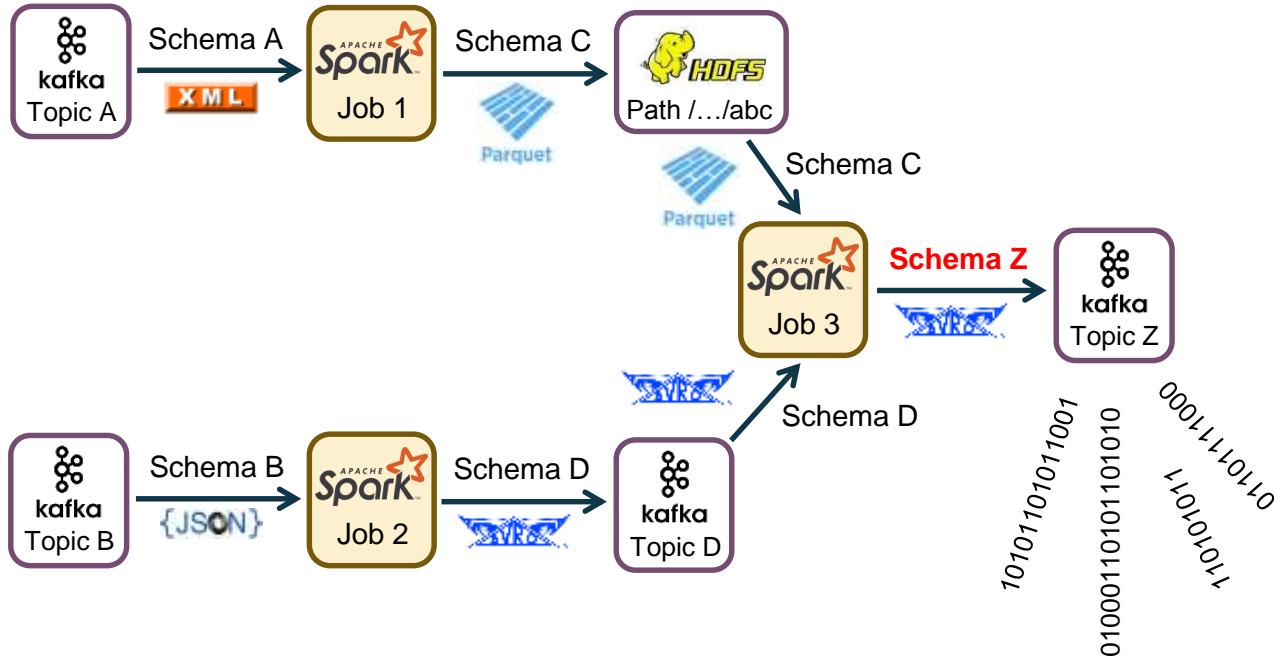
Data Flow



Transformations



Schemas and Formats






- To make Spark BCBS (Clarity) compliant
- To communicate with business people



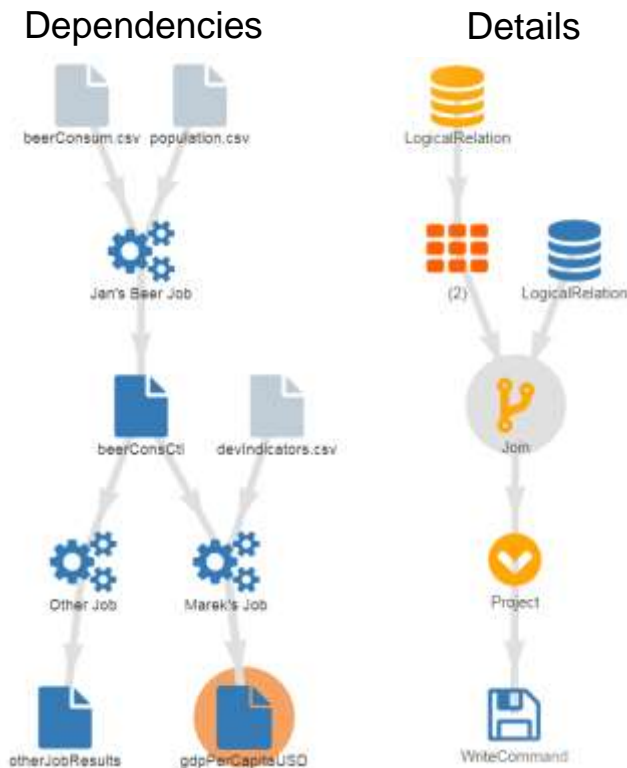
- To make Spark BCBS (Clarity) compliant
- To communicate with business people
- Online documentation of — Job dependencies

Dependencies

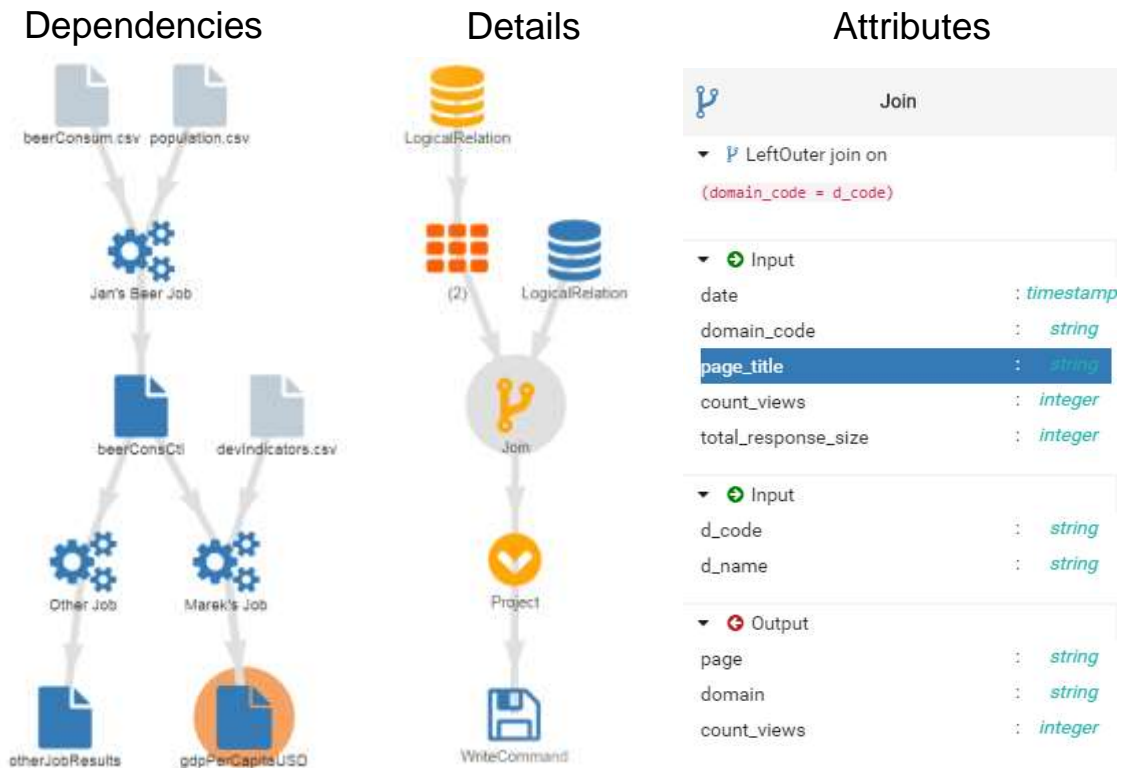




- To make Spark BCBS (Clarity) compliant
- To communicate with business people
- Online documentation of
 - Job dependencies
 - Particular Spark SQL jobs

