

AquaLang: A Streaming Dataflow Programming Language

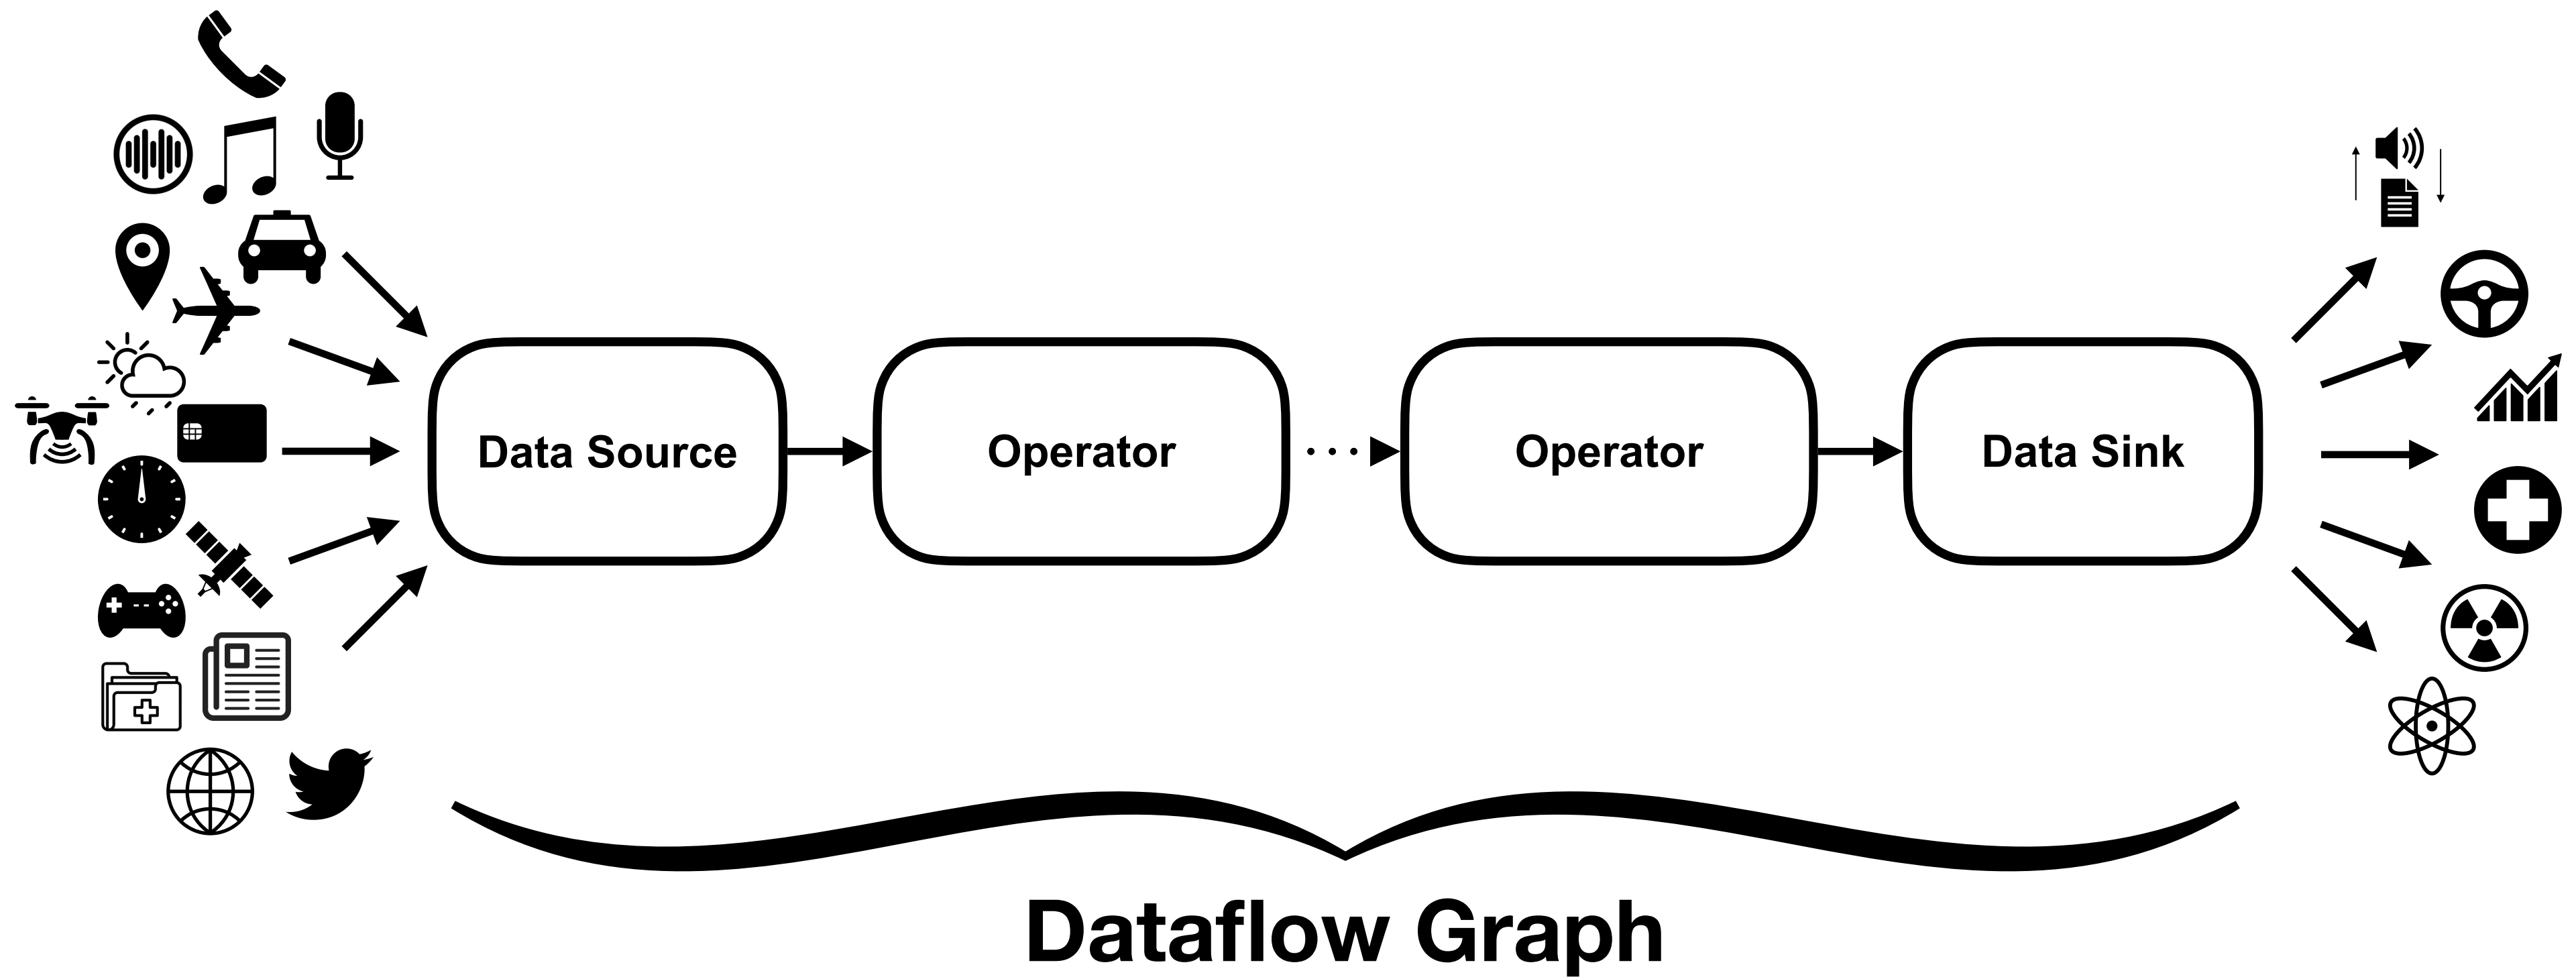
Klas Segeljakt^{1,2} (klasseg@kth)
Seif Haridi^{1,2} (haridi@kth.se)
Paris Carbone^{1,2} (parisc@kth.se)

¹ **Data Systems Lab** at KTH Royal Institute of Technology

² RISE Research Institutes of Sweden

HYTRADBOI 25

What is Dataflow Programming?




The Prospect of Dataflow Programming

The Prospect of Dataflow Programming



The Prospect of Dataflow Programming

User



Data Analytics


"What"

Complexity

"How"

The Prospect of Dataflow Programming

User



Data Analytics

"What"

Fault tolerance

Consistency Scalability

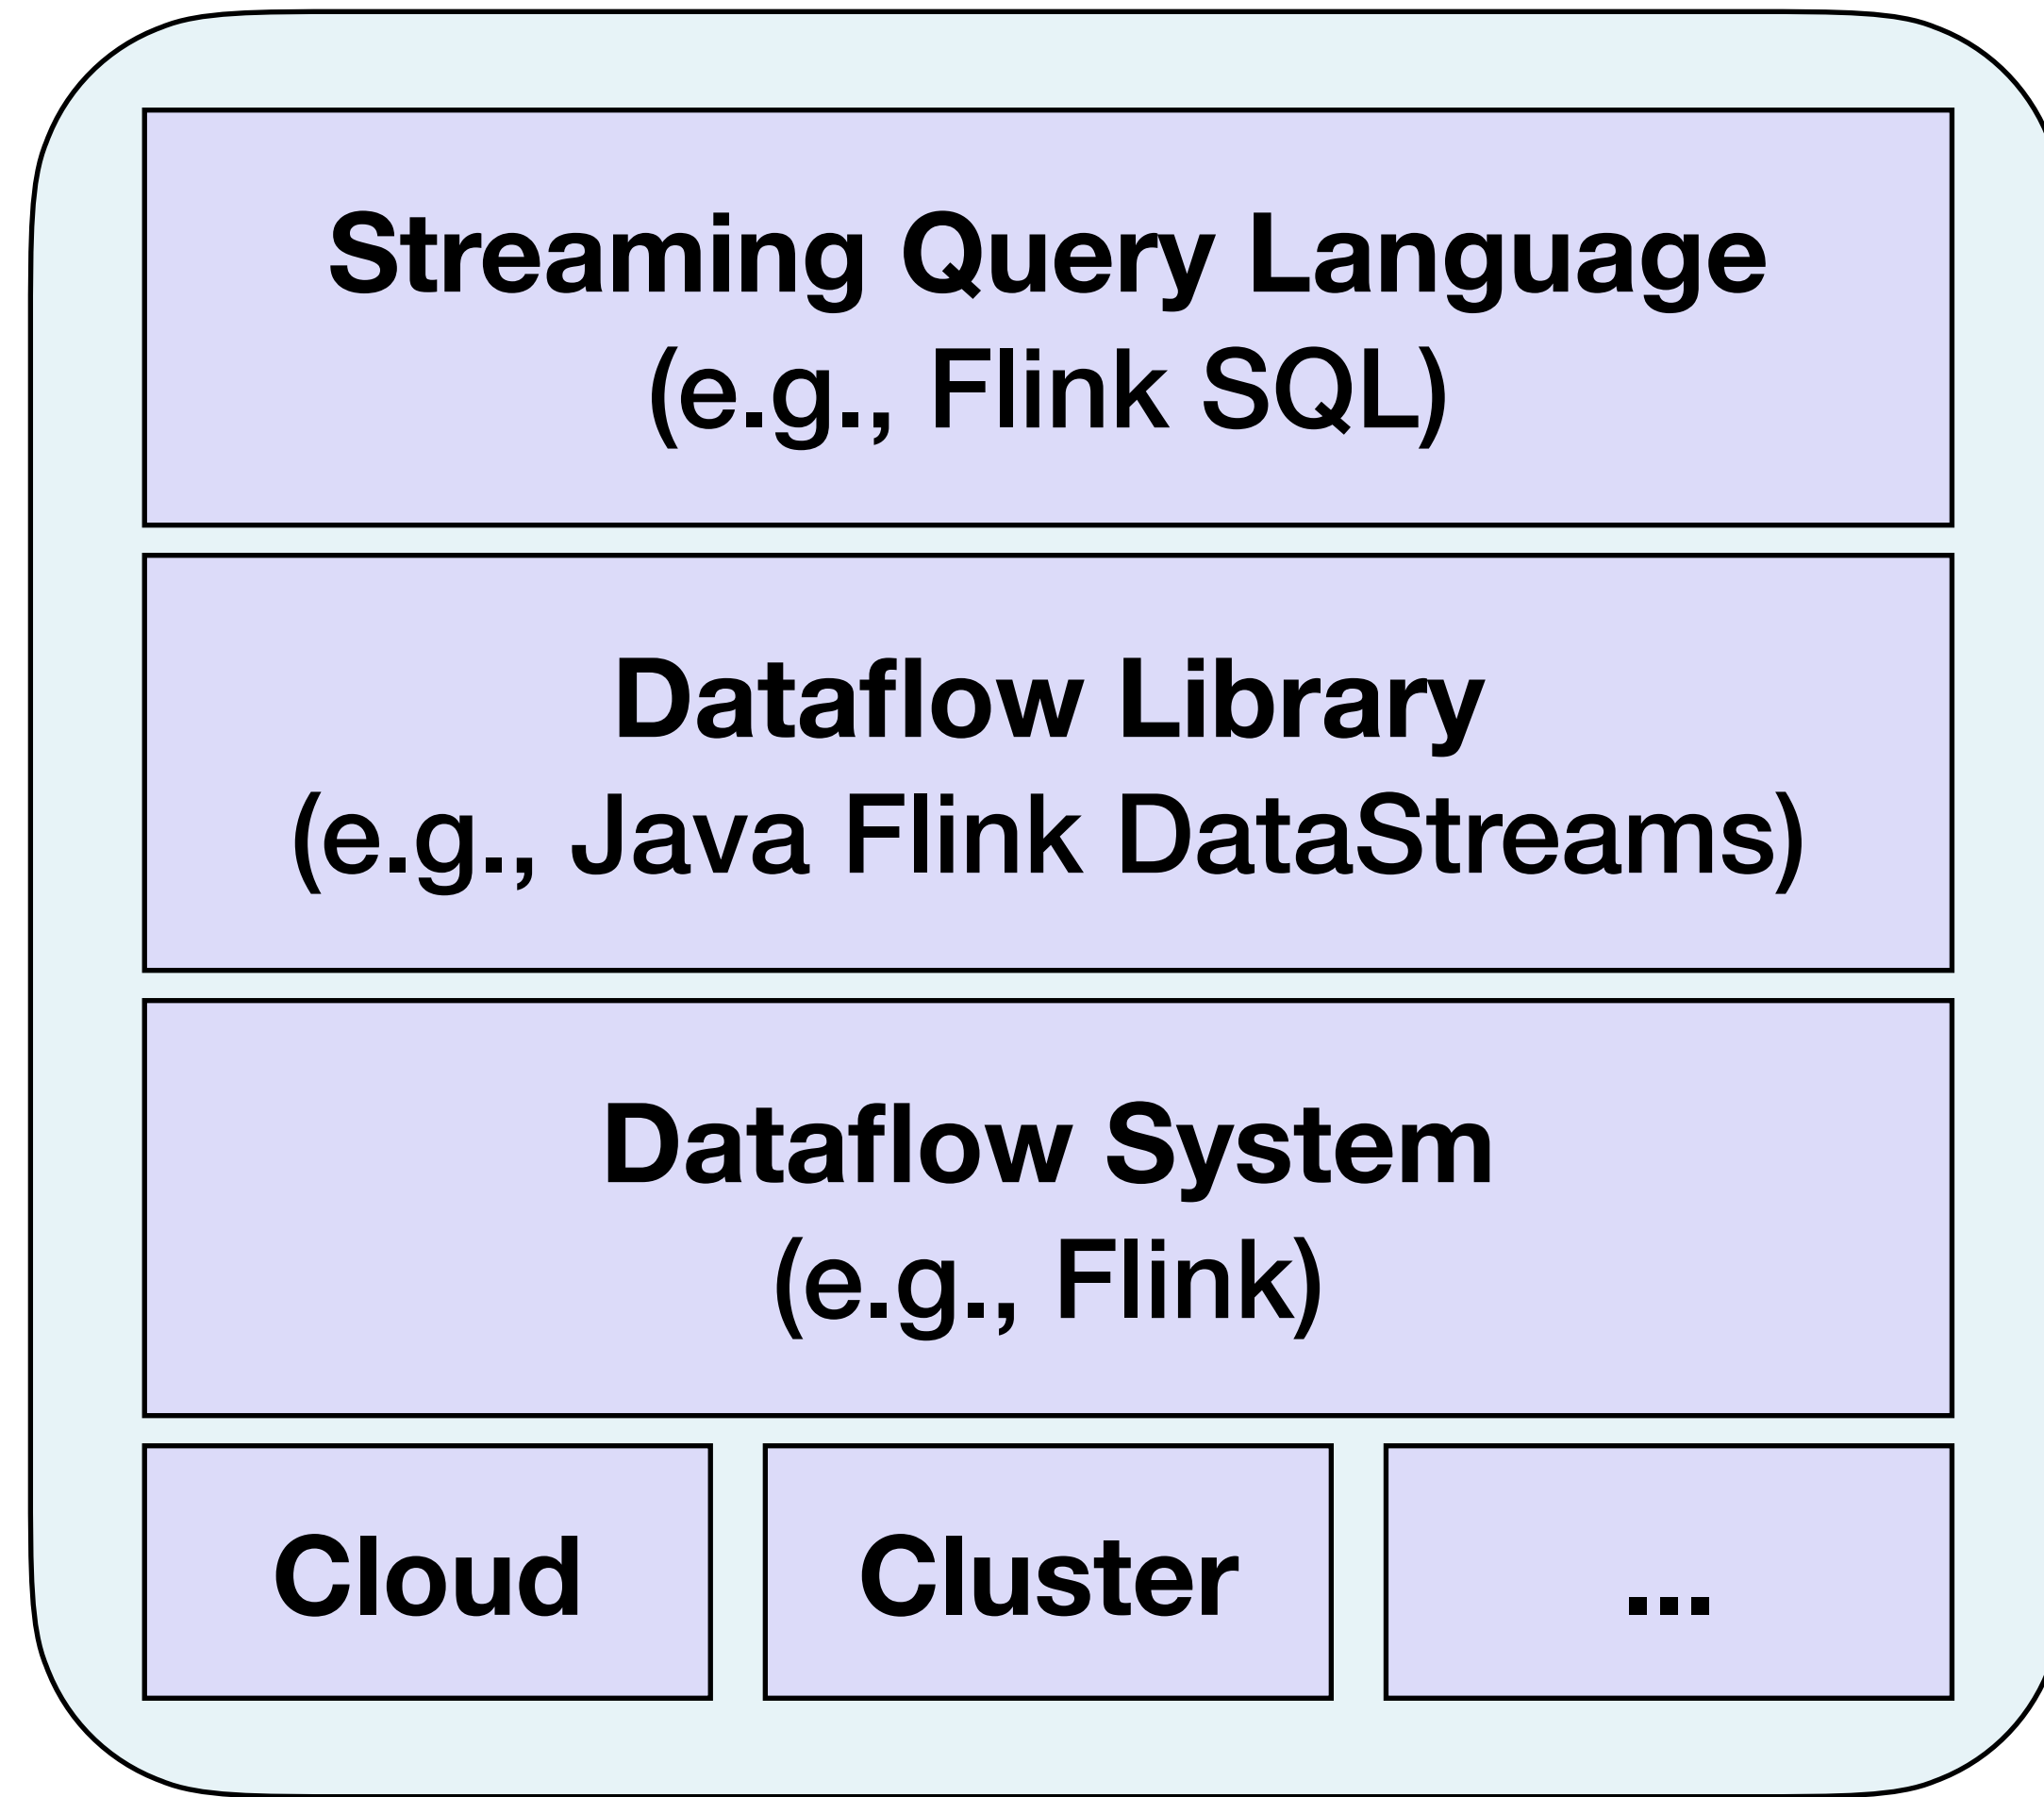
Complexity

Parallelism Serialisation


...

"How"

What exists today



What exists today

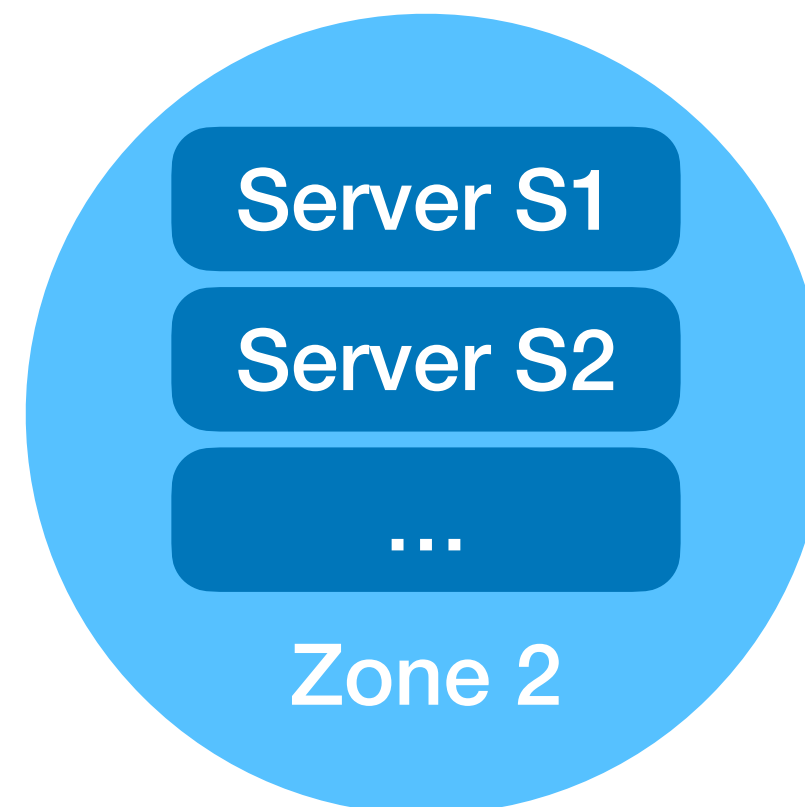




Why AquaLang?