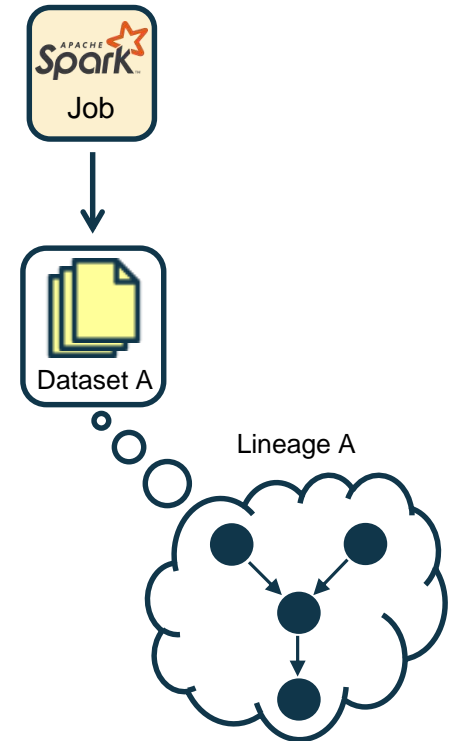


- To make Spark BCBS (Clarity) compliant
- To communicate with business people
- Online documentation of
 - Job dependencies
 - Spark SQL job details
 - Attributes occurring in the logic



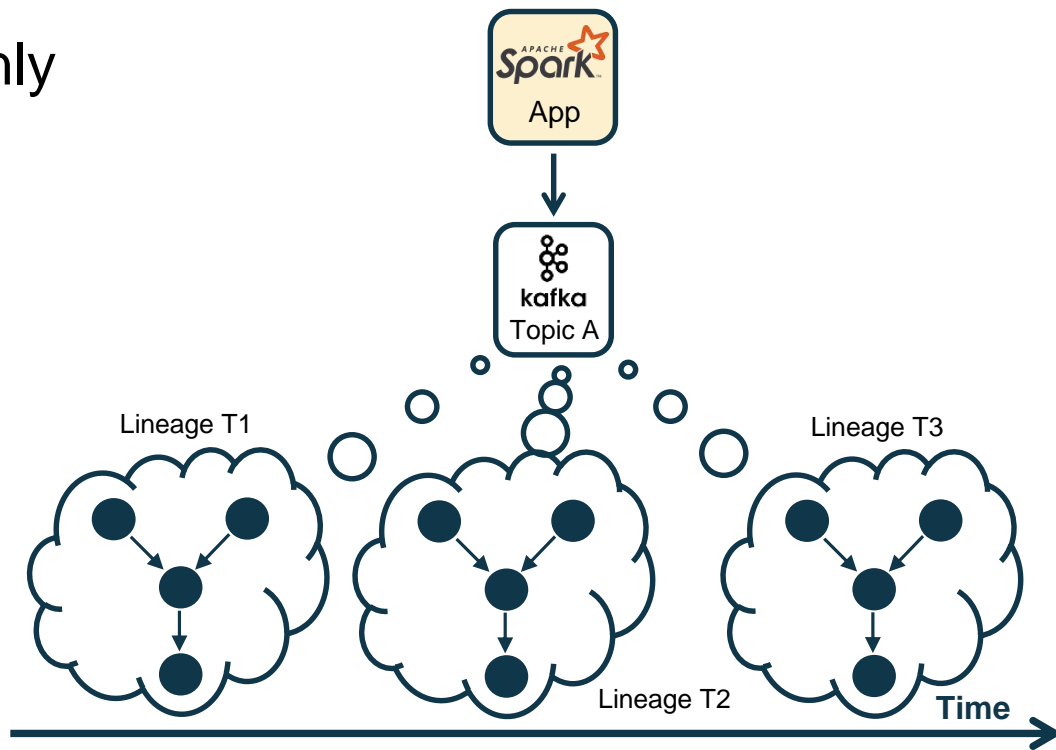
Lineage Tracking of Batch Jobs

- Dataset-oriented
- Leverages execution plans
- Structured APIs only
 - SQL
 - Dataframes
 - Datasets
- UDFs and lambdas are considered as black boxes