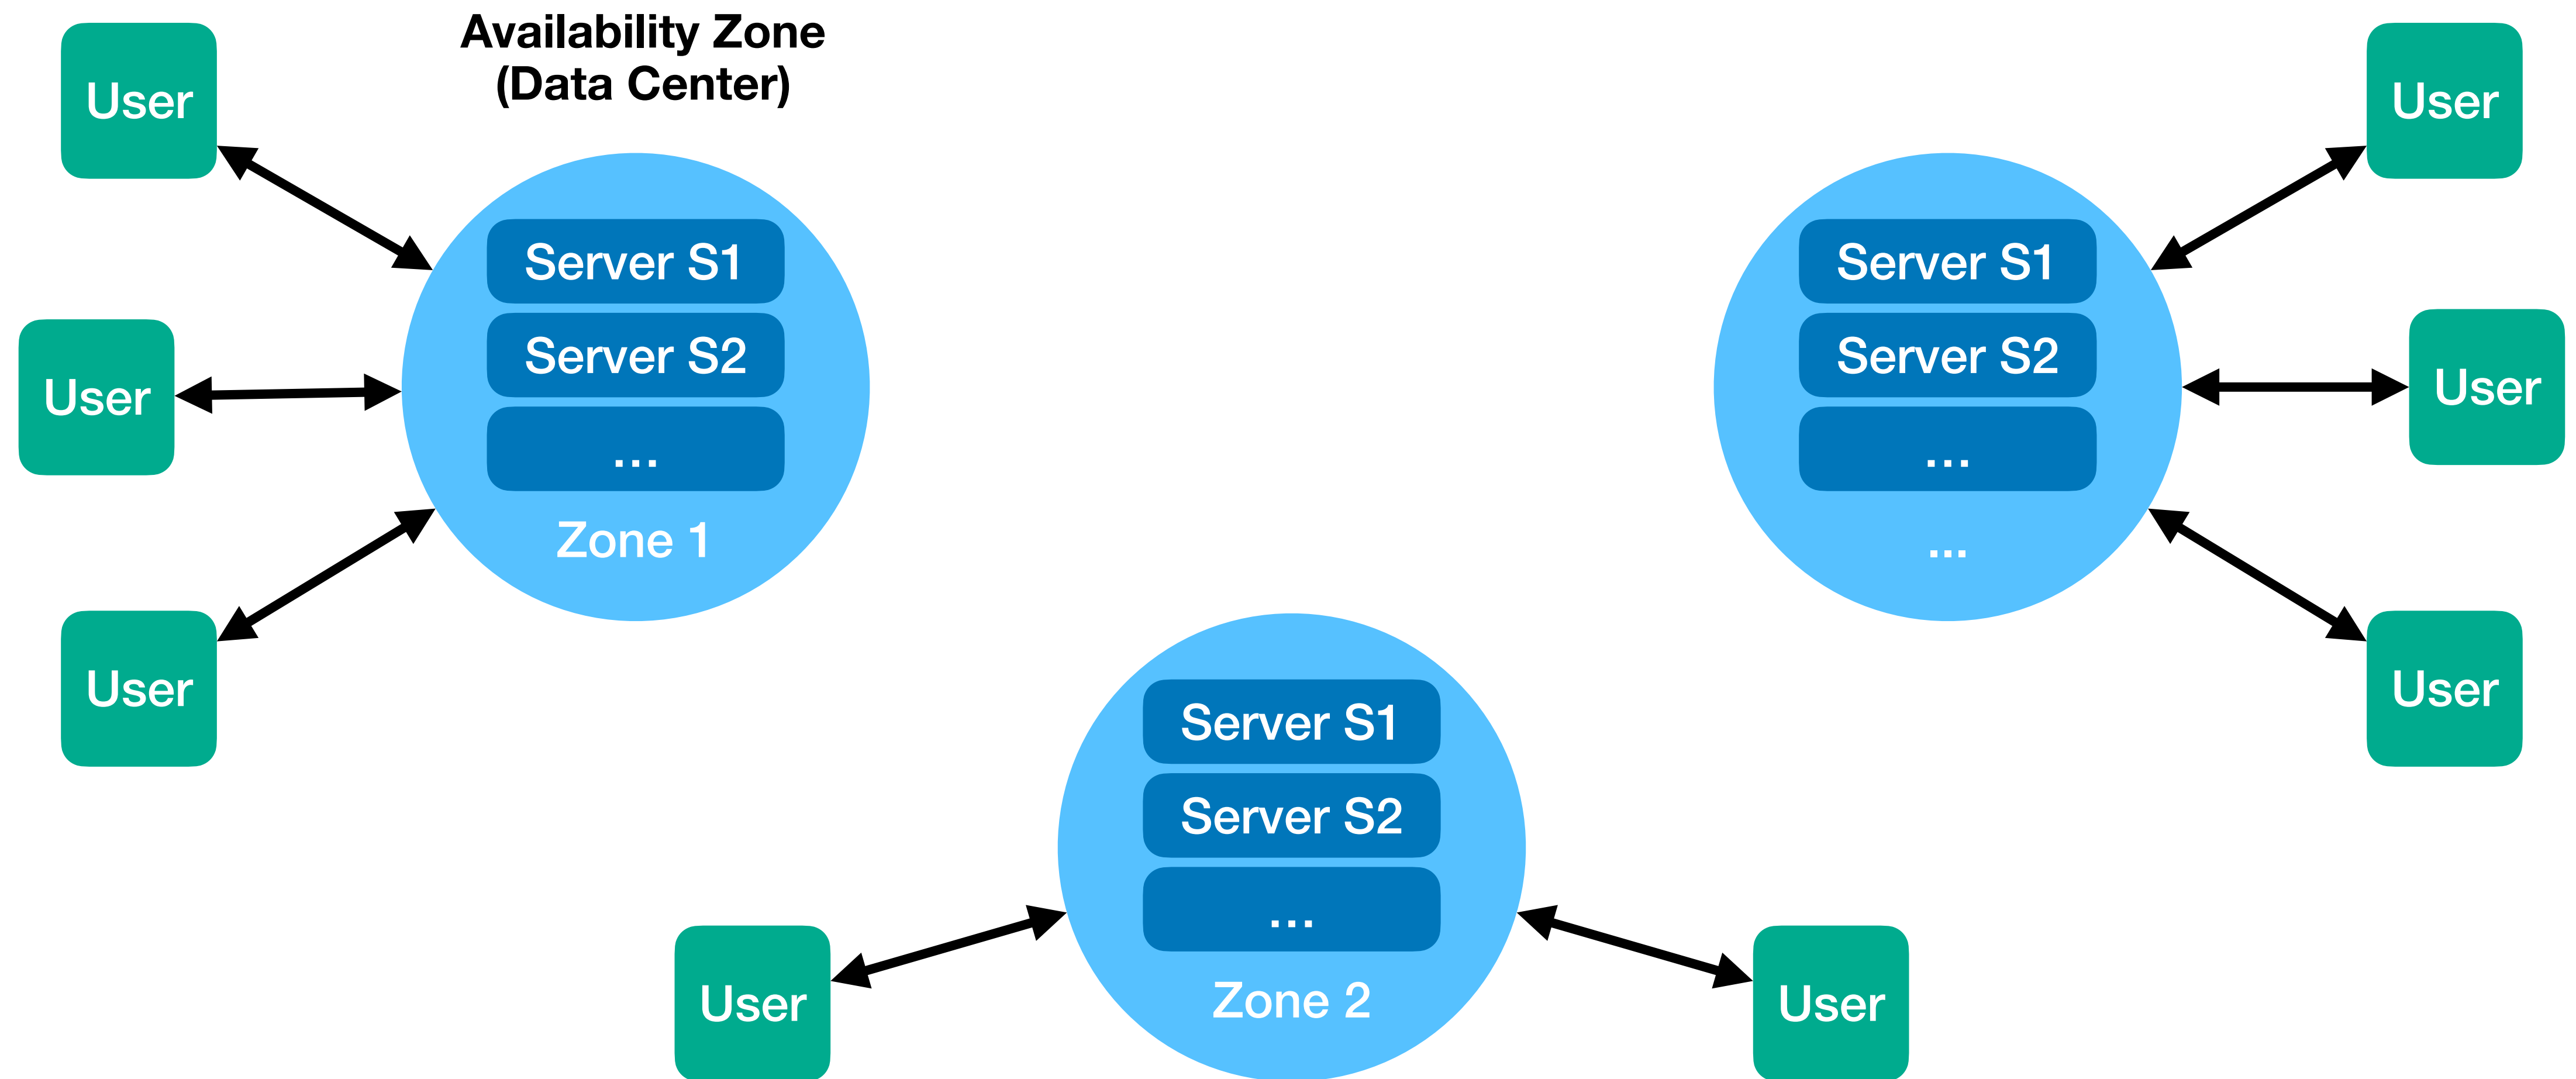
Example Application

Example Application

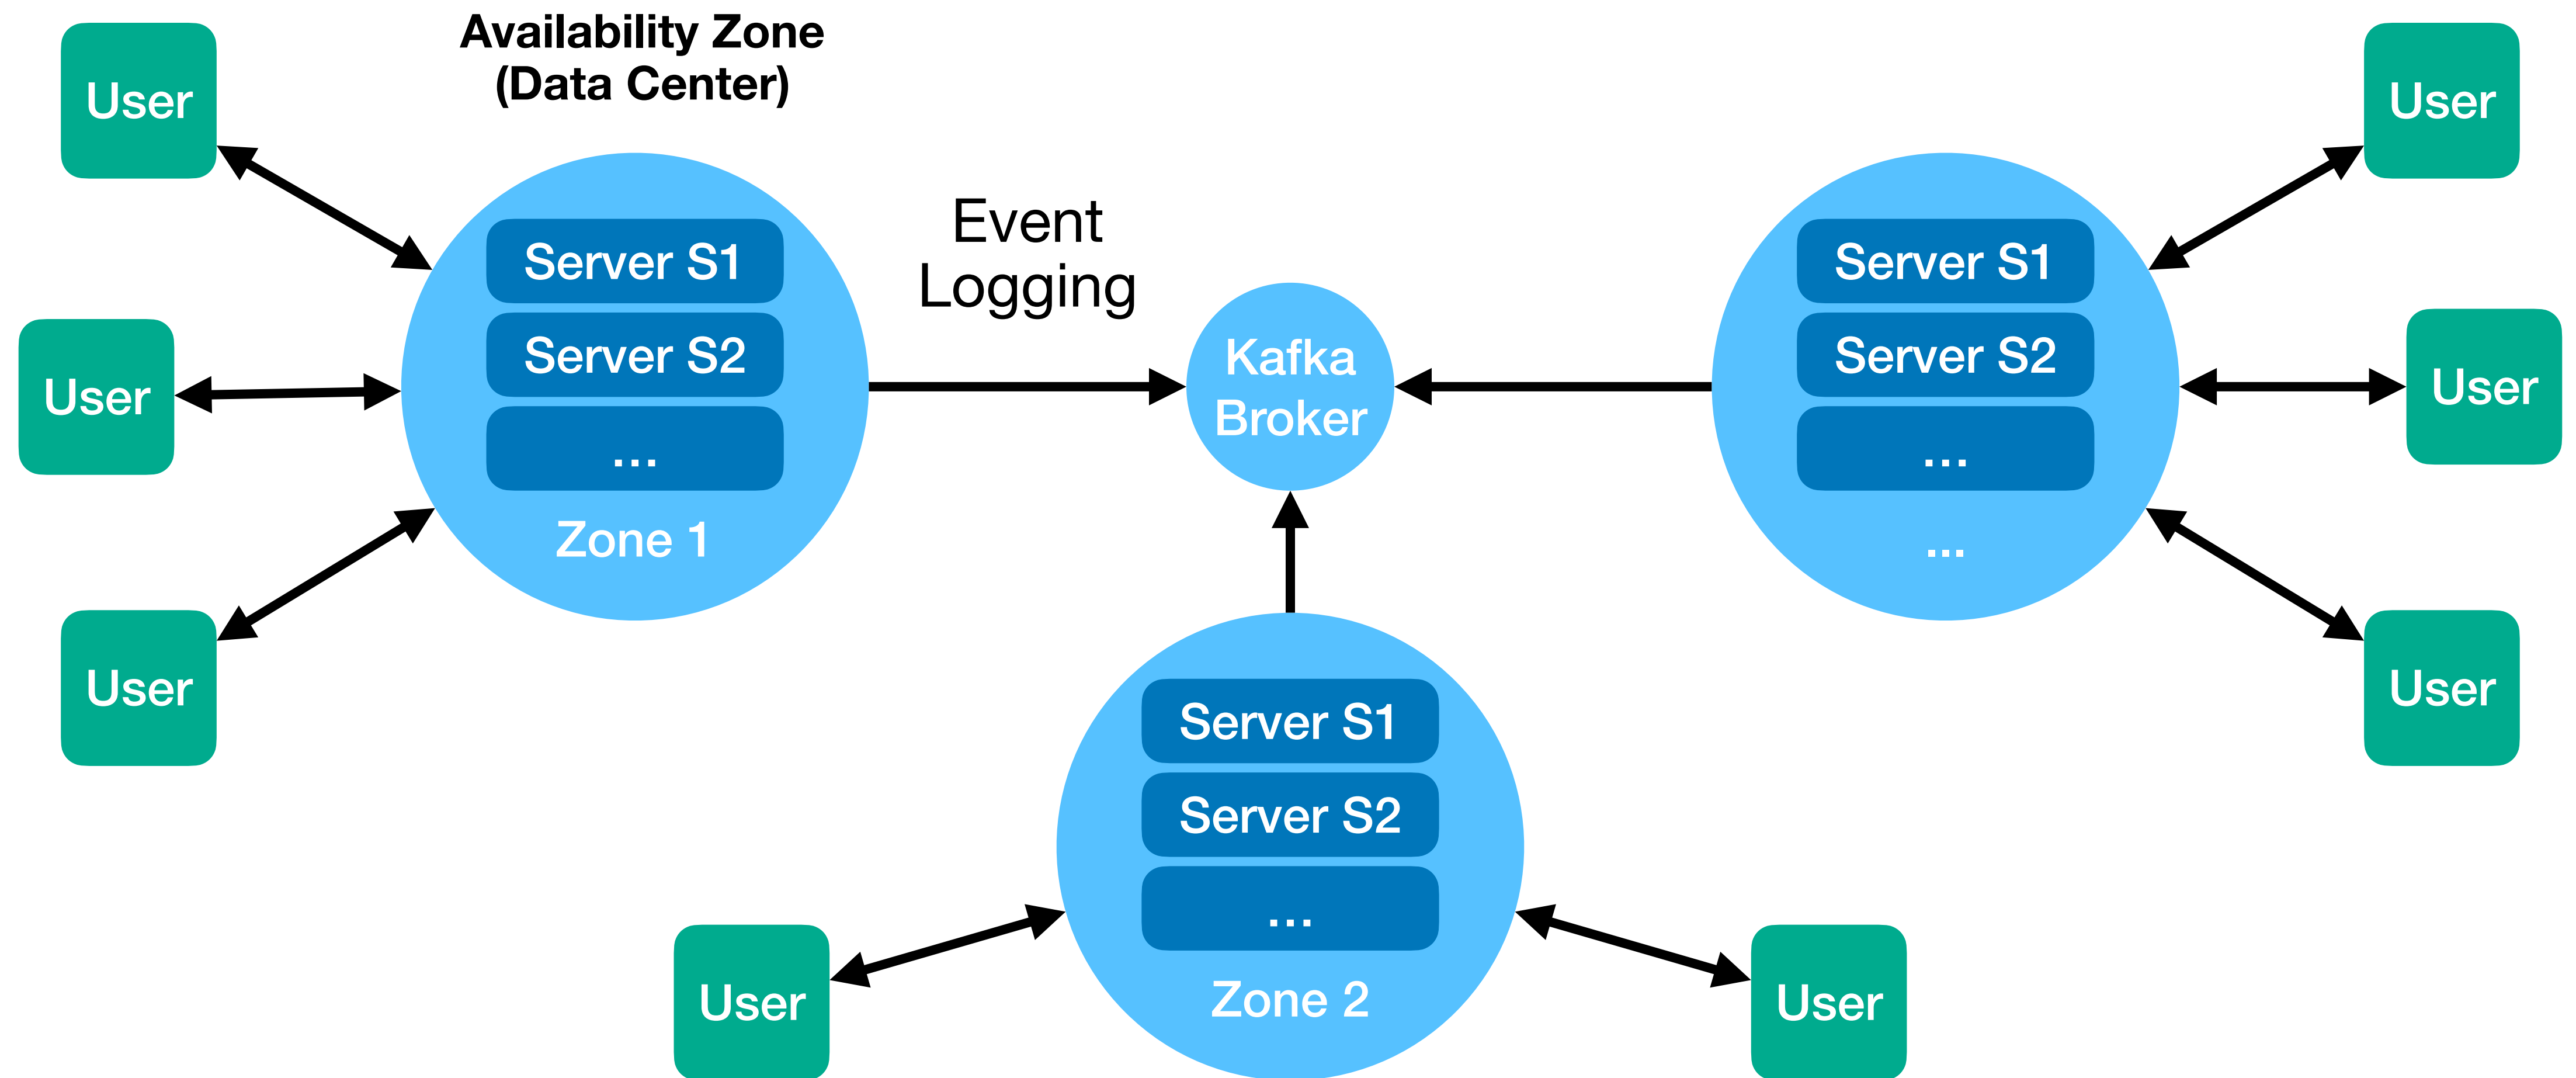
Availability Zone
(Data Center)



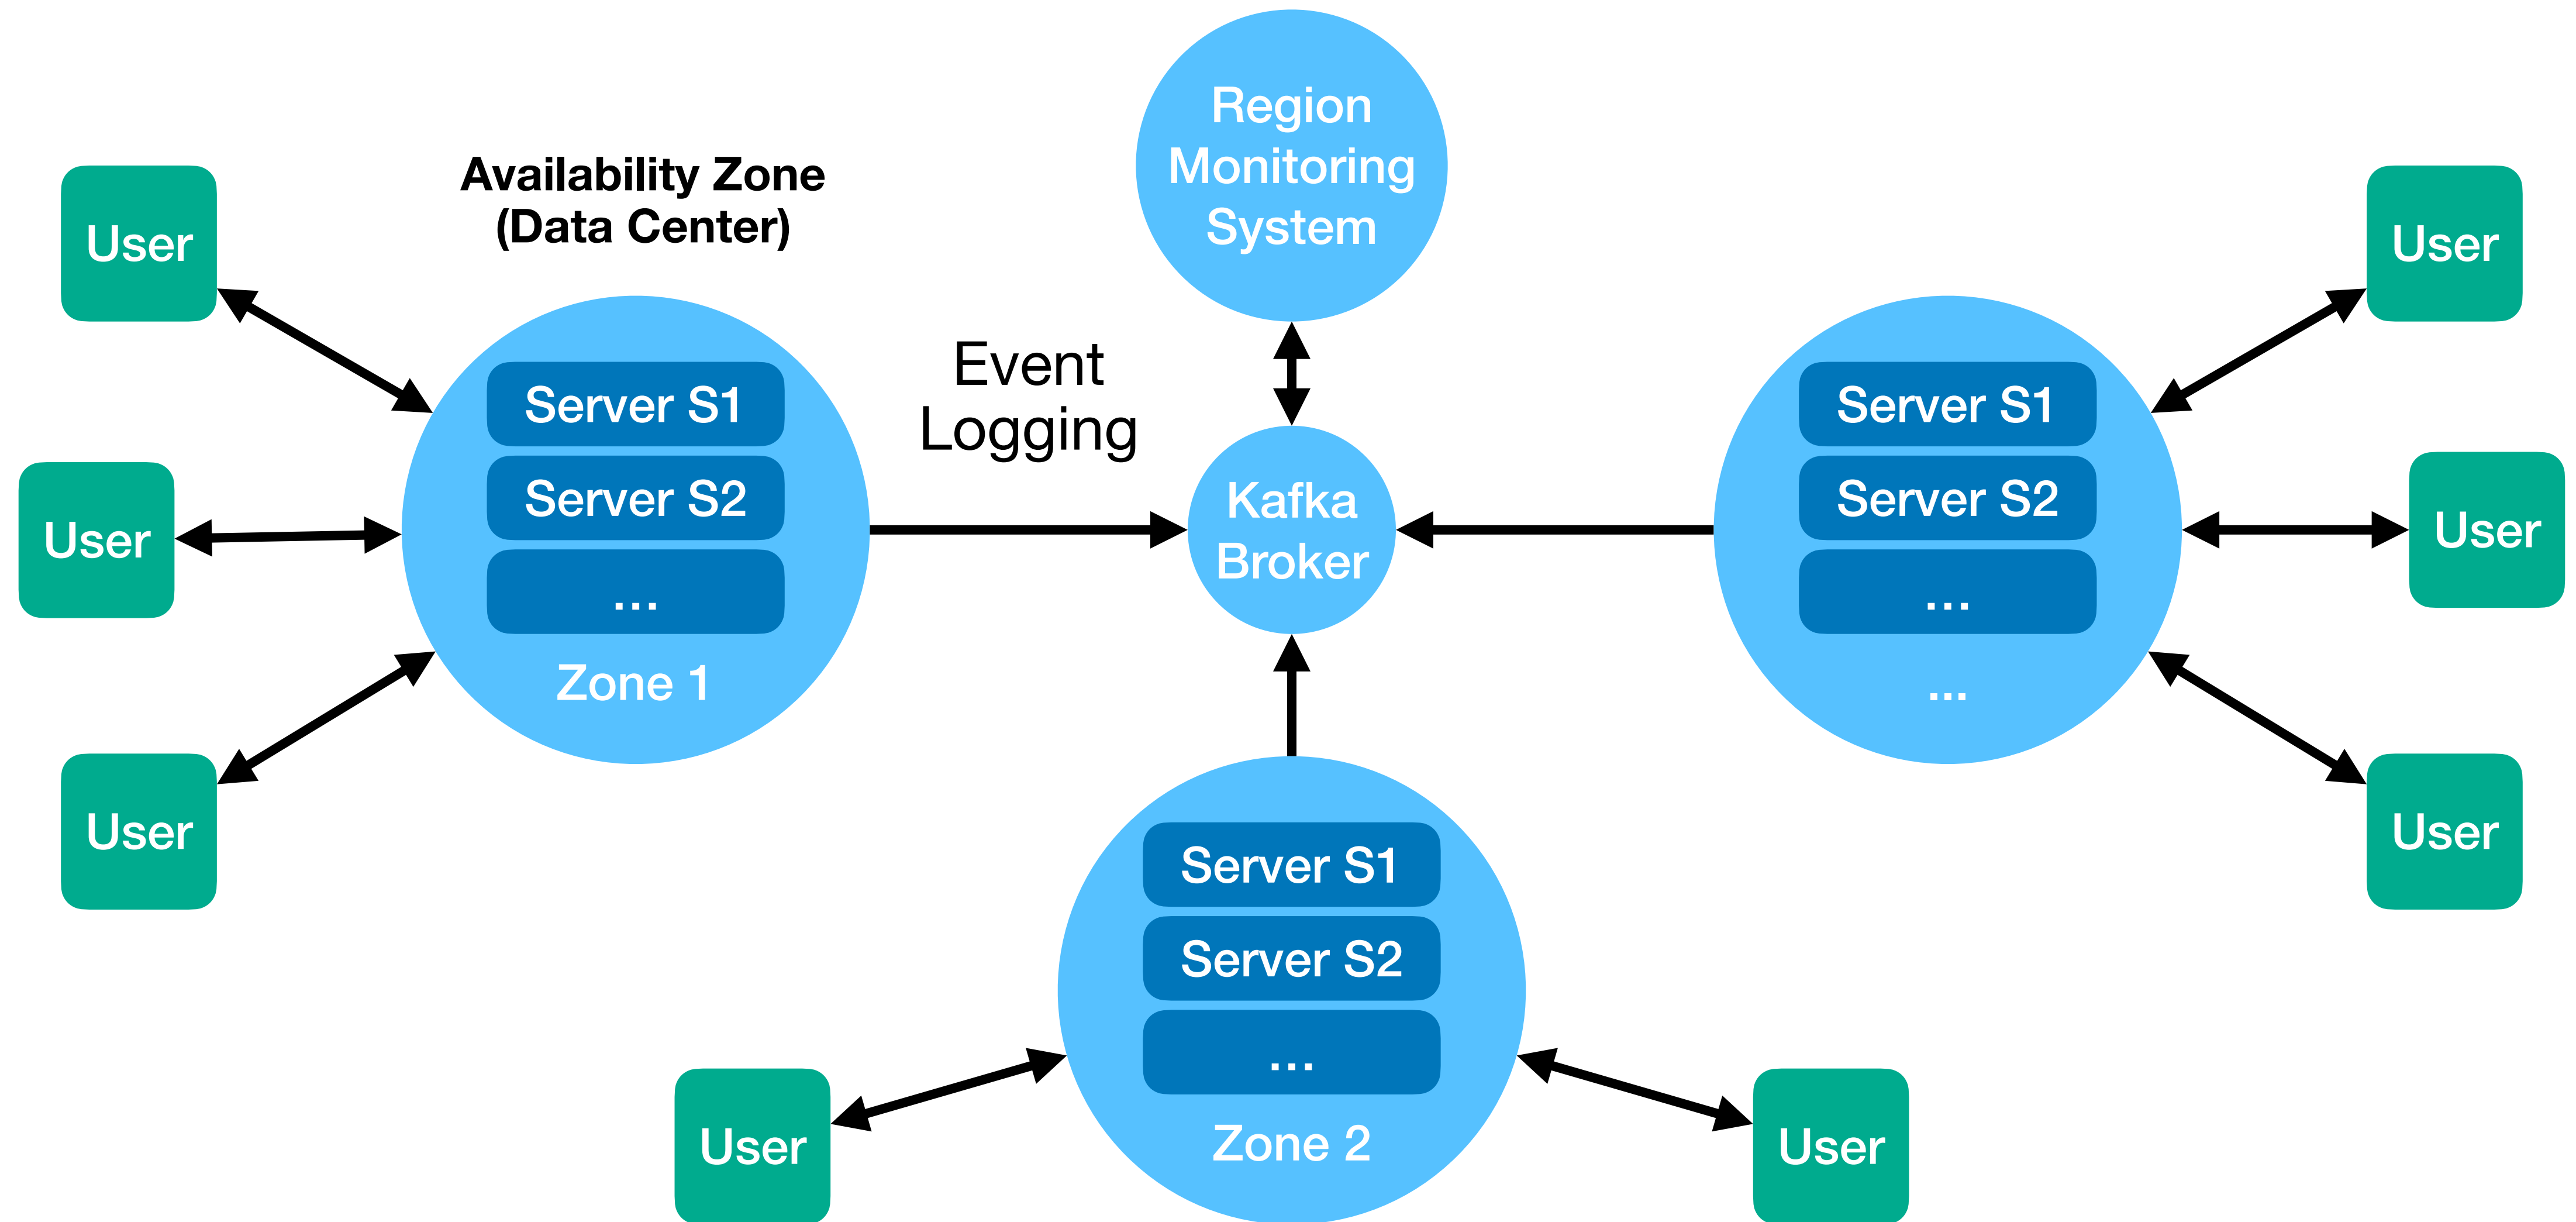
Example Application



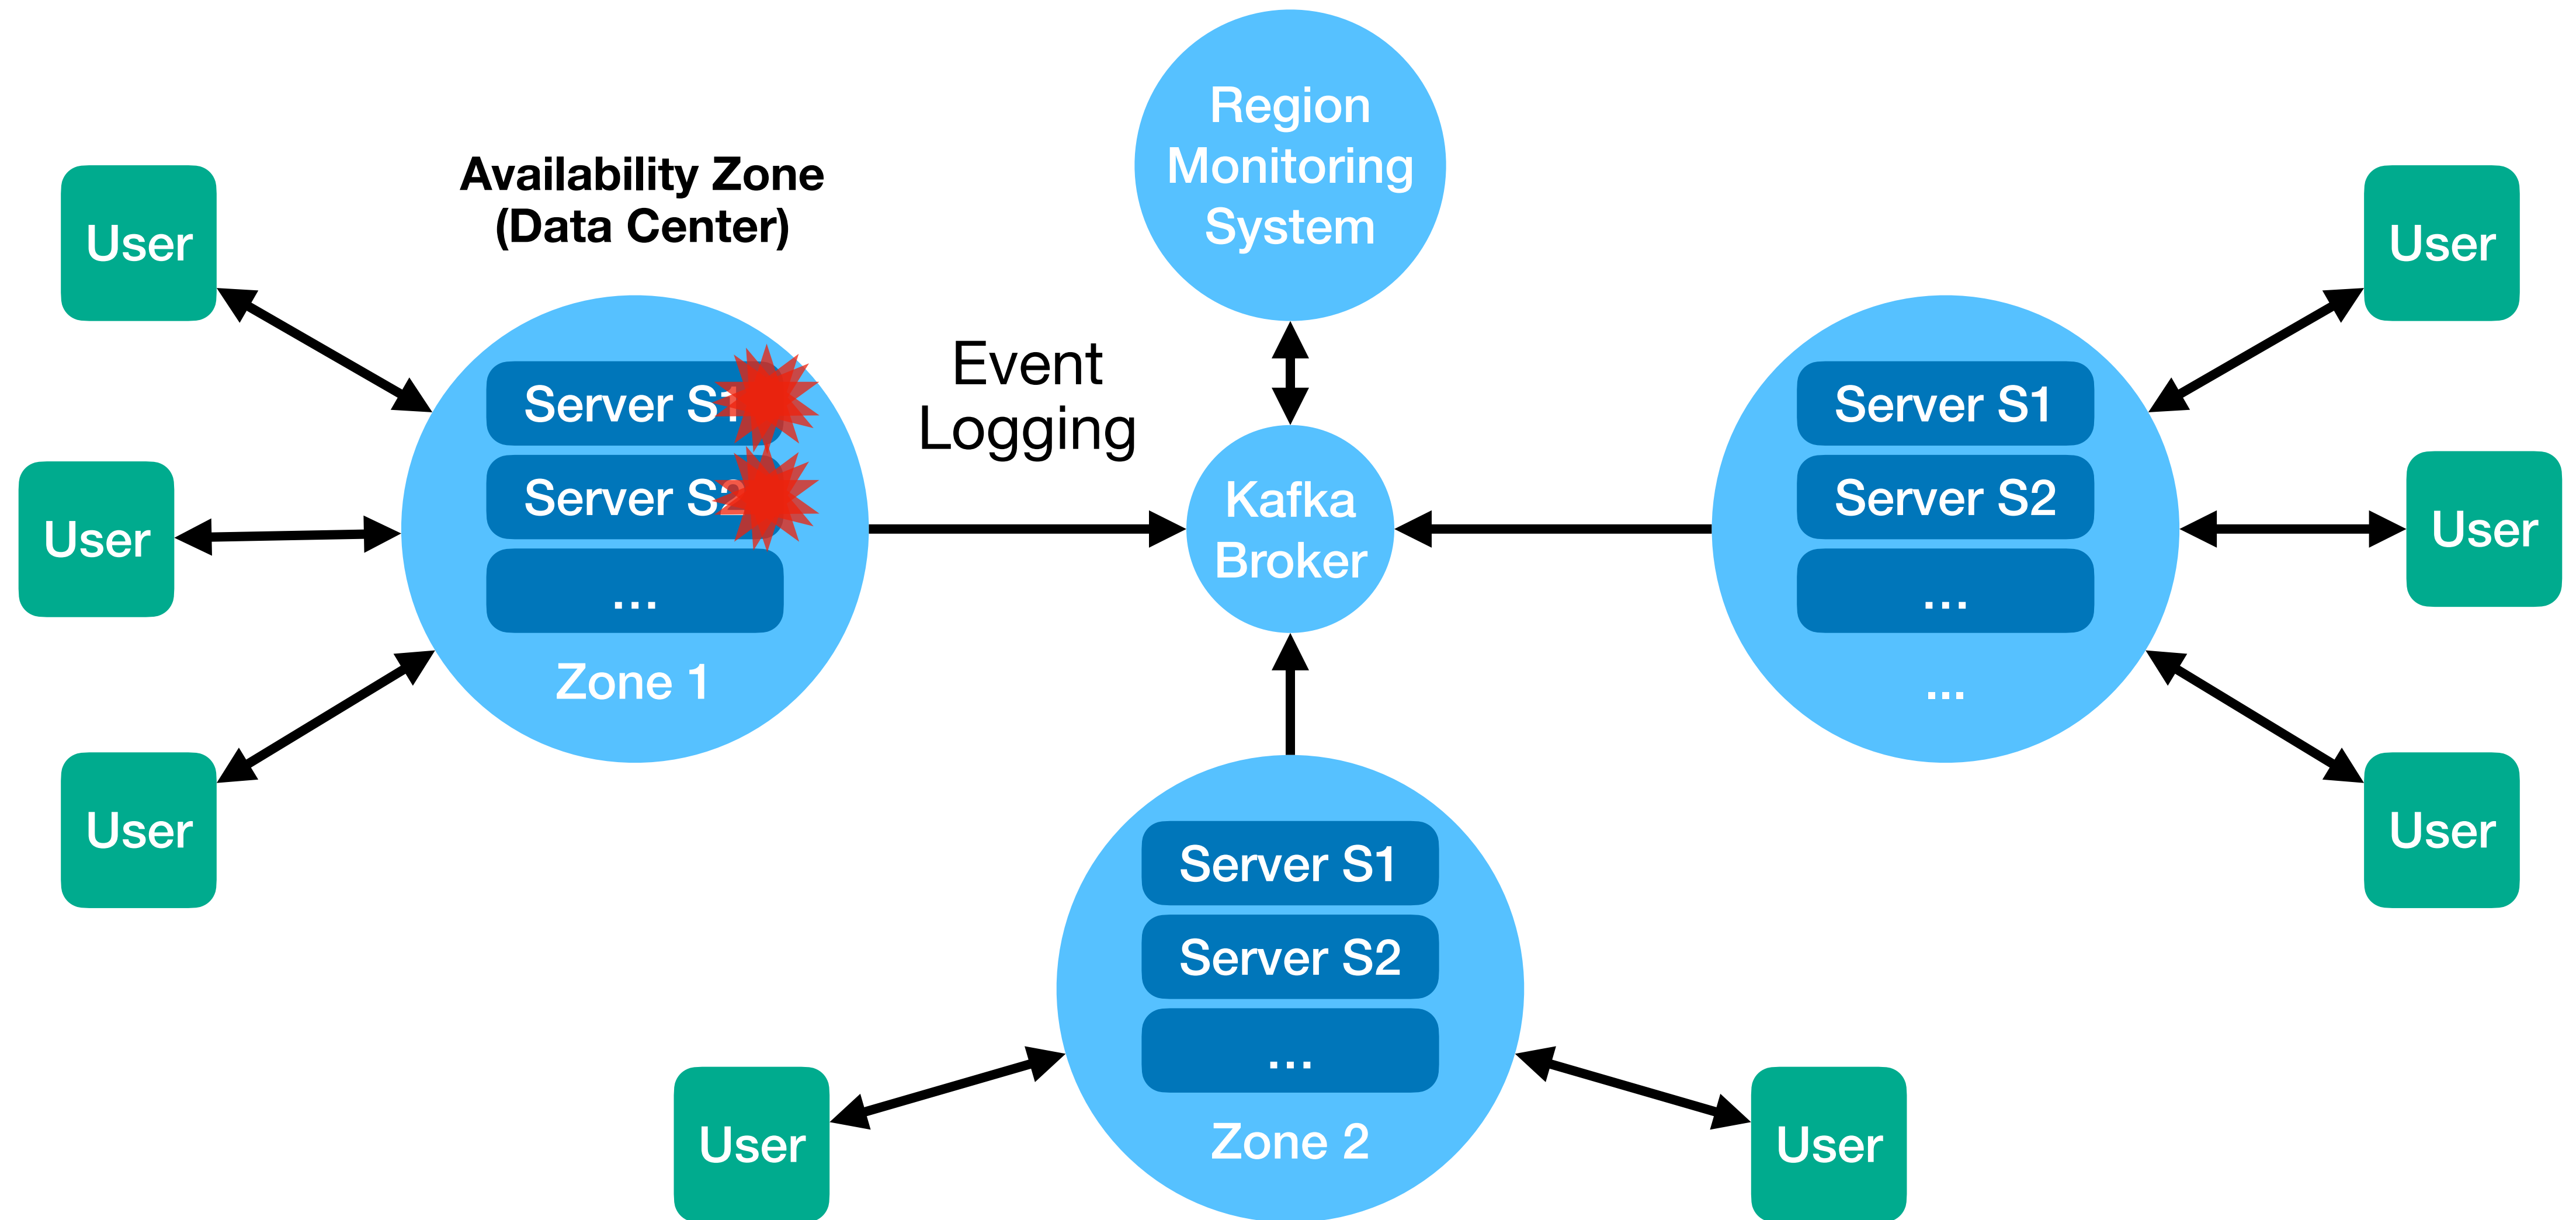
Example Application



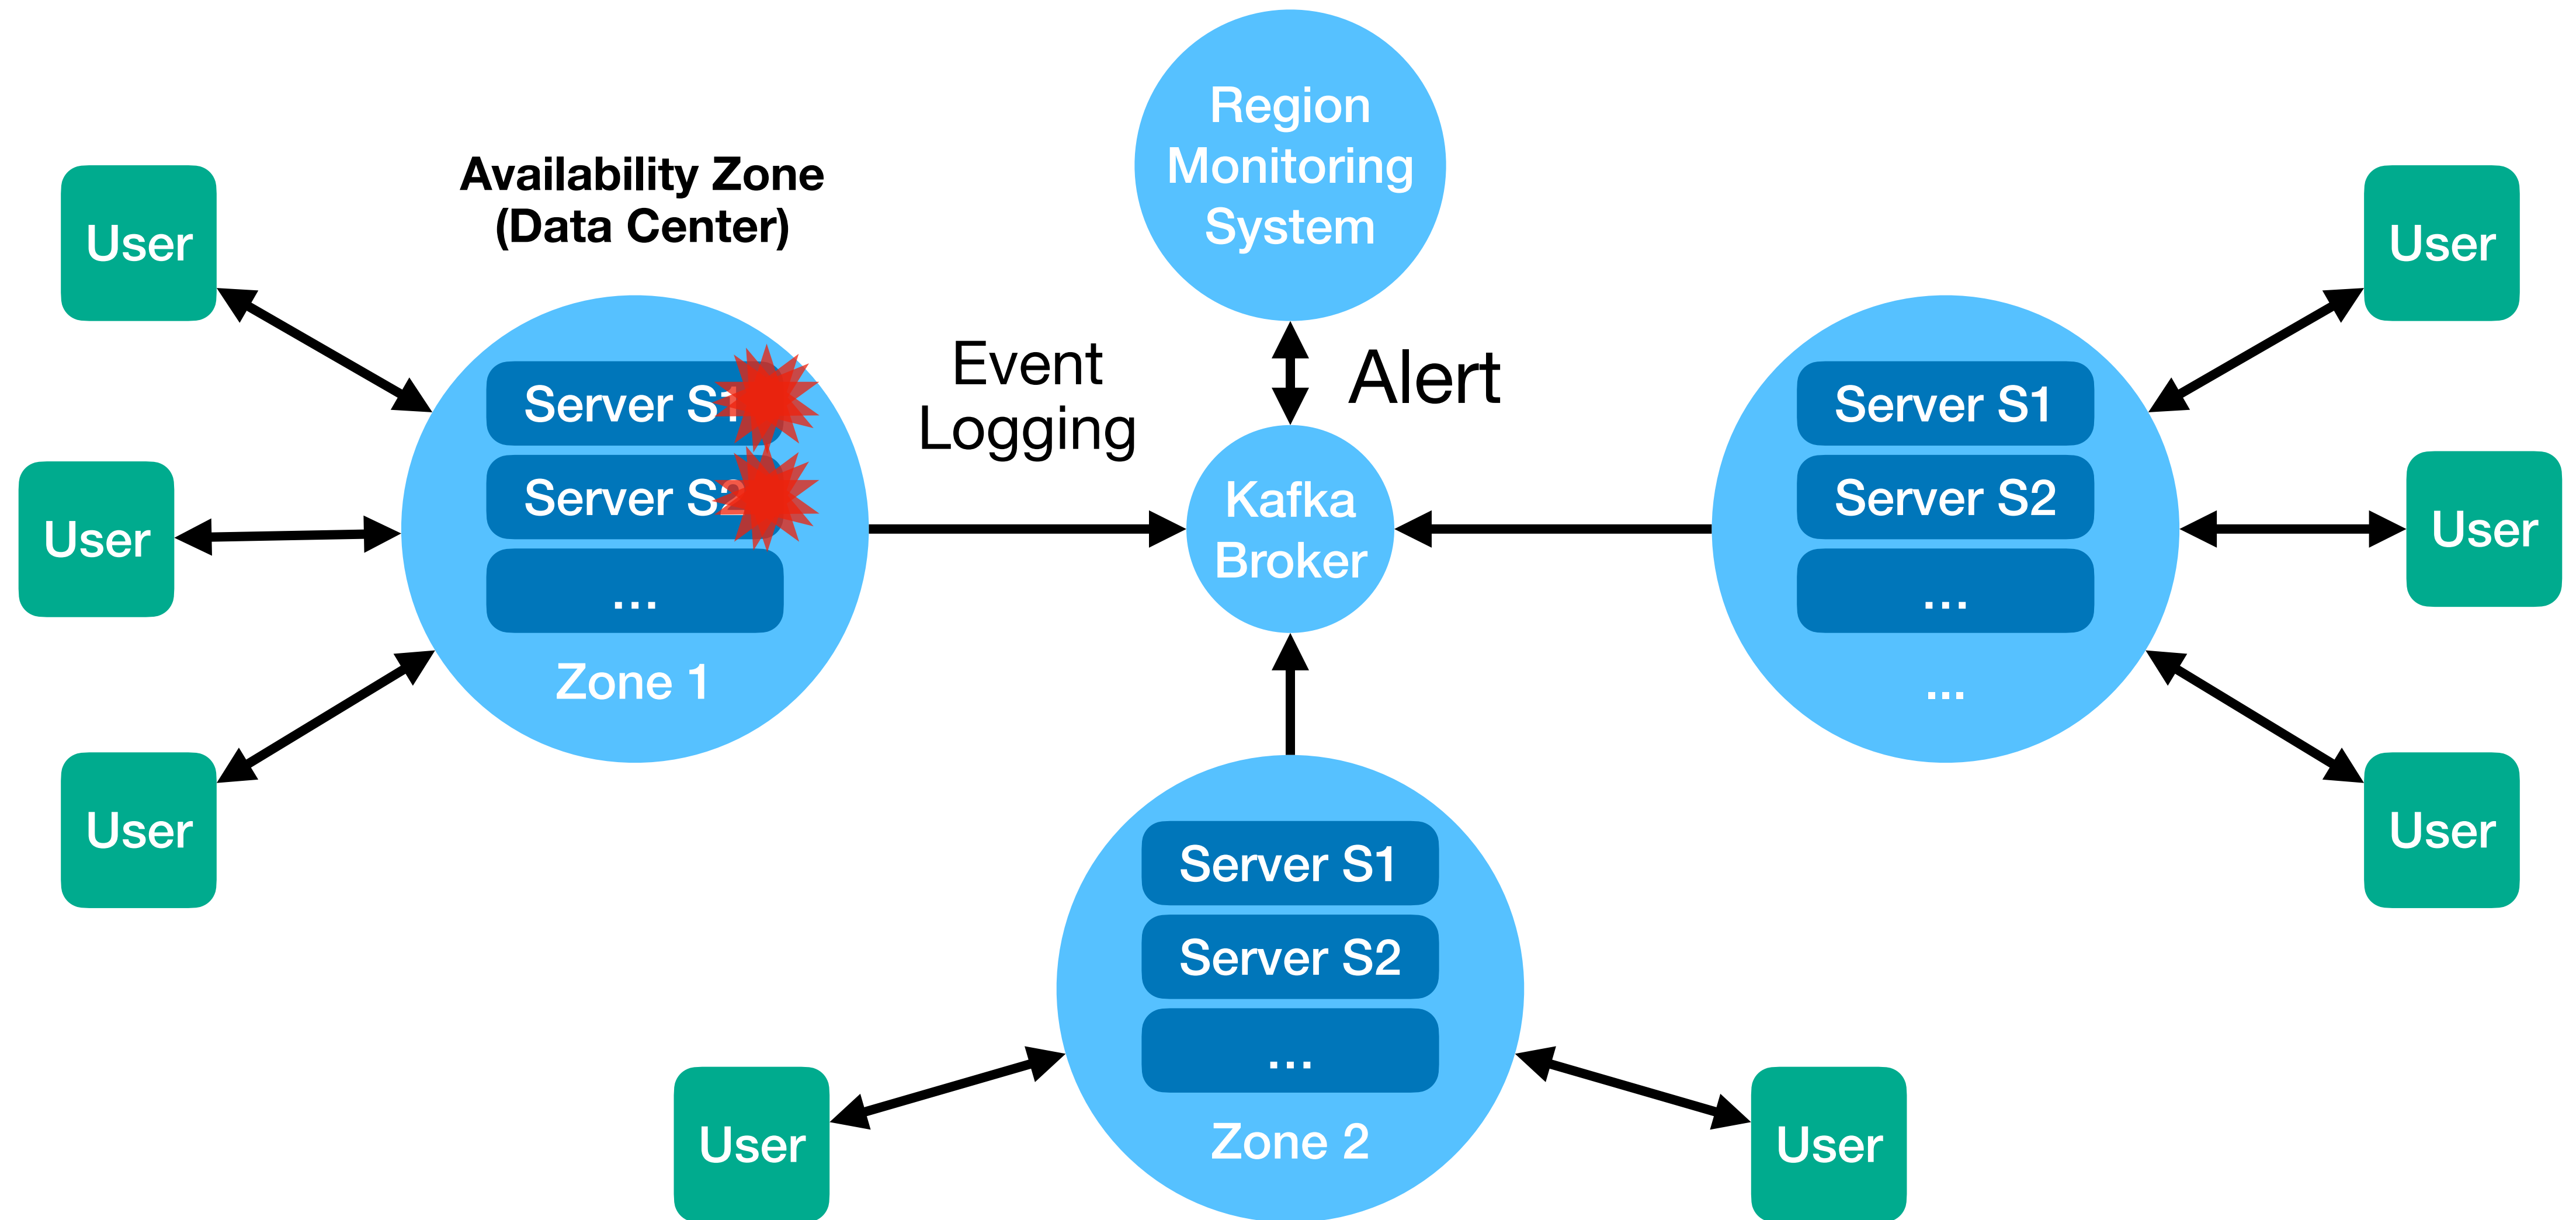
Example Application



Example Application



Example Application



Example Application - Log Events

Timestamp	Separator	Event Type	Error Code	Error Severity	Zone	Server	Compressed Message
2020-04-15	08:23:24	INFO			Z1	S2	0EGKKY ...
2020-04-15	08:23:30	ERROR	500	5	Z2	S1	9AE13K ...
2020-04-15	08:24:55	WARN			Z3	S3	mBvdqr ...
...							

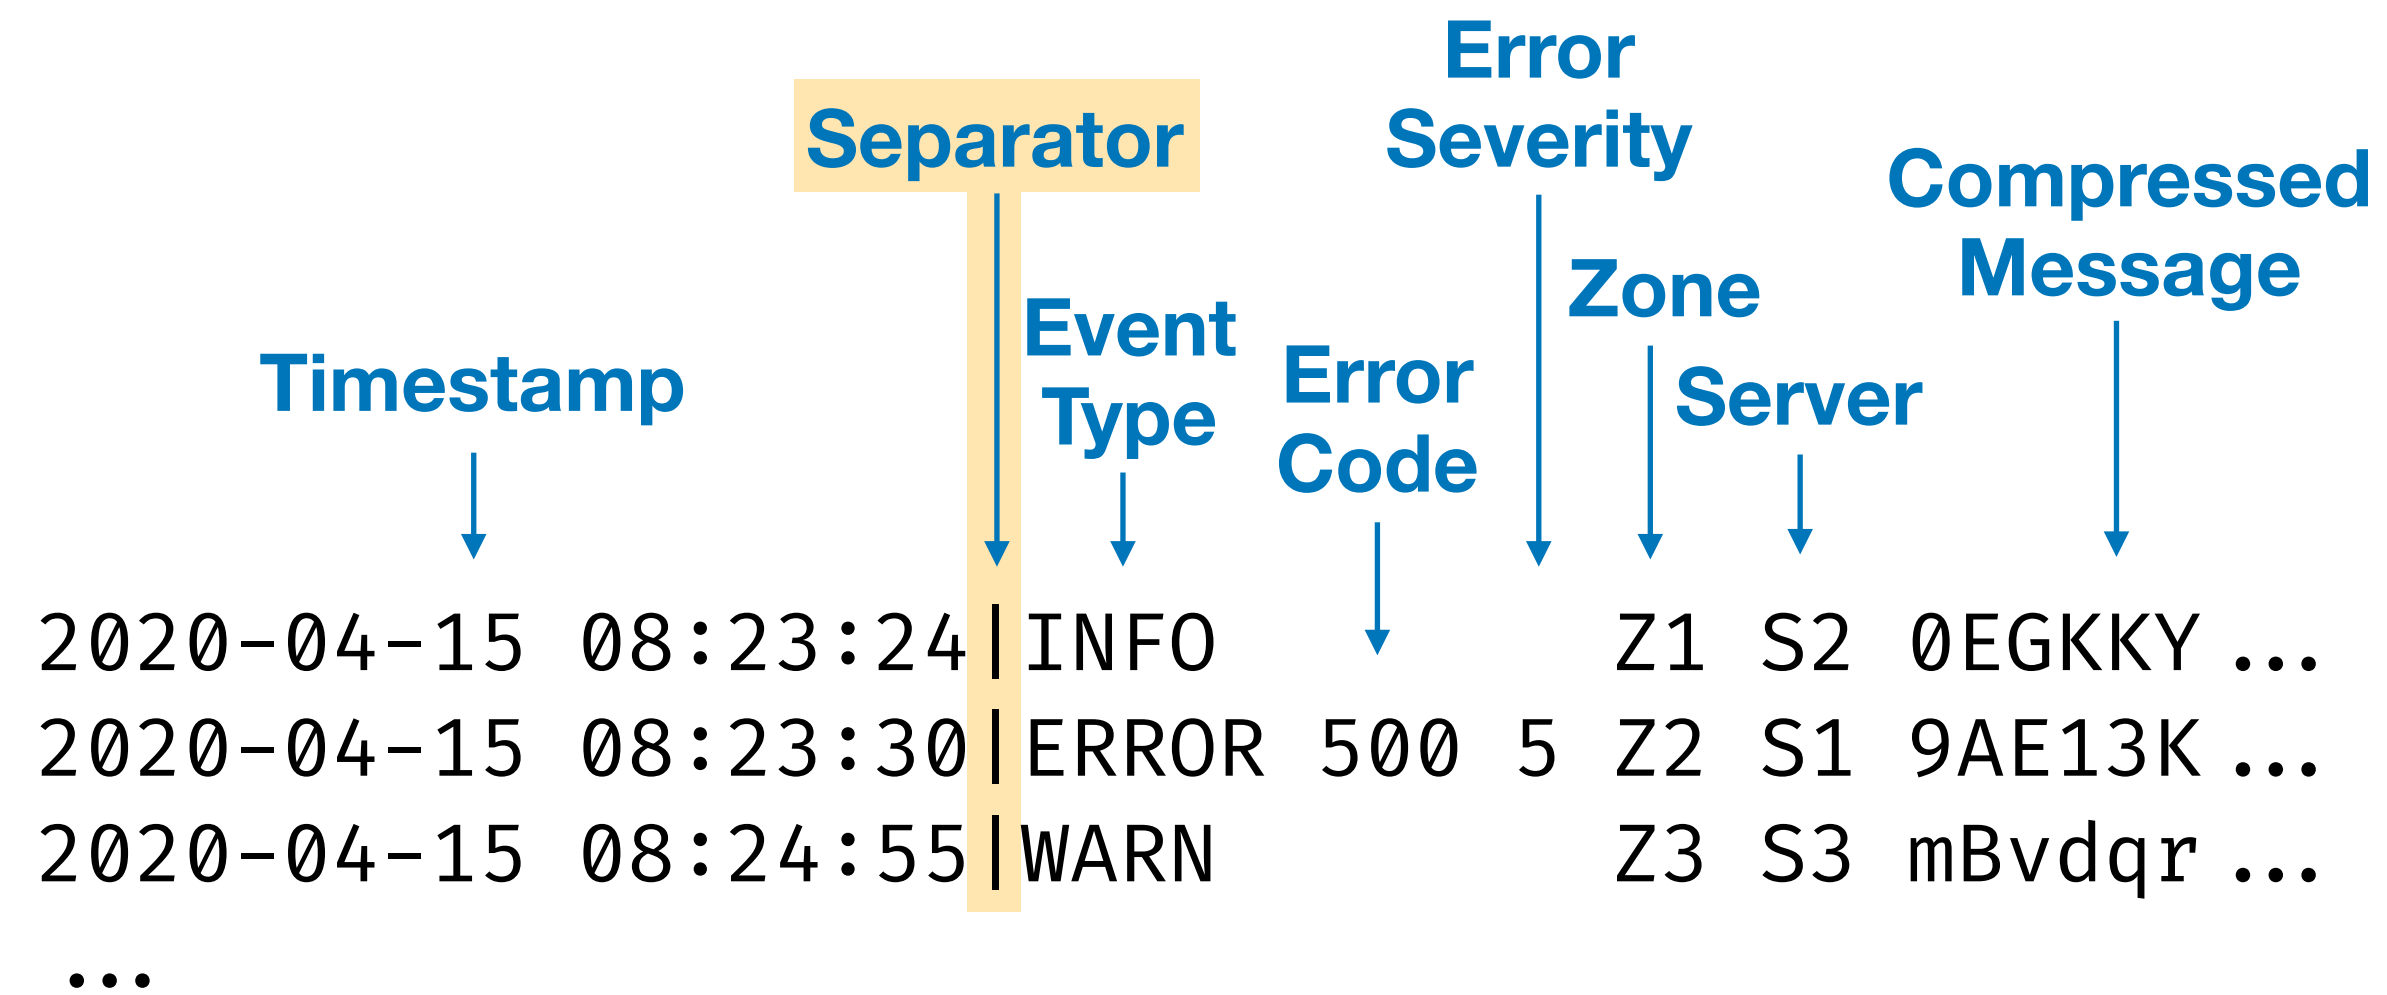
Example Application - Log Events

The diagram illustrates the structure of log events. Labels with arrows point to specific fields in the log entries:

- Timestamp** points to the date and time portion of the first two columns.
- Separator** points to the vertical bar character.
- Event Type** points to the text in the third column.
- Error Code** points to the numeric value in the fourth column.
- Error Severity** points to the numeric value in the fifth column.
- Zone** points to the text in the sixth column.
- Server** points to the text in the seventh column.
- Compressed Message** points to the alphanumeric string in the eighth column.

2020-04-15 08:23:24		INFO			Z1	S2	0EGKKY ...
2020-04-15 08:23:30		ERROR	500	5	Z2	S1	9AE13K ...
2020-04-15 08:24:55		WARN			Z3	S3	mBvdqr ...
...							

Example Application - Log Events



Example Application - Log Events

Timestamp	Separator	Event Type	Error Code	Error Severity	Zone	Server	Compressed Message
2020-04-15 08:23:24		INFO			Z1	S2	0EGKKY ...
2020-04-15 08:23:30		ERROR	500	5	Z2	S1	9AE13K ...
2020-04-15 08:24:55		WARN			Z3	S3	mBvdqr ...
...							