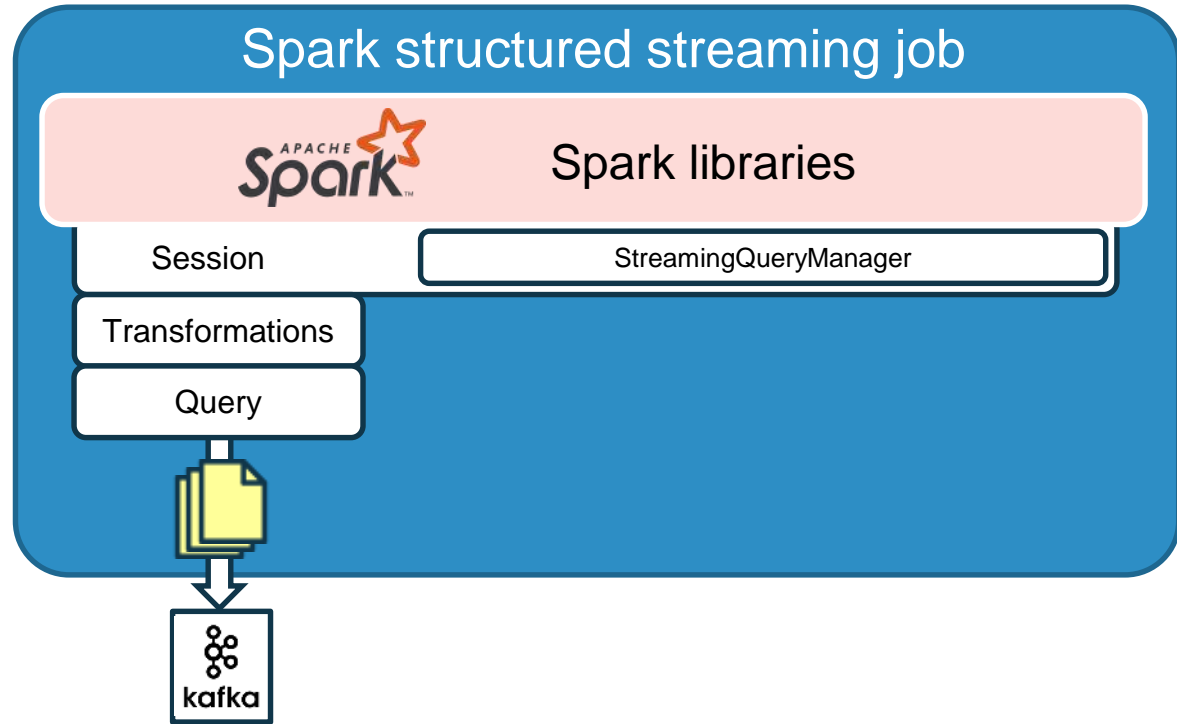
Lineage Tracking of Streaming Jobs

- Structured Streaming only
- Source-oriented (topic)
- Evolves in time



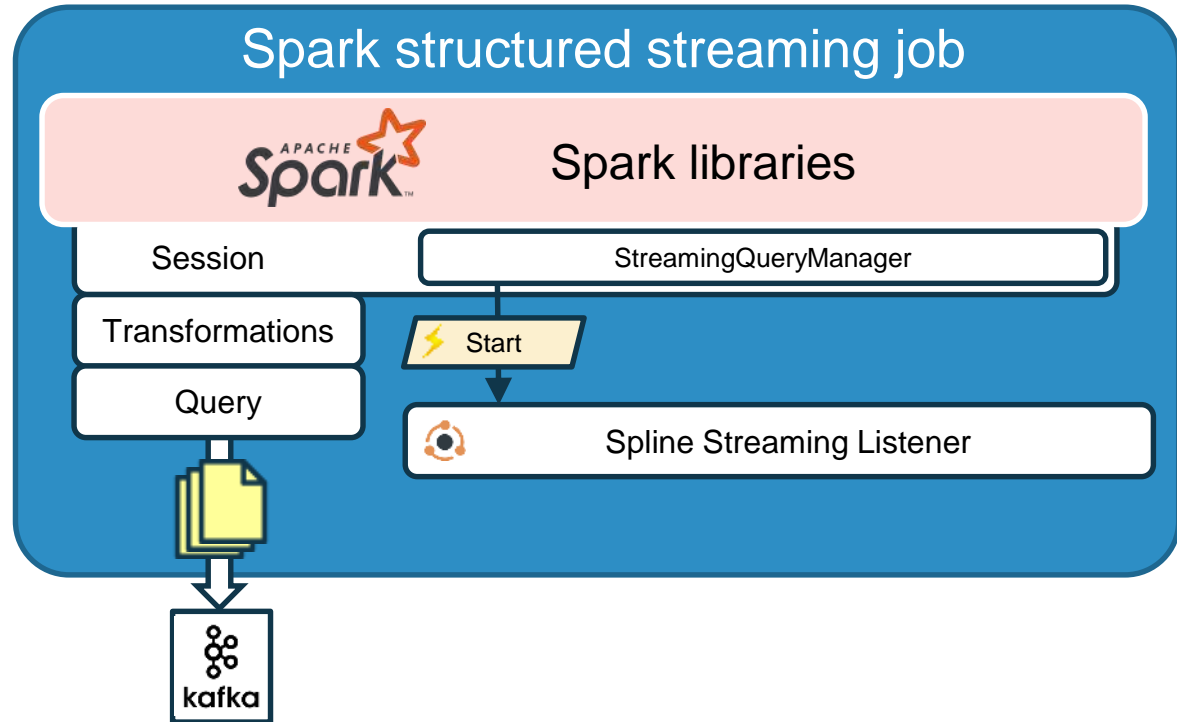
Structured Streaming Support

- StreamingQueryManager



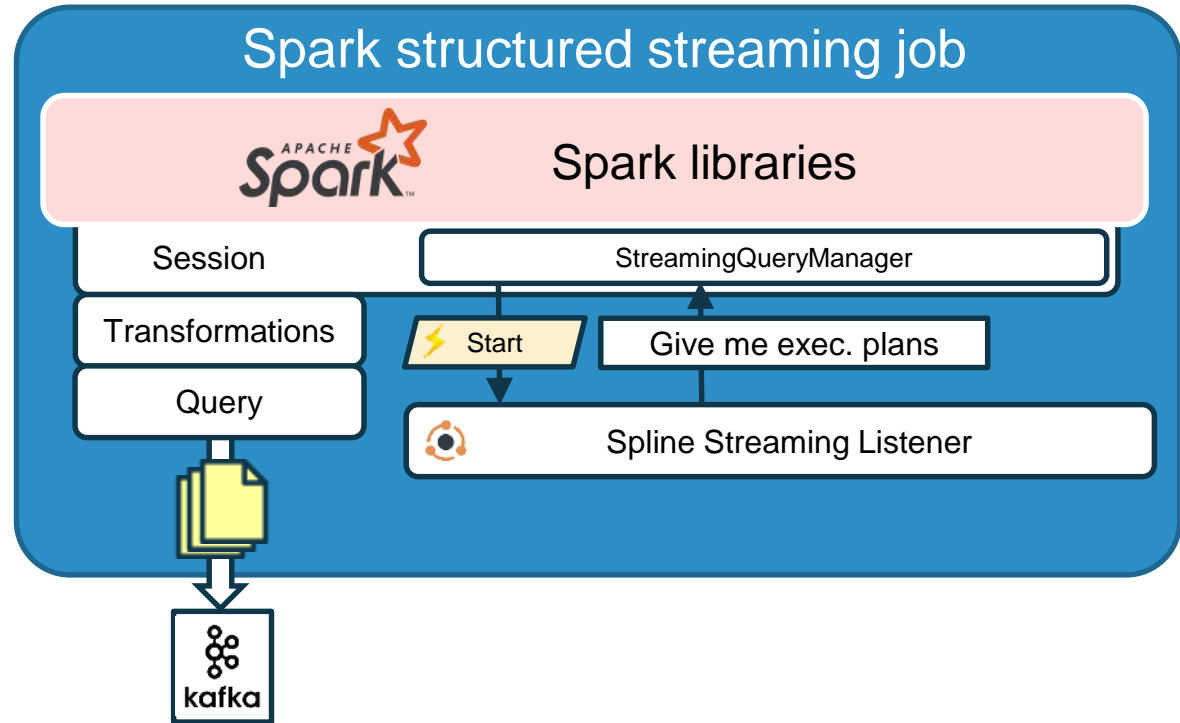
Structured Streaming Support

- StreamingQueryManager
 - Information about start



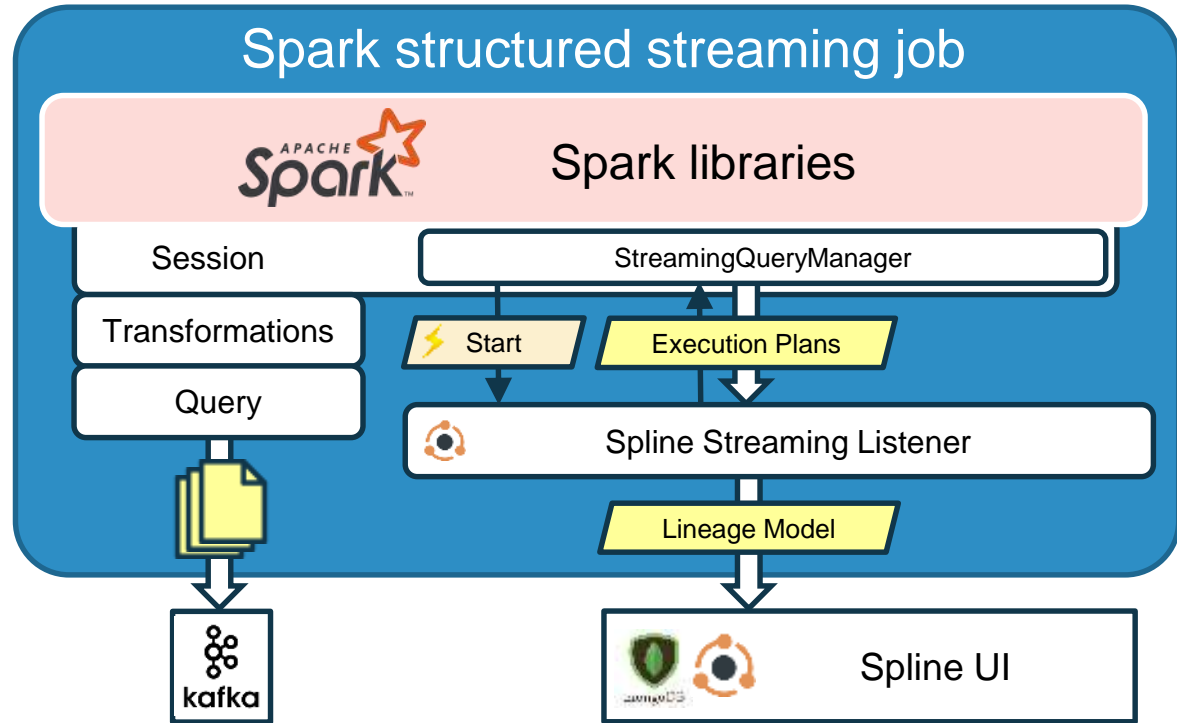
Structured Streaming Support

- StreamingQueryManager
 - Information about start
 - Can provide execution plans



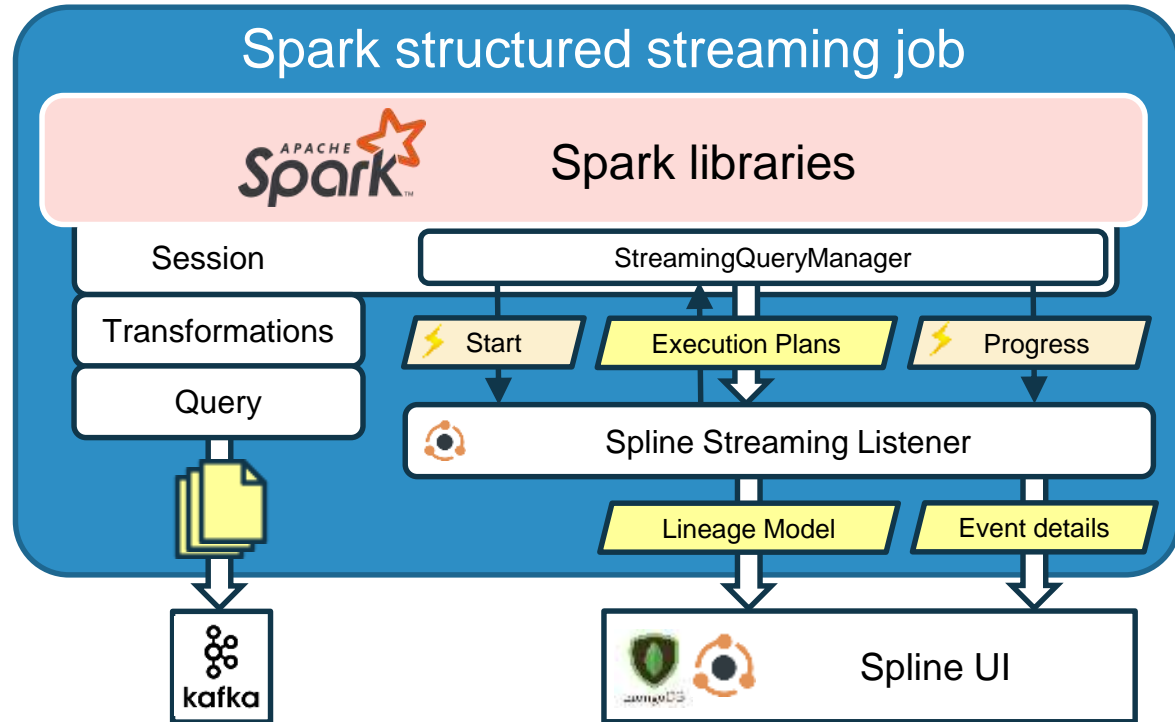
Structured Streaming Support

- StreamingQueryManager
 - Information about start
 - Can provide execution plans



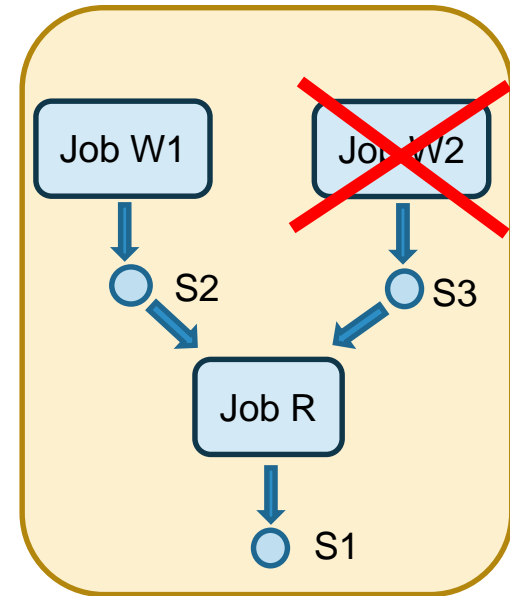
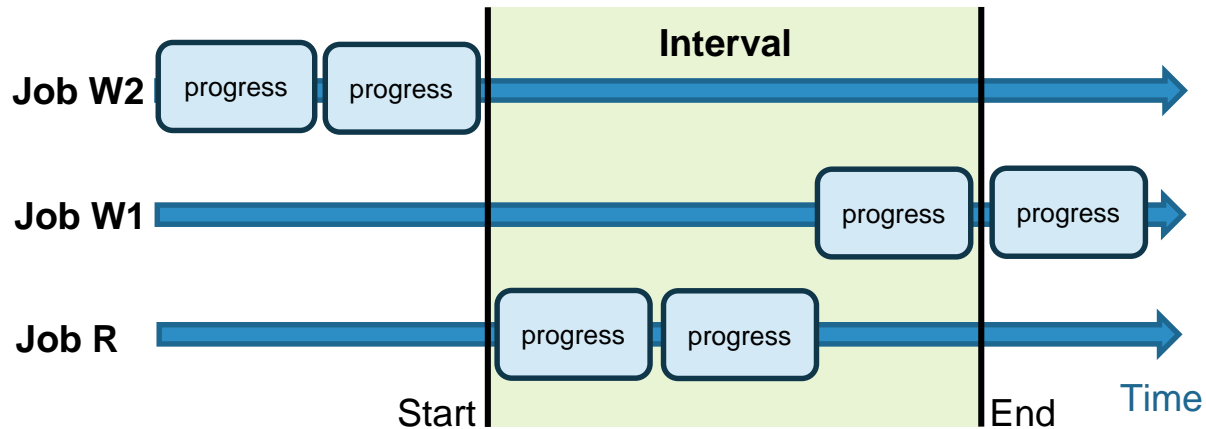
Structured Streaming Support

- StreamingQueryManager
 - Information about start
 - Can provide execution plans
 - Information about progress
 - MicroBatch



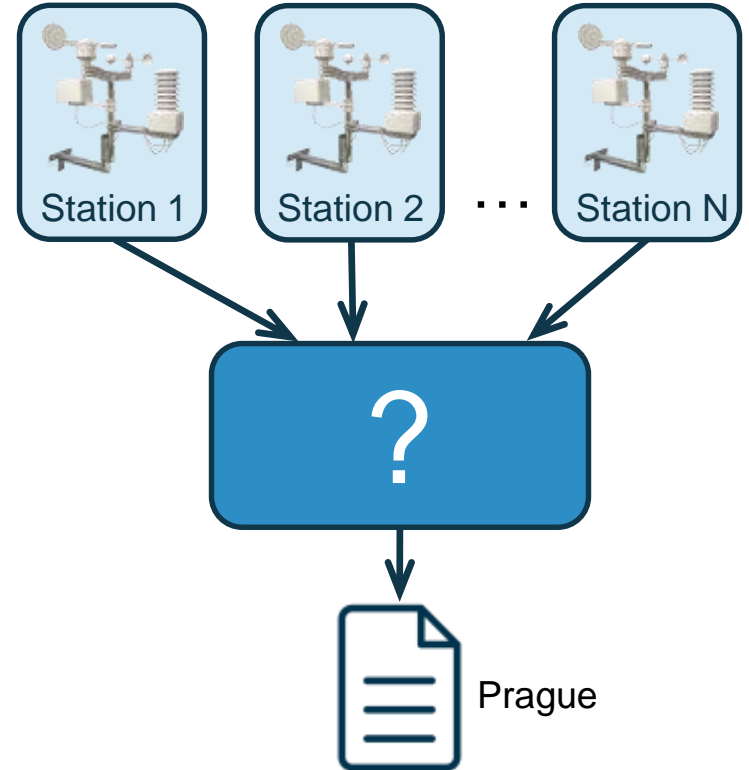
Interval View

- Displays data flow in fixed interval



Demo – Use Case


What is temperature per hour in Prague?