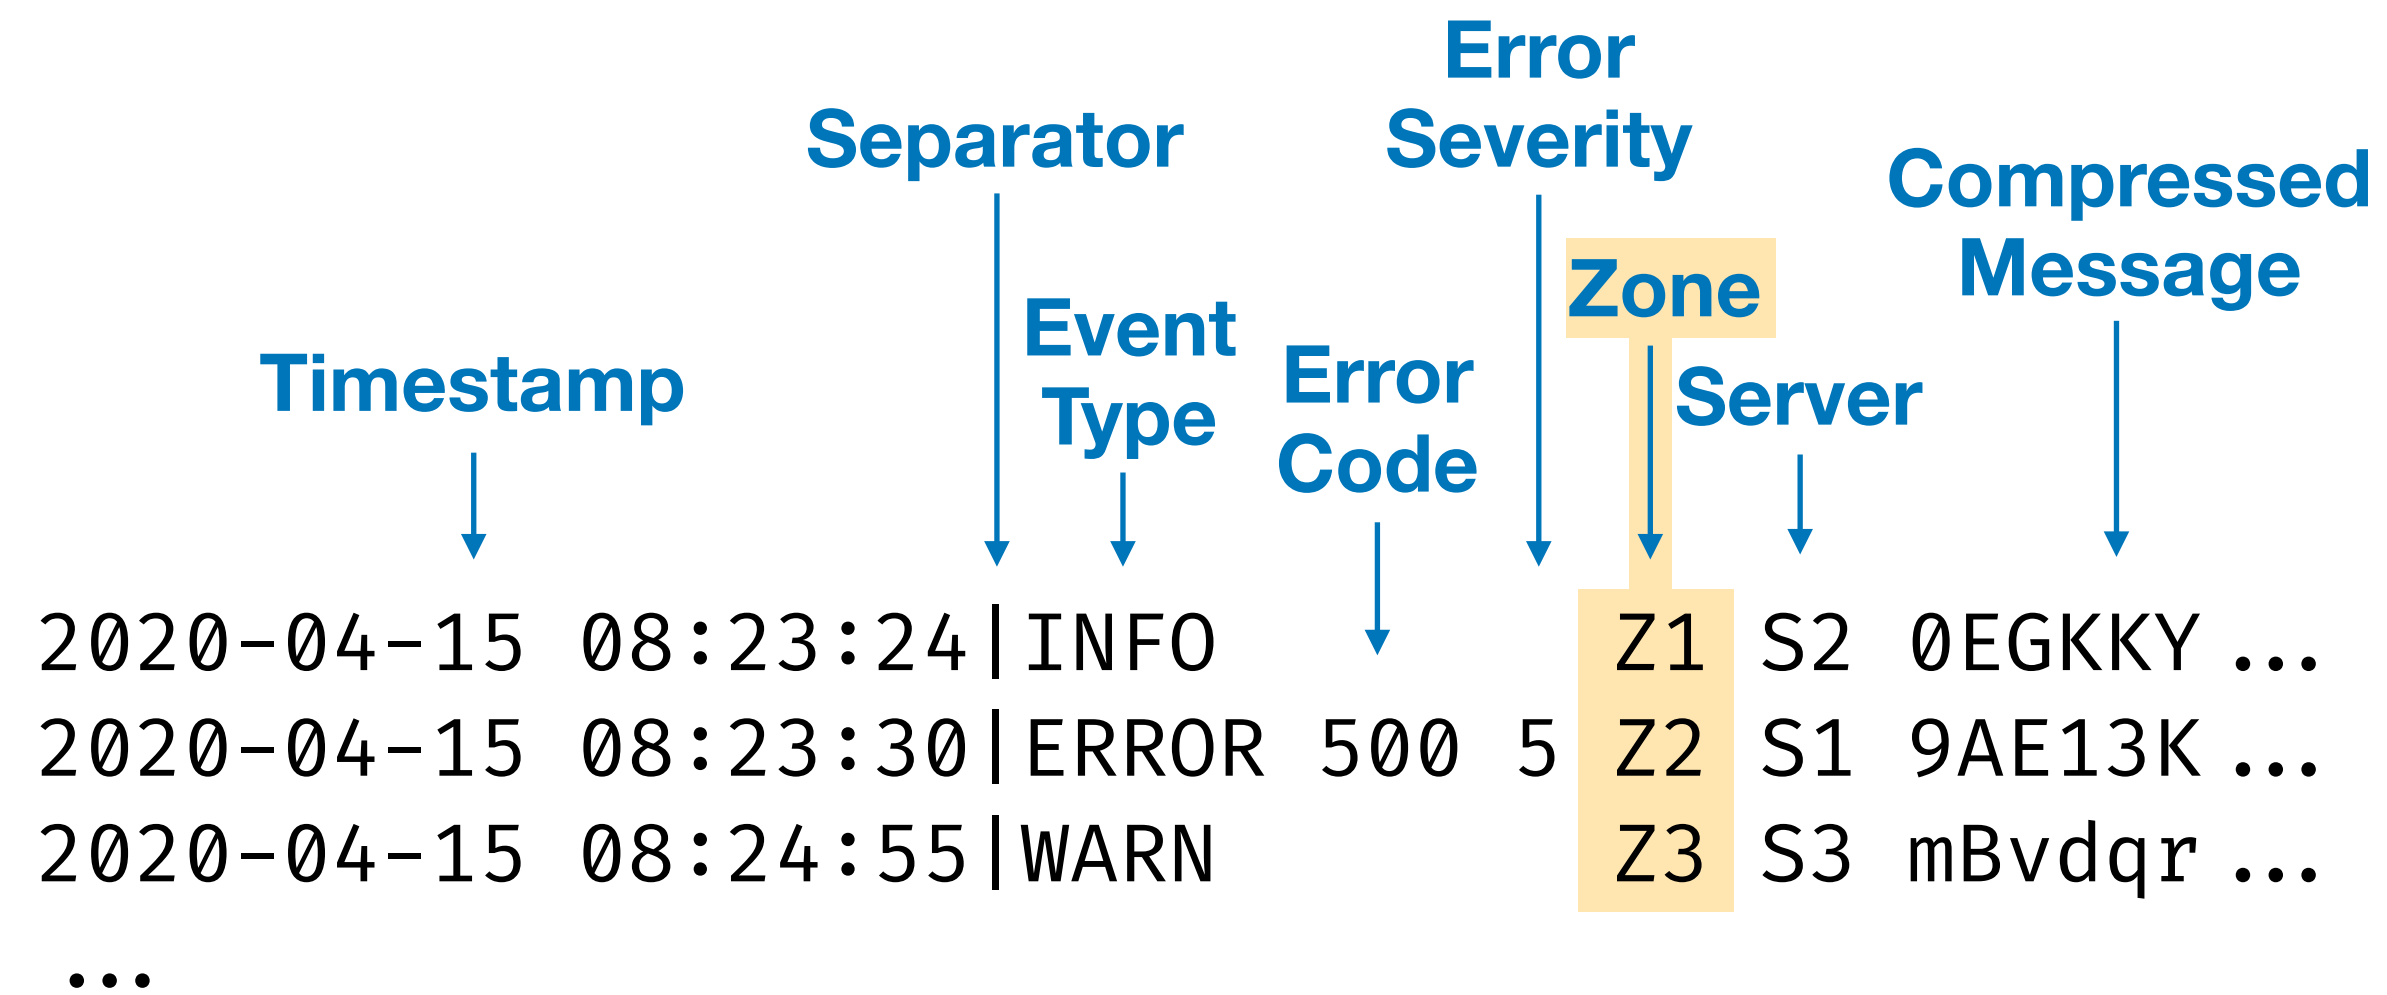
Example Application - Log Events

Timestamp	Separator	Event Type	Error Code	Error Severity	Zone	Server	Compressed Message
2020-04-15	08:23:24	INFO			Z1	S2	0EGKKY ...
2020-04-15	08:23:30	ERROR	500	5	Z2	S1	9AE13K ...
2020-04-15	08:24:55	WARN			Z3	S3	mBvdqr ...
...							

Example Application - Log Events

Timestamp	Separator	Event Type	Error Code	Error Severity	Zone	Server	Compressed Message
2020-04-15	08:23:24	INFO			Z1	S2	0EGKKY ...
2020-04-15	08:23:30	ERROR	500	5	Z2	S1	9AE13K ...
2020-04-15	08:24:55	WARN			Z3	S3	mBvdqr ...
...							

Example Application - Log Events



Example Application - Log Events

Timestamp	Separator	Event Type	Error Code	Error Severity	Zone	Server	Compressed Message
2020-04-15	08:23:24	INFO			Z1	S2	0EGKKY ...
2020-04-15	08:23:30	ERROR	500	5	Z2	S1	9AE13K ...
2020-04-15	08:24:55	WARN			Z3	S3	mBvdqr ...
...							

Example Application - Log Events


Timestamp	Separator	Event Type	Error Code	Error Severity	Zone	Server	Compressed Message
2020-04-15	08:23:24	INFO			Z1	S2	0EGKKY ...
2020-04-15	08:23:30	ERROR	500	5	Z2	S1	9AE13K ...
2020-04-15	08:24:55	WARN			Z3	S3	mBvdqr ...
...							

Example Application - Dataflow

Example Application - Dataflow

**Kafka
Event Source**


Example Application - Dataflow



Example Application - Dataflow




Example Application - Dataflow



Example Application - Dataflow



Example Application - Dataflow



Example Application - Dataflow



Example Application - Flink SQL

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)  
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)  
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
```


Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```


Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```


Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```

Example Application - Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, WATERMARK FOR ts AS ts - INTERVAL '5' SECONDS)
2 WITH ('connector'='kafka', 'topic'='eventlog', 'format'='csv', 'csv.field-delimiter'='|', ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, window_start TIMESTAMP(3), window_end TIMESTAMP(3))
5 WITH ('connector'='kafka', 'topic'='alerts', 'format'='json', ... );
6
7 > CREATE VIEW ErrorEvents AS
8 SELECT ts,
9         CAST(SPLIT_INDEX(line, ' ', 1) AS INT) AS code,
10        CAST(SPLIT_INDEX(line, ' ', 2) AS INT) AS severity,
11        SPLIT_INDEX(line, ' ', 3) AS zone,
12        SPLIT_INDEX(line, ' ', 4) AS server,
13        SPLIT_INDEX(line, ' ', 5) AS message
14 FROM RawEvents
15 WHERE SPLIT_INDEX(line, ' ', 0) = 'ERROR';
16
17 > INSERT INTO Alerts
18 SELECT zone, AVG(severity) AS avg_severity,
19        window_start, window_end,
20 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
21                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
22 WHERE code = 500
23 GROUP BY window_start, window_end, zone
24 HAVING AVG(severity) > 2;
```

How to isolate the parsing code for reuse and testing?

Example Application - UDFs

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ... )
2 WITH ( ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ... )
5 WITH ( ... );
6
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7         zone STRING, server STRING,
8         message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Example Application - UDFs

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Example Application - UDFs

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ... )
2 WITH ( ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ... )
5 WITH ( ... );
6
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7         zone STRING, server STRING,
8         message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Example Application - UDFs

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ... )
2 WITH ( ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ... )
5 WITH ( ... );
6
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7         zone STRING, server STRING,
8         message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```


Example Application - UDFs

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ... )
2 WITH ( ... );
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ... )
5 WITH ( ... );
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Example Application - UDFs

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```


Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Problem:
Separate compilation

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > I
15 SEL $ avg_severity,
16 end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18 INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Problem:
Cannot pushdown
predicate into UDF

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7         zone STRING, server STRING,
8         message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Problem:
Cannot pushdown
projection into UDF

Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<severity INT, code INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         int eventType = Integer.parseInt(parts[0]);
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(severity, code, zone, server, message));
20         }
21     }
22 }
23 }
```

Problem:
Impedance mismatch

Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18               INTERVAL '2' MINUTES, INTERVAL '10' M
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             System.out.println("Parsed line: " + line);
20             collect(Row.of(code, severity, zone, server, message));
21         }
22     }
23 }
24 }
```

Problem:
Undefined behaviour

Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18               INTERVAL '2' MINUTES, INTERVAL '10' M
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             System.out.println("Parsed line: " + line);
20             collect(Row.of(code, severity, zone, server, message));
21         }
22     }
23 }
24 }
```

Problem:
Undefined behaviour

Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```


Example Application - Problems?

Flink SQL

```
1 > CREATE TABLE RawEvents(ts TIMESTAMP(3), line STRING, ...)
2 WITH (...);
3
4 > CREATE TABLE Alerts(zone STRING, avg_severity INT, ...)
5 WITH (...);
6
7 > ADD JAR 'file:///UDFs/target/UDFs-1.0.jar';
8 > CREATE FUNCTION Parse AS 'com.example.Parse' LANGUAGE JAVA;
9
10 > CREATE VIEW ErrorEvents AS
11 SELECT ts, parsed.*
12 FROM RawEvents, LATERAL TABLE(Parse(line)) AS parsed;
13
14 > INSERT INTO Alerts
15 SELECT zone, AVG(severity) AS avg_severity,
16        window_start, window_end
17 FROM TABLE(HOP(TABLE ErrorEvents, DESCRIPTOR(ts),
18                INTERVAL '2' MINUTES, INTERVAL '10' MINUTES))
19 WHERE code = 500
20 GROUP BY window_start, window_end, zone
21 HAVING AVG(severity) > 2;
```

Flink Java API (Maven project)

```
1 package com.example;
2 import org.apache.flink.*;
3
4 @FunctionHint(
5     input = @DataTypeHint("STRING"),
6     output = @DataTypeHint("ROW<code INT, severity INT,
7                             zone STRING, server STRING,
8                             message STRING>"))
9 public class Parse extends TableFunction<Row> {
10     public void eval(String line) {
11         String[] parts = line.split(" ");
12         String eventType = parts[0];
13         if (eventType.equals("ERROR")) {
14             int code = Integer.parseInt(parts[1]);
15             int severity = Integer.parseInt(parts[2]);
16             String zone = parts[3];
17             String server = parts[4];
18             String message = parts[5];
19             collect(Row.of(code, severity, zone, server, message));
20         }
21     }
22 }
23 }
```

Problem:
Close coupling

Other options?

Other options?

Flink DataStream API

```
1 StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
2
3 env.addSource(new FlinkKafkaConsumer<>("eventlog", ... ))
4     .assignTimestampAndWatermarks( ... )
5     .flatMap(new Parse())
6     .filter(event → event.getCode() = 500)
7     .map(event → Tuple2.of(event.getZone(), event.getSeverity()))
8     .keyBy(tuple → tuple.f0)
9     .window(SlidingEventTimeWindows.of(Time.minutes(10), Time.minutes(1)))
10    .apply((zone, window, errors, out) → {
11        double avgSeverity = errors.map(error → error.severity).average();
12        out.collect(Tuple4.of(window.getStart(), window.getEnd(), zone, avgSeverity));
13    })
14    .map(tuple → tuple.f0 + "," + tuple.f1 + "," + tuple.f2 + "," + tuple.f3)
15    .addSink(new FlinkKafkaProducer<>("alerts", ... ));
16
17 env.execute();
```

Other options?

Flink DataStream API

```
1 StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
2
3 env.addSource(new FlinkKafkaConsumer<>("eventlog", ... ))
4     .assignTimestampAndWatermarks( ... )
5     .flatMap(new Parse())
6     .filter(event → event.getCode() = 500)
7     .map(event → Tuple2.of(event.getZone(), event.getSeverity()))
8     .keyBy(tuple → tuple.f0)
9     .window(SlidingEventTimeWindows.of(Time.minutes(10), Time.minutes(1)))
10    .apply((zone, window, errors, out) → {
11        double avgSeverity = errors.map(error → error.severity).average();
12        out.collect(Tuple4.of(window.getStart(), window.getEnd(), zone, avgSeverity));
13    })
14    .map(tuple → tuple.f0 + "," + tuple.f1 + "," + tuple.f2 + "," + tuple.f3)
15    .addSink(new FlinkKafkaProducer<>("alerts", ... ));
16
17 env.execute();
```

Other options?

Flink DataStream API

```
1 StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
2
3 env.addSource(new FlinkKafkaConsumer<>("eventlog", ... ))
4     .assignTimestampAndWatermarks( ... )
5     .flatMap(new Parse())
6     .filter(event → event.getCode() = 500)
7     .map(event → Tuple2.of(event.getZone(), event.getSeverity()))
8     .keyBy(tuple → tuple.f0)
9     .window(SlidingEventTimeWindows.of(Time.minutes(10), Time.minutes(1)))
10    .apply((zone, window, errors, out) → {
11        double avgSeverity = errors.map(error → error.severity).average();
12        out.collect(Tuple4.of(window.getStart(), window.getEnd(), zone, avgSeverity));
13    })
14    .map(tuple → tuple.f0 + "," + tuple.f1 + "," + tuple.f2 + "," + tuple.f3)
15    .addSink(new FlinkKafkaProducer<>("alerts", ... ));
16
17 env.execute();
```

Other options?

Flink DataStream API

```
1 StreamExecutionEnvironment env = StreamExecutionEnvironment.  
2  
3 env.addSource(new FlinkKafkaConsumer<>("eventlog", ... ))  
4   .assignTimestampAndWatermarks( ... )  
5   .flatMap(new Parse())  
6   .filter(event → event.getCode() = 500)  
7   .map(event → Tuple2.of(event.getZone(), event.getSeverity()))  
8   .keyBy(tuple → tuple.f0)  
9   .window(SlidingEventTimeWindows.of(Time.minutes(10), Time.minutes(1)))  
10  .apply((zone, window, events, out) → {  
11      double avgSeverity = events.map(event → event.severity).average();  
12      out.collect(Tuple4.of(window.getStart(), window.getEnd(), zone, avgSeverity));  
13  })  
14  .map(tuple → tuple.f0 + "," + tuple.f1 + "," + tuple.f2 + "," + tuple.f3)  
15  .addSink(new FlinkKafkaProducer<>("alerts", ... ));  
16  
17 env.execute();
```

Runtime Error: The return type of function 'org.example.Main\$
\$Lambda\$263/0x00000008002e3040@36074e47' could not be
determined automatically, due to **type erasure**. You can give
type information hints by using the **returns(...)** method on the
result of the transformation call.

Other options?

Flink DataStream API