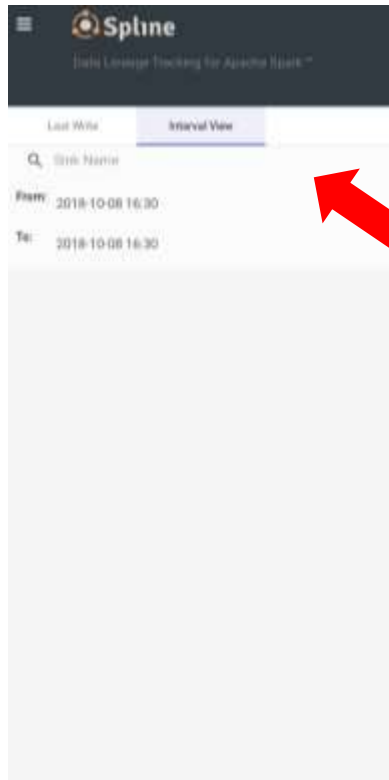
Demo – Use Case Output



2018-09-24



Demo – Select Interval View




Welcome to Spline - Data Lineage Tracking for Apache Spark™

Please select a dataset from the left-hand side menu to display the data lineage for

Spline v1.4.0 (20181027)

Demo – Select Interval



Last View Interval View

🔍 Sink Name

From: 2018-09-24 00:00

To: ◀ ▶ **SEP 2018** ▶ ⌵

Mo	Tu	We	Th	Fr		
Sa	Su	1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	1	2	3

Time: Current Time

Hours:

Minutes:


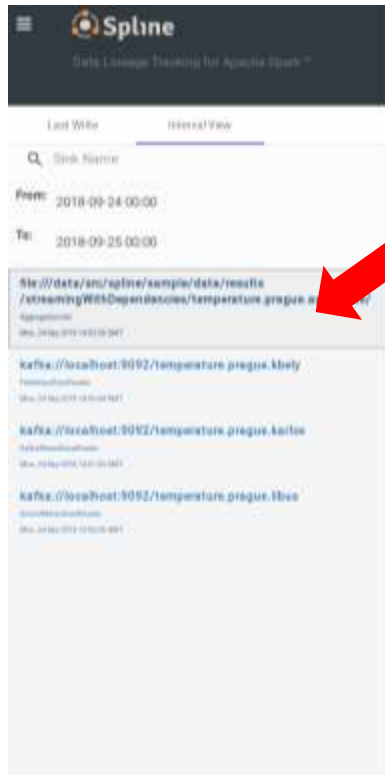


Welcome to Spline - Data Lineage Tracking for Apache Spark™

Please select a subject from the left hand side menu to display the data lineage for

Spline v4.2.0-09/24/18

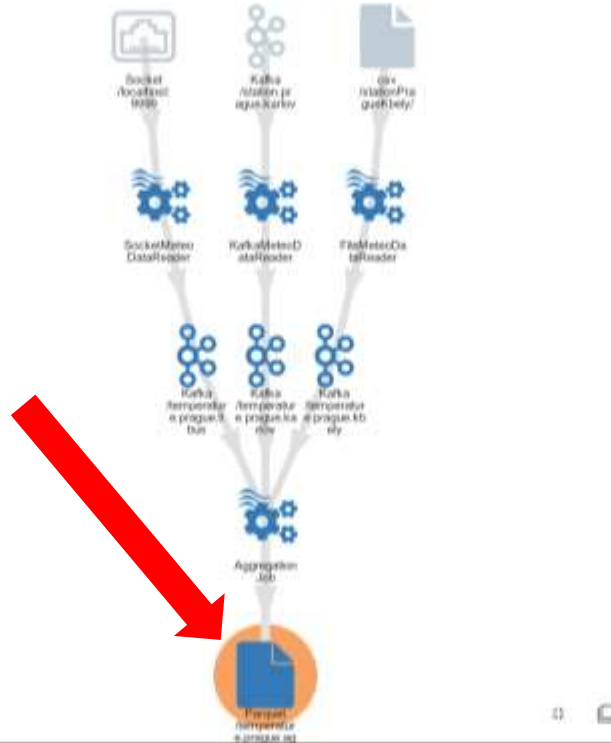
Demo – Select Sink



Demo – Find Highlighted Sink

The screenshot shows the Spline web interface. At the top, there is a header with the Spline logo and the tagline "Data Lineage Tracking for Apache Spark™". Below the header, there are two tabs: "Last Write" and "Interest View". A search bar labeled "Sink Name" is present. The main content area displays a list of sinks with the following details:

- From:** 2018-09-24 00:00
- To:** 2018-09-25 00:00
- Sink Name:** file:///data/spl/spine/sample/data/results/streamingWithDependencies/temperature.prgue.aggregate/
- Source:** kafka://localhost:9092/temperature.prgue.kbety
- Source:** kafka://localhost:9092/temperature.prgue.kafwa
- Source:** kafka://localhost:9092/temperature.prgue.kbwa



The screenshot shows the details of a specific sink in the Spline web interface. The sink name is "file:///data/spl/spine/sample/data/results/streamingWithDependencies/temperature.prgue.aggregate/". The type is "Parquet" and the timestamp is "2018-09-24, 4:03:58 PM". The schema is displayed as follows:

Schema:	Value
date	date
hour	string
temperature	double

Demo – Review The Lineage

Spline
Data Lineage Tracking for Apache Spark™

Last Write: 2018-09-24 00:00
Interact View

File Name: Sink Name

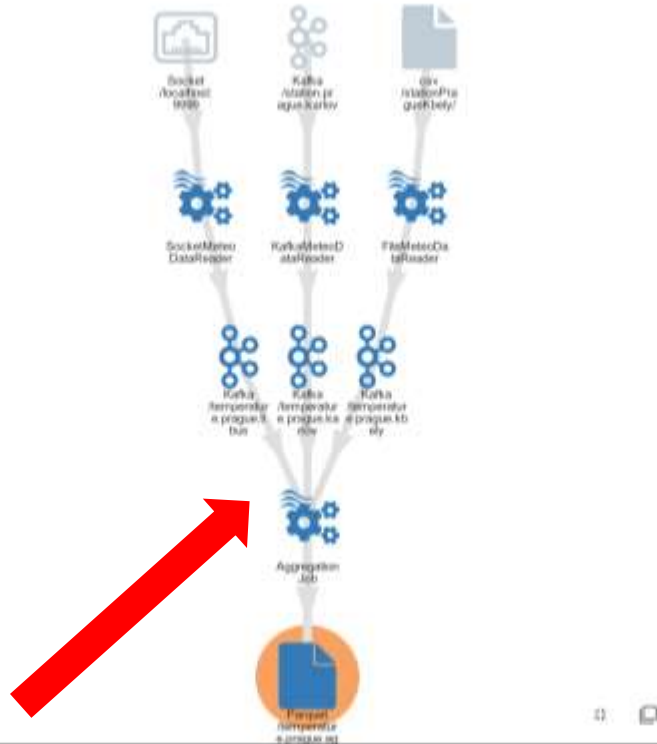
From: 2018-09-24 00:00
To: 2018-09-25 00:00

File: `file:///data/spl/sp/sample/data/resultstreamingWithDependencies/temperature.prgue.aggregate/`
Aggregation Job
File: `file:///data/spl/sp/sample/data/resultstreamingWithDependencies/temperature.prgue.aggregate/`

Kafka://localhost:9092/temperature.prgue.kbety
TemperatureDataLoader
File: `file:///data/spl/sp/sample/data/resultstreamingWithDependencies/temperature.prgue.kbety`

Kafka://localhost:9092/temperature.prgue.kafwa
KafkaTemperatureDataLoader
File: `file:///data/spl/sp/sample/data/resultstreamingWithDependencies/temperature.prgue.kafwa`

Kafka://localhost:9092/temperature.prgue.kbwa
KafkaTemperatureDataLoader
File: `file:///data/spl/sp/sample/data/resultstreamingWithDependencies/temperature.prgue.kbwa`



File: `file:///data/spl/sp/sample/data/resultstreamingWithDependencies/temperature.prgue.aggregate/`

Type: Parquet
Timestamp: Sep 24, 2018, 4:03:58 PM

Schema:

- date: date
- hour: string
- temperature: double

Spline v1.4.2-2018092402

Demo – Change The Interval

Spline
Early Licensee: Testing for Apache Spark™

Last View: Interval View

Search: Sink Name


From: 2018-09-24 00:00
To: 2018-09-25 00:00

file:///data/spl/spline/sample/data/results/streamingWithDependencies/temperature.prgue.aggregate/
Aggregation Job
Sep 24, 2018 03:58 PM

kafka://localhost:9092/temperature.prgue.kbety
TemperatureKbety
Sep 24, 2018 03:58 PM

kafka://localhost:9092/temperature.prgue.kafwa
TemperatureKafwa
Sep 24, 2018 03:58 PM

kafka://localhost:9092/temperature.prgue.kbwa
TemperatureKbwa
Sep 24, 2018 03:58 PM



file:///data/spl/spline/sample/data/results/streamingWithDependencies/temperature.prgue.aggregate/

Type: Parquet
Timestamp: Sep 24, 2018, 4:03:58 PM

Schema:

date | date
hour | hour
temperature | temperature

Spline v1.4.2-090913422

Demo – Observe New Lineage

The screenshot displays the Spline interface for observing data lineage. On the left, a search bar shows the file path: `file:///data/emr/spline/sample/data/results/streamingWithDependencies/temperature.pregus.aggregate/`. A red arrow points from this search result to a central lineage diagram. The diagram illustrates the data flow: two source files, `csv/temperature.csv` and `csv/temperature.csv`, are processed by `KafkaMessageDistributer` and `KafkaMessageLoader` respectively. These feed into three Kafka topics: `kafka-temperature-1`, `kafka-temperature-2`, and `kafka-temperature-3`. The data then flows through an `Aggregation Job` to produce the final `Parquet/temperature.pregus.aggregate.parquet` file. On the right, a detailed view of the output file shows its type as `Parquet`, timestamp as `Sep 24, 2018, 4:03:58 PM`, and schema with columns: `date` (date), `hour` (string), and `temperature` (double).

Demo – Select A Job

The screenshot displays the Spline web interface. On the left, a sidebar shows a search bar and a list of jobs. The main area features a data pipeline diagram with nodes for Kafka, Hive, and Aggregation Job. A red arrow points to the 'Aggregation Job' node. On the right, a 'Job Details' panel shows the application ID, timestamp, and a list of sources.

Spline
Early Licensee: Testing for Apache Spark™

Last View: **Interest View**

Search: Sink Name

From: 2018-09-24 16:00

To: 2018-09-25 00:00

File://data/om/spline/sample/data/results/streamingWithDependencies/temperature.pregus.aggregate/AggregationJob

Kafka://localhost:9092/temperature.pregus.thety

Kafka://localhost:9092/temperature.pregus.kafkas

Kafka://localhost:9092/temperature.pregus.thety

Kafka://localhost:9092/temperature.pregus.thety

Kafka://localhost:9092/temperature.pregus.thety

Aggregation Job

Parquet/temperature.pregus.thety

AggregationJob

Application ID: local-1527797832018

Timestamp: Sep 24, 2018, 4:03:58 PM

Type: Stream

Show detailed lineage graph

Sources:

- Kafka - kafka://localhost:9092/temperature.pregus.thety (0ba2716f-2ea7-4509-a911-e04709115a27)
- Kafka - kafka://localhost:9092/temperature.pregus.kafkas (4aac21a3-8119-4046-b294-162393917793)
- Kafka - kafka://localhost:9092/temperature.pregus.thety (e02673d5-8d21-4acc-810e-a091bd46554e)

Demo – Drill Down

The image shows a Spline interface with a data pipeline diagram on the left and a detailed view of an 'Aggregation Job' on the right. A red arrow points from the 'Show detailed lineage graph' button in the job details to the pipeline diagram.

Spline Interface (Left):

- Header: Spline, Early License Testing for Apache Spark™
- Navigation: Last View, Interest View
- Search: Spline Name
- From: 2018-09-24 16:00
- To: 2018-09-25 00:00
- File path: file:///data/.../streamingWithDependencies/temperature.pregus.aggregate/
- Job details for 'Kafka://localhost:9092/temperature.pregus.thety' and 'Kafka://localhost:9092/temperature.pregus.kafka'.

Data Pipeline Diagram (Center):

- Inputs: 'Kafka://localhost:9092/temperature.pregus.thety' and 'file:///data/.../streamingWithDependencies/temperature.pregus.thety'.
- Processors: 'KafkaMessageDistributor' and 'FileMessageDistributor'.
- Intermediate: Three 'Kafka://localhost:9092/temperature.pregus.thety' jobs.
- Output: 'Aggregation Job'.
- Final Output: 'Parquet/temperature.pregus.thety'.

Aggregation Job Details (Right):

- Application ID: local-1537797832018
- Timestamp: Sep 24, 2018, 4:03:58 PM
- Type: Stream
- Button: Show detailed lineage graph (indicated by a red arrow)
- Sources:
 - Kafka - kafka://localhost:9092/temperature.pregus.thety (file:///data/.../streamingWithDependencies/temperature.pregus.thety)
 - Kafka - kafka://localhost:9092/temperature.pregus.kafka (file:///data/.../streamingWithDependencies/temperature.pregus.kafka)
 - Kafka - kafka://localhost:9092/temperature.pregus.thety (file:///data/.../streamingWithDependencies/temperature.pregus.thety)

Demo – Review Job Details

The screenshot displays the Spline web interface for reviewing a job. The top left shows the Spline logo and the text "Early Licensee: Testing for Apache Spark™". Below this, there are tabs for "Last Wife" and "Interest View". A search bar labeled "Sink Name" is present. The main content area lists job details:

- From:** 2018-09-24 16:00
- To:** 2018-09-25 00:00
- Job Name:** `file:///data/om/spline/sample/data/results/streamingWithDependencies/temperature.pregus.aggregate/Aggregate`
- Job ID:** `file:///data/om/spline/sample/data/results/streamingWithDependencies/temperature.pregus.kbety`
- Job ID:** `file:///data/om/spline/sample/data/results/streamingWithDependencies/temperature.pregus.kafka`

The central part of the interface shows a vertical pipeline diagram with various nodes connected by arrows. The right-hand panel is titled "Summary" and "Aggregate", showing configuration options for "Input window" and "Output window". A red arrow points to the "Aggregate" section in the right panel.

Demo – Select An Operation

The screenshot displays the Spline interface for a data pipeline. On the left, a sidebar shows the pipeline details, including the sink name, start and end times, and the path to the aggregate operation. The main area shows a vertical flow of operations: a source, several filters, a projection, an aggregate operation, and a sink. A red arrow points to the 'Aggregate' operation in the flow. The right sidebar is open to the configuration of the 'Aggregate' operation, showing 'Group by' and 'Aggregations' sections.

Spline
Early Licensee Tracking for Apache Spark™

Last Write Interval View

Link Name

From: 2018-09-24 16:00

To: 2018-09-25 00:00

`s3://data-us-east-1-samples/sample/data/results/streamingWSSDependencies/temperature_pregus_aggregate/`

`kafka://localhost:9092/temperature_pregus_kbety`

`kafka://localhost:9092/temperature_pregus_kafwa`


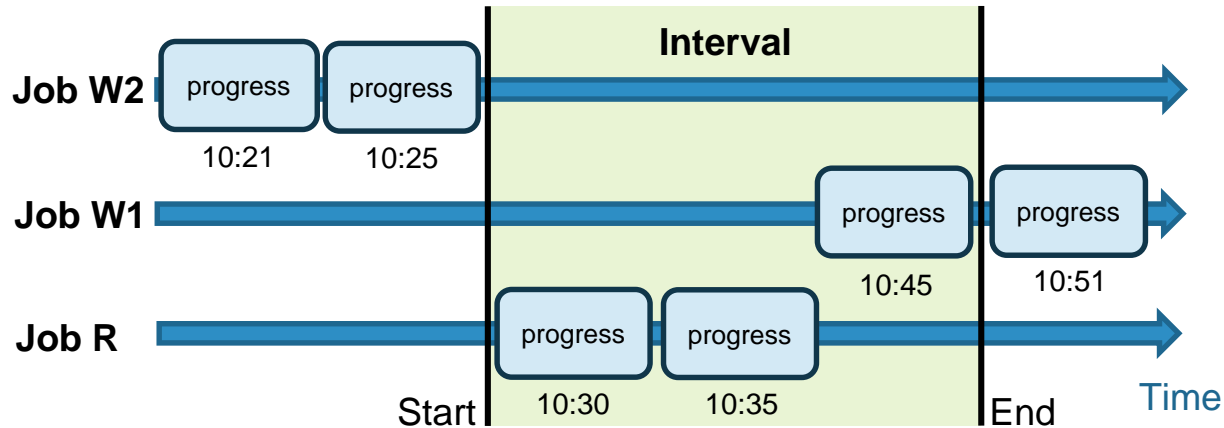
Projection Aggregate Filter

Aggregate

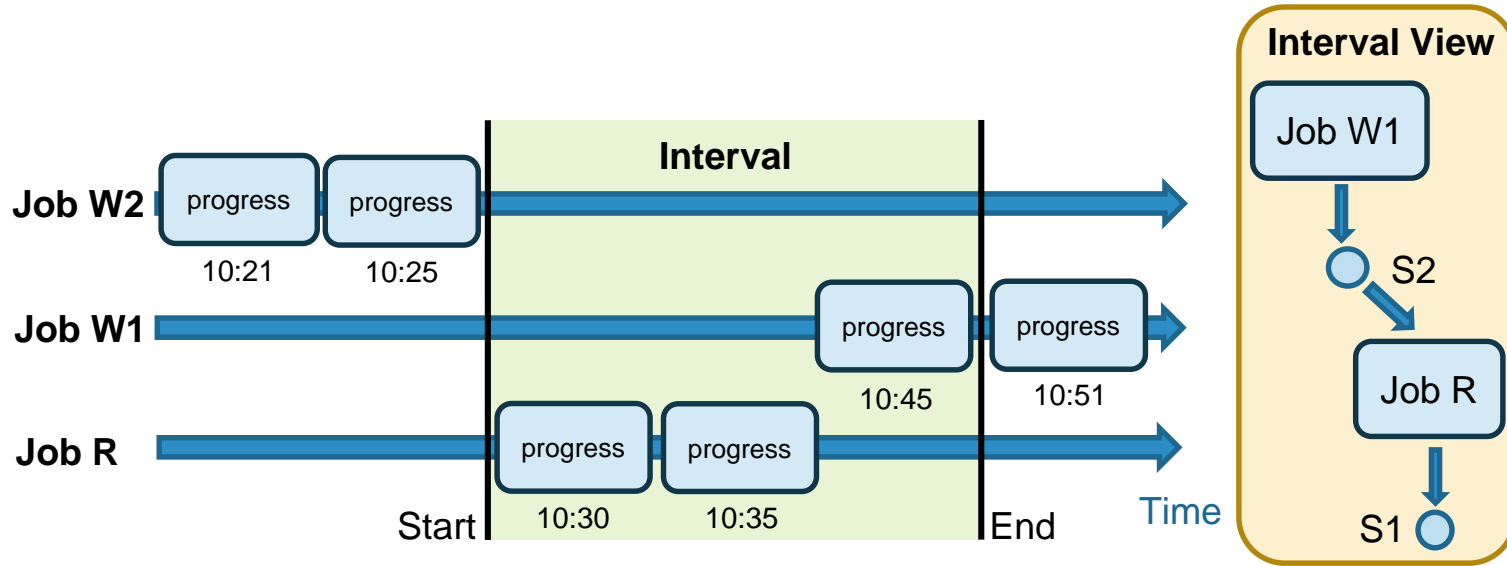
- Group by
 - window
- Aggregations
 - `count(*) OVER (PARTITION BY window)`
 - `sum(temperature) OVER (PARTITION BY window)`
- Input
 - window
 - time
 - temperature
- Output
 - window
 - temperature

Spline v1.2.3 (2018/09/24)

Interval View Limitations

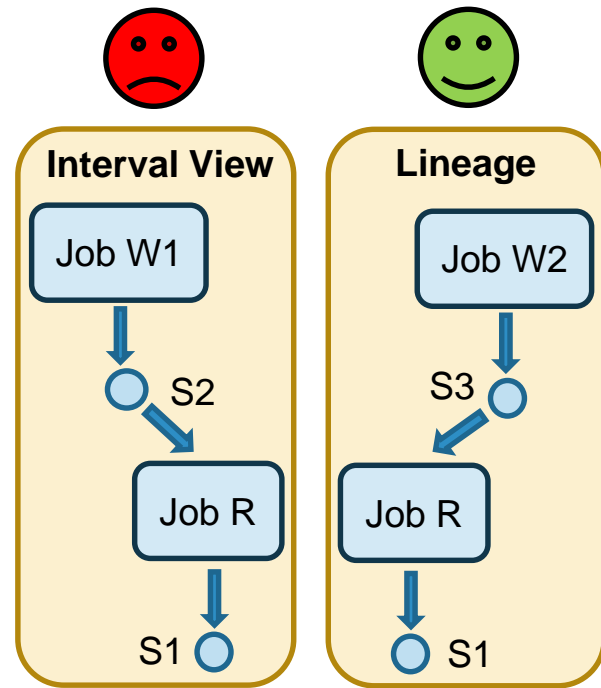
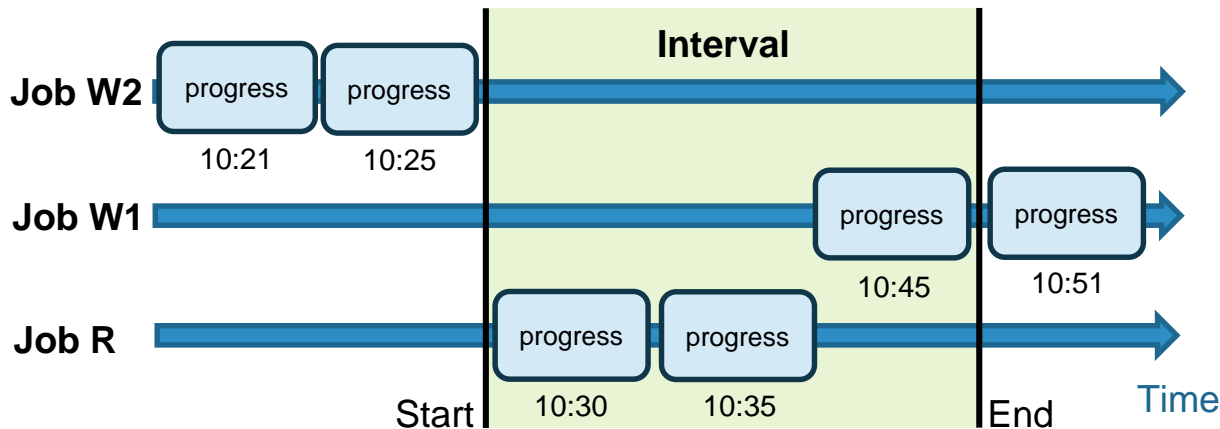


Interval View Limitations



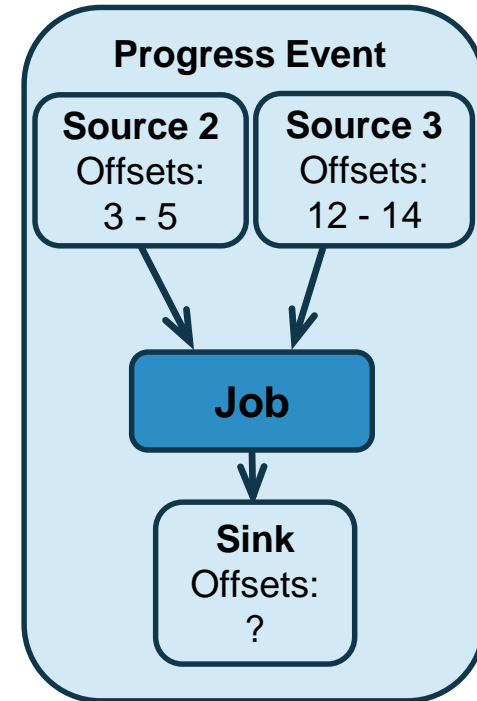
Interval View Limitations

- Edge case (delayed read, early write)
 - **Job W1** should be linked
 - **Job W2** should not be linked





Beyond The Interval View

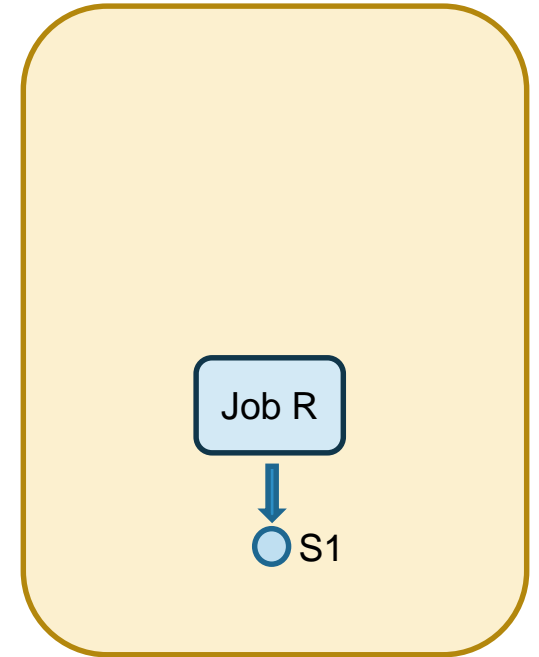

- Instead of timestamp use **addresses** of rows
- SS has addresses (**offsets**) on each source, but not on sinks
- Most sinks are also sources and thus *could* return **offsets**



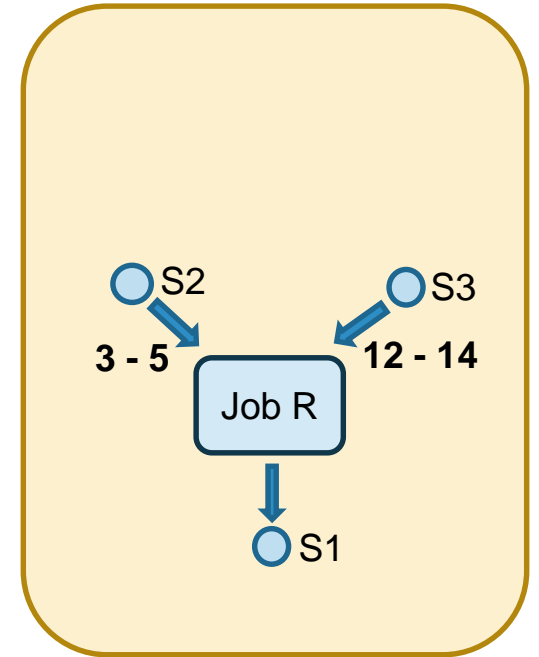
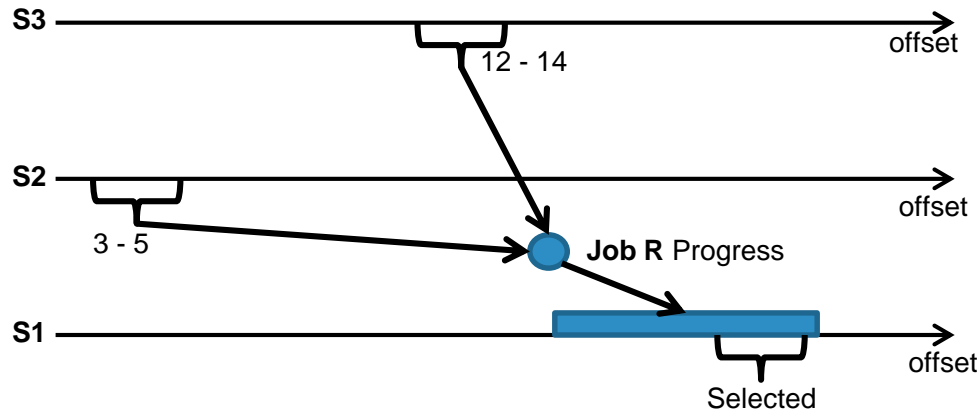
Offset-Based Linking




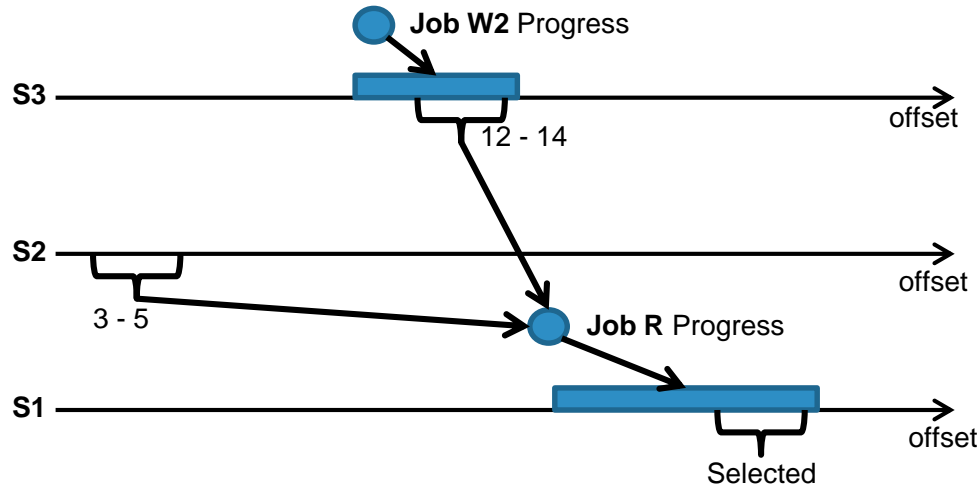
Offset-Based Linking



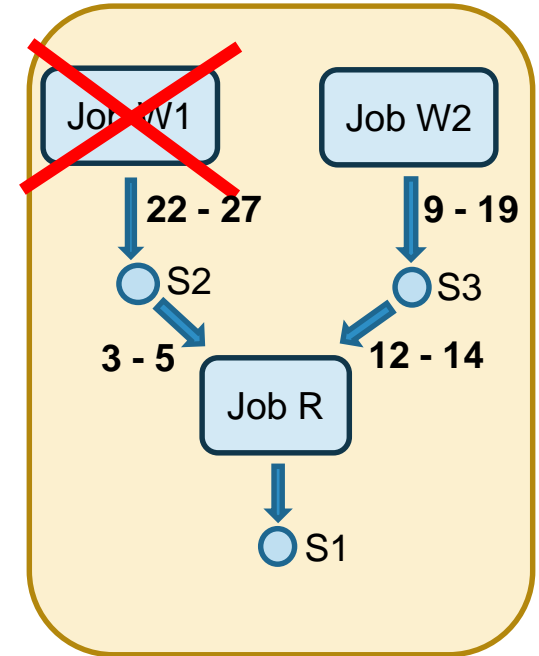

Offset-Based Linking



Offset-Based Linking





Offset-Based Linking



Offset-Based Linking

- Jobs are linked when progress offsets overlap
- Offset timestamp doesn't matter



Conclusion

- Spline: data lineage tracking tool
- New support for Structured Streaming
- Demo POC: Interval View
- Proposed generalization: offset-based linking

Future Plans

- Release Interval View in Spline
- After changes to Spark:
 - Offset based linking for micro-batch streaming
 - Continuous streaming support
- Support for dataset checkpoints

Questions

- Now is a good time
- Or feel free to contact us
 - Marek Novotny
 - mn.mikke@gmail.com
 - Vaclav Kosar
 - admin@vaclavkosar.com
- github.com/AbsaOSS/spline