```
1 StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
2
3 env.addSource(new FlinkKafkaConsumer<>("eventlog", ... ))
4     .assignTimestampAndWatermarks( ... )
5     .flatMap(new Parse())
6     .filter(event → event.getCode() = 500)
7     .map(event → Tuple2.of(event.getZone(), event.getSeverity()))
8     .returns(TypeInformation.of(new TypeHint<Tuple2<String, Integer>>() {}))
9     .keyBy(tuple → tuple.f0)
10    .window(SlidingEventTimeWindows.of(Time.minutes(10), Time.minutes(1)))
11    .apply((zone, window, events, out) → {
12        double avgSeverity = events.map(event → event.severity).average();
13        out.collect(Tuple4.of(window.getStart(), window.getEnd(), zone, avgSeverity));
14    })
15    .map(tuple → tuple.f0 + "," + tuple.f1 + "," + tuple.f2 + "," + tuple.f3)
16    .addSink(new FlinkKafkaProducer<>("alerts", ... ));
17
18 env.execute();
```

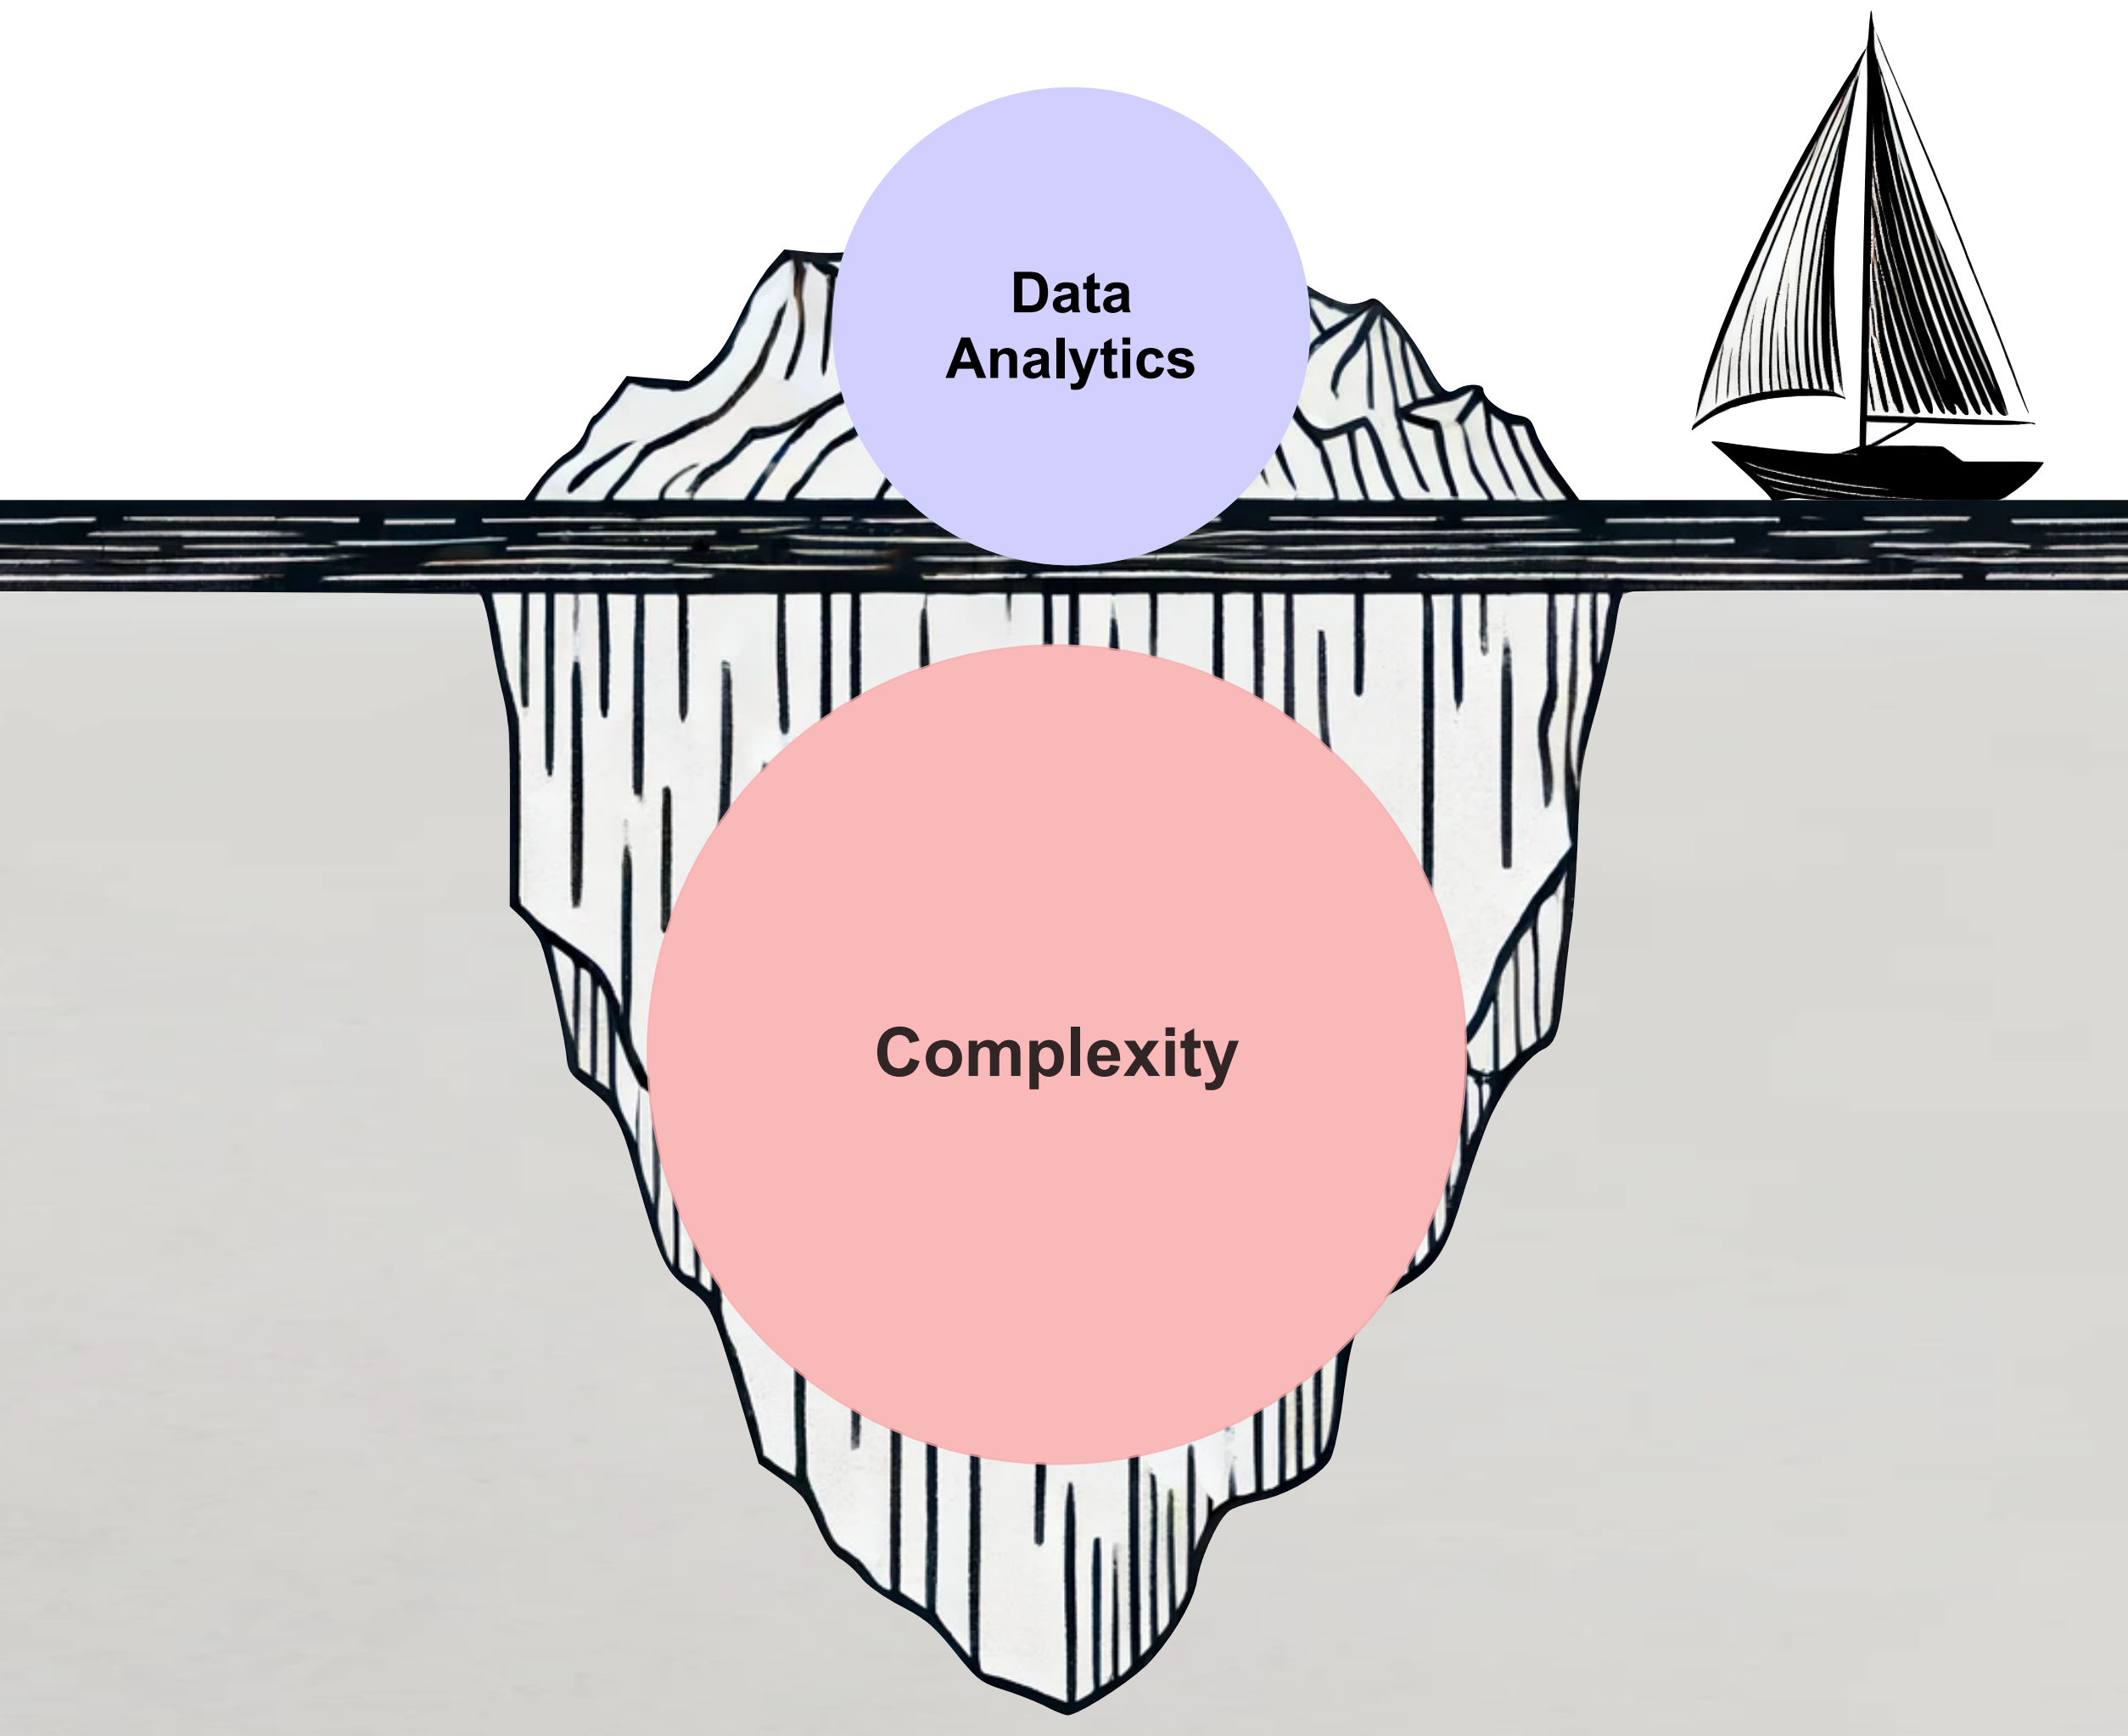
The Reality of Dataflow Systems

Expectation

User

**Data
Analytics**

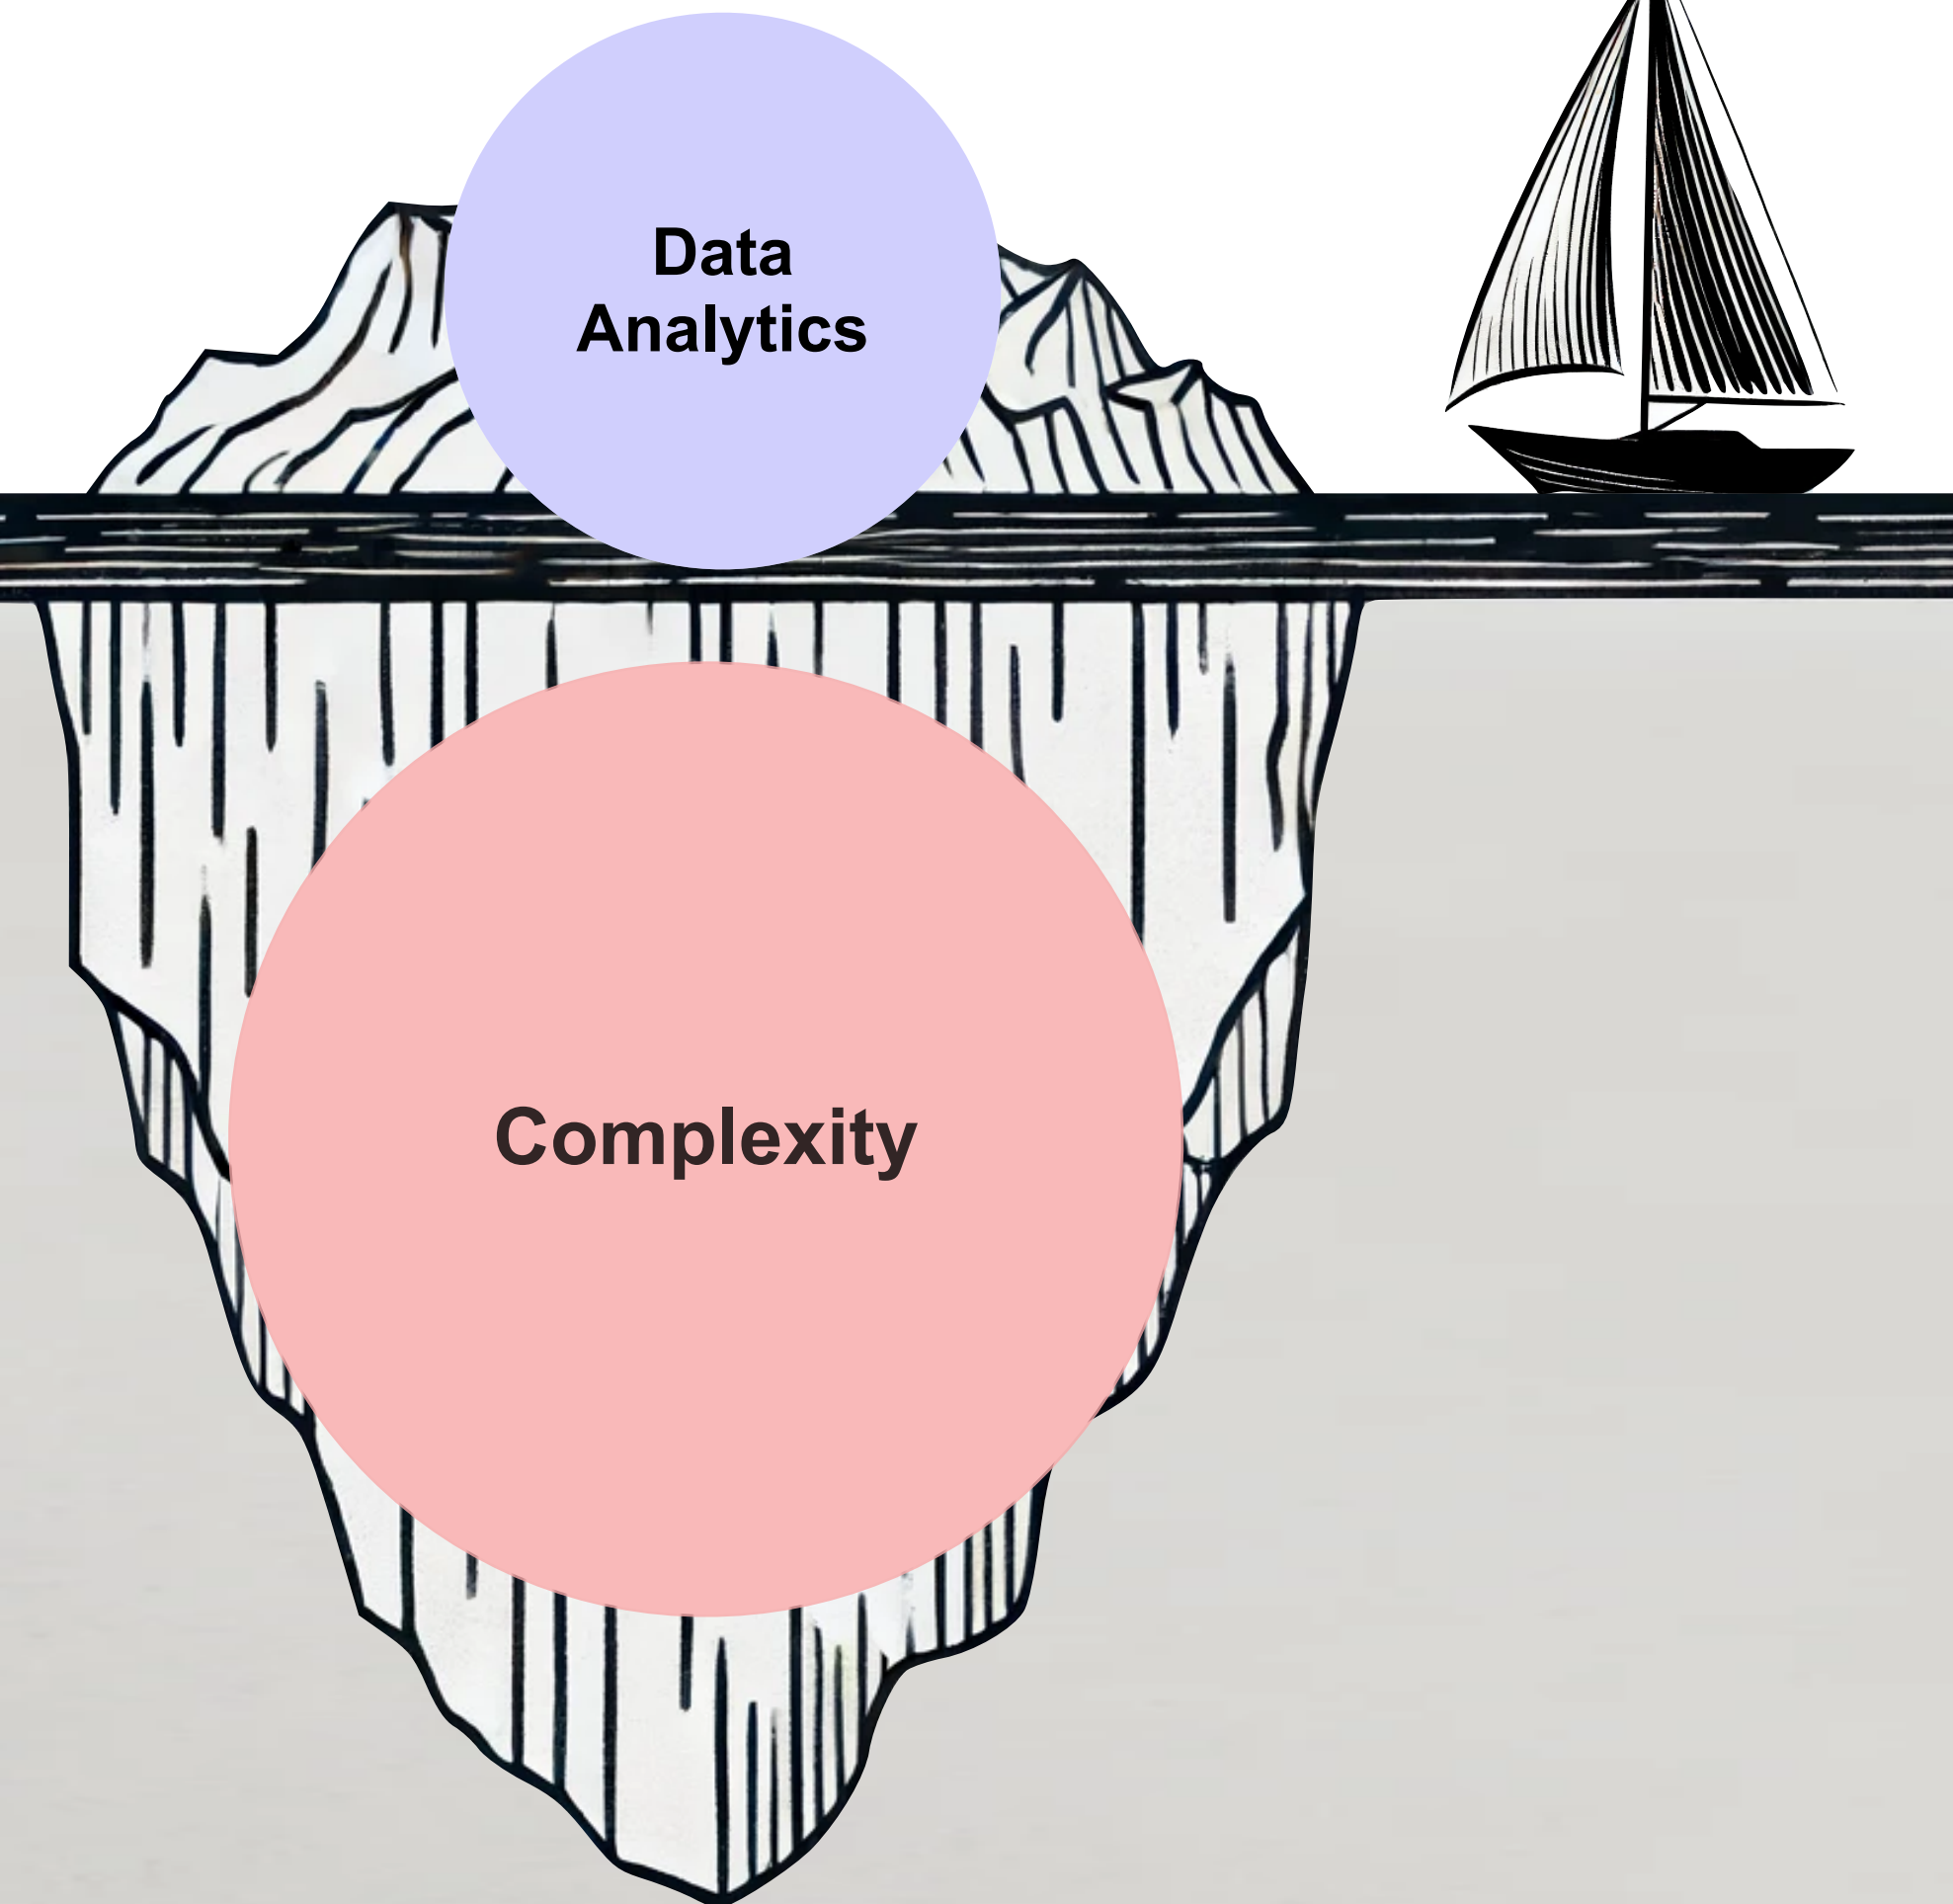
Complexity




The Reality of Dataflow Systems

Expectation

User



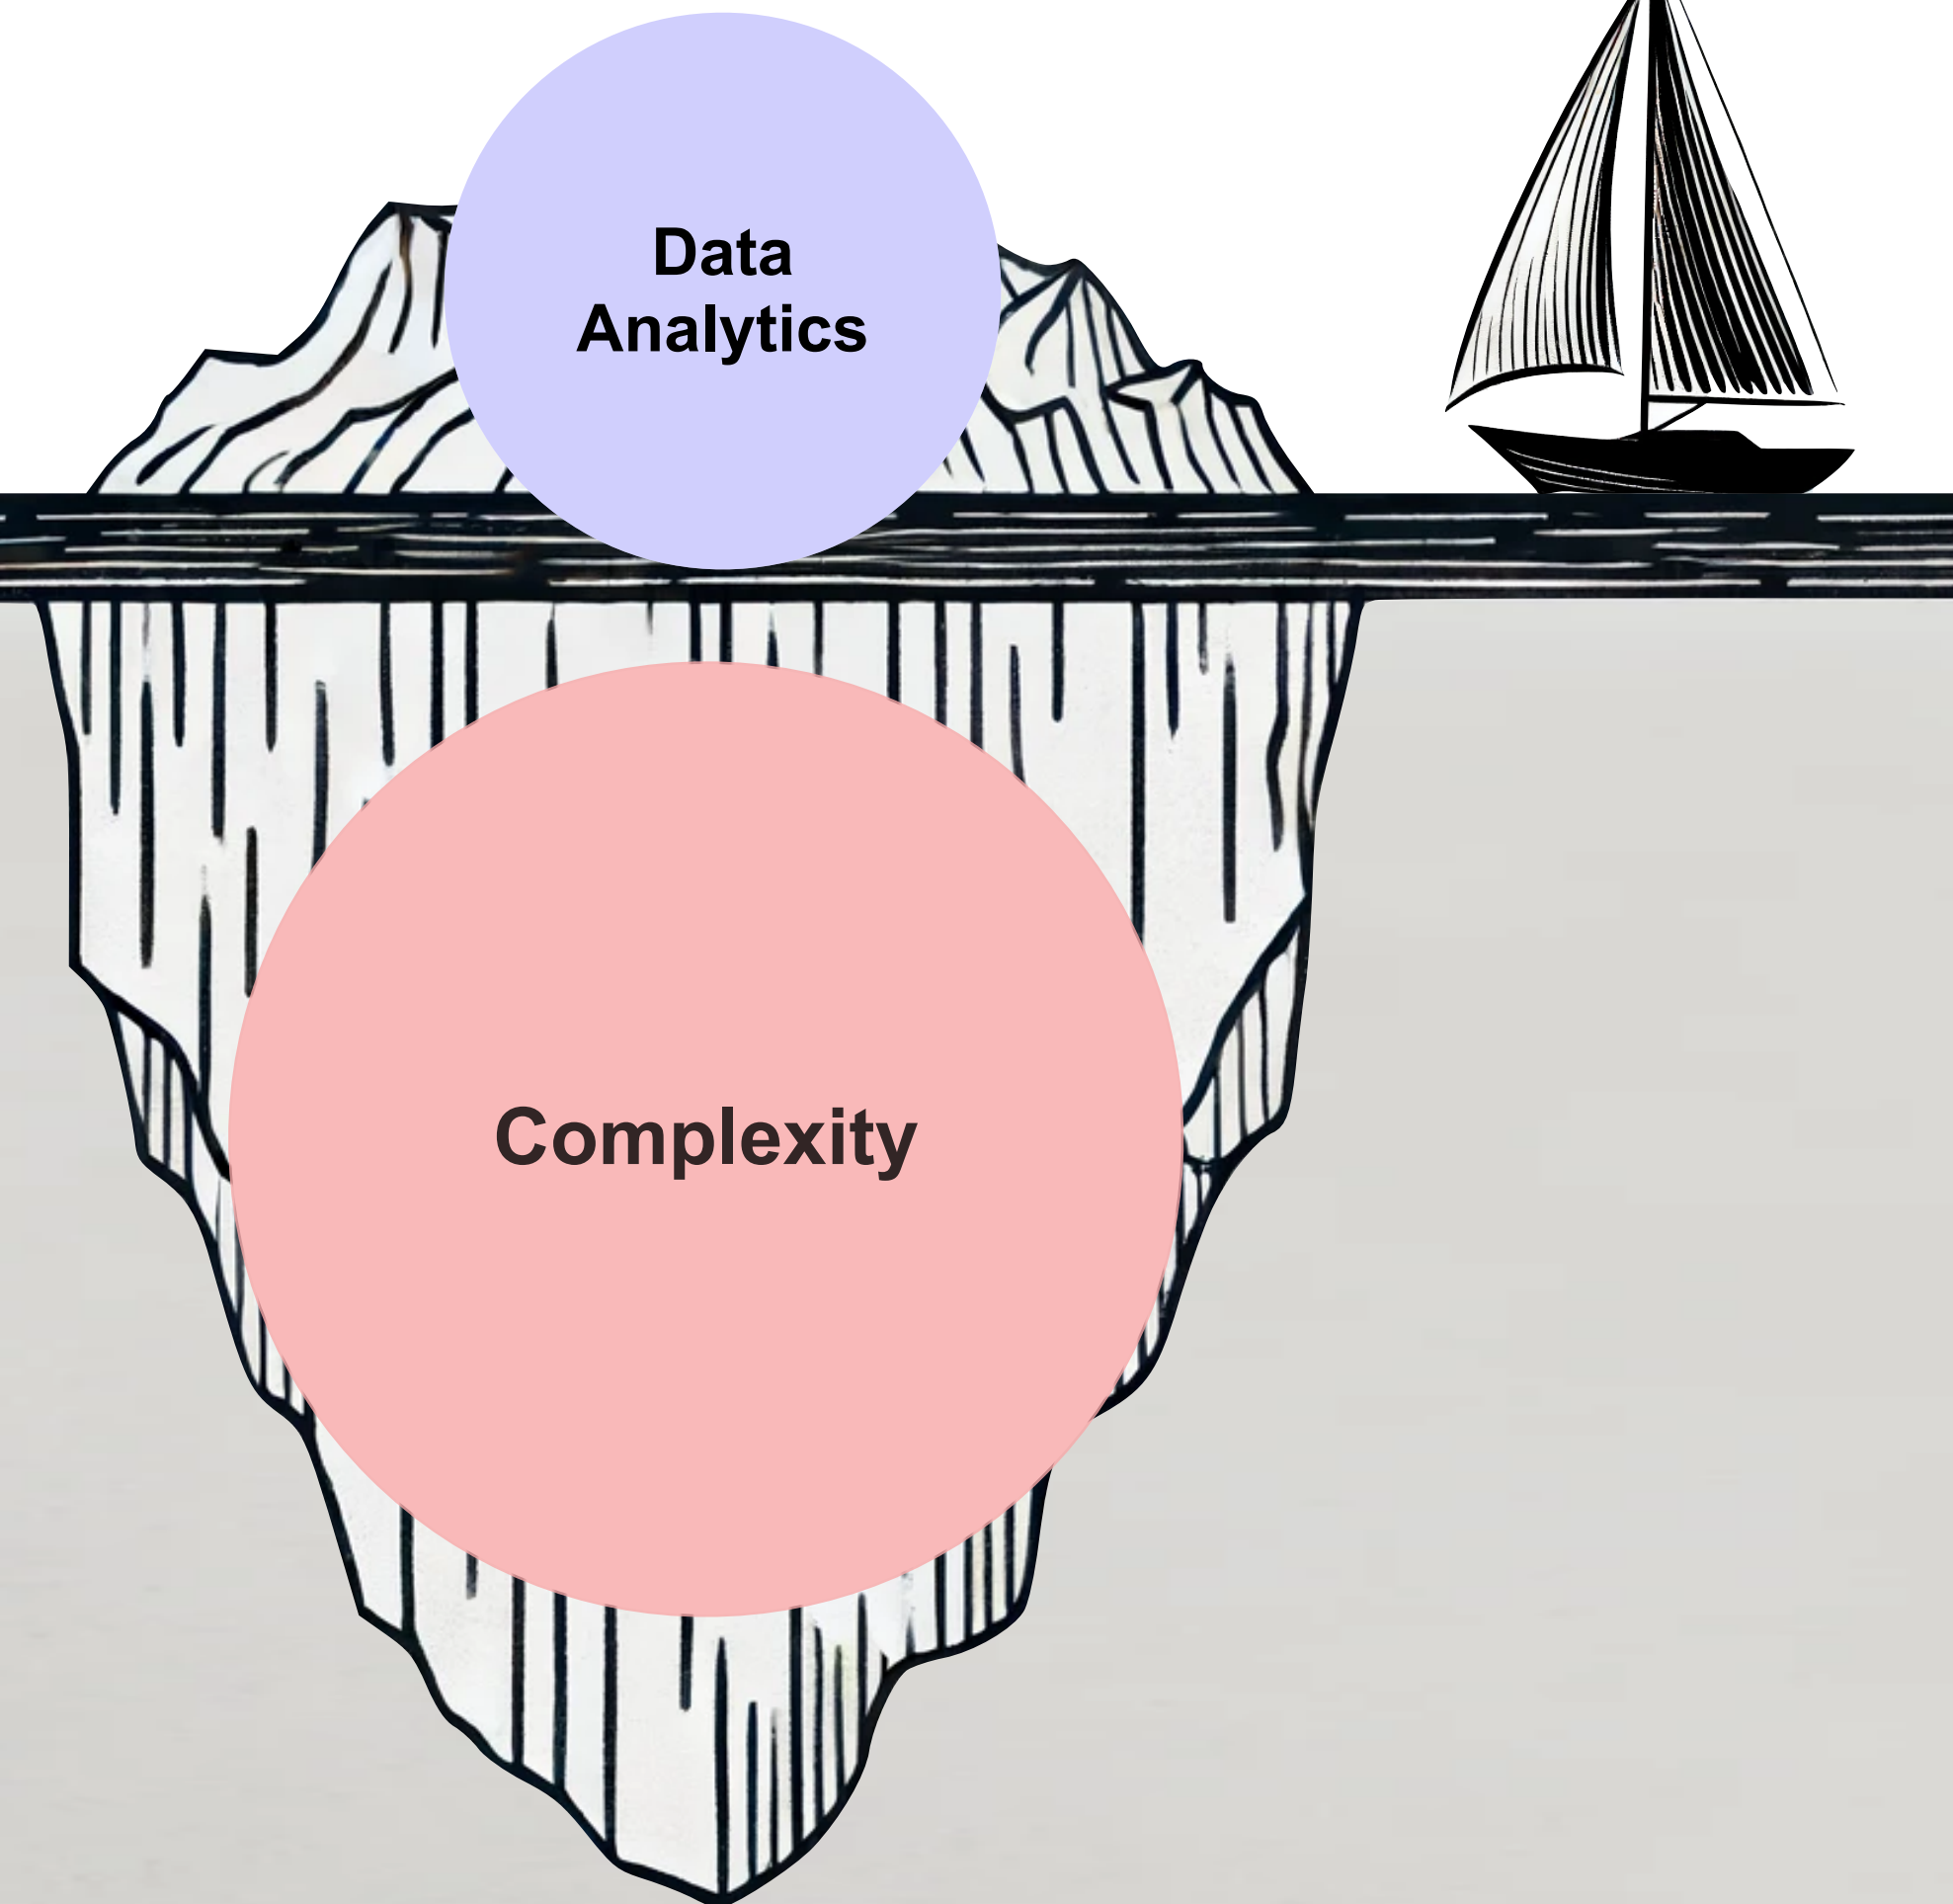
Reality



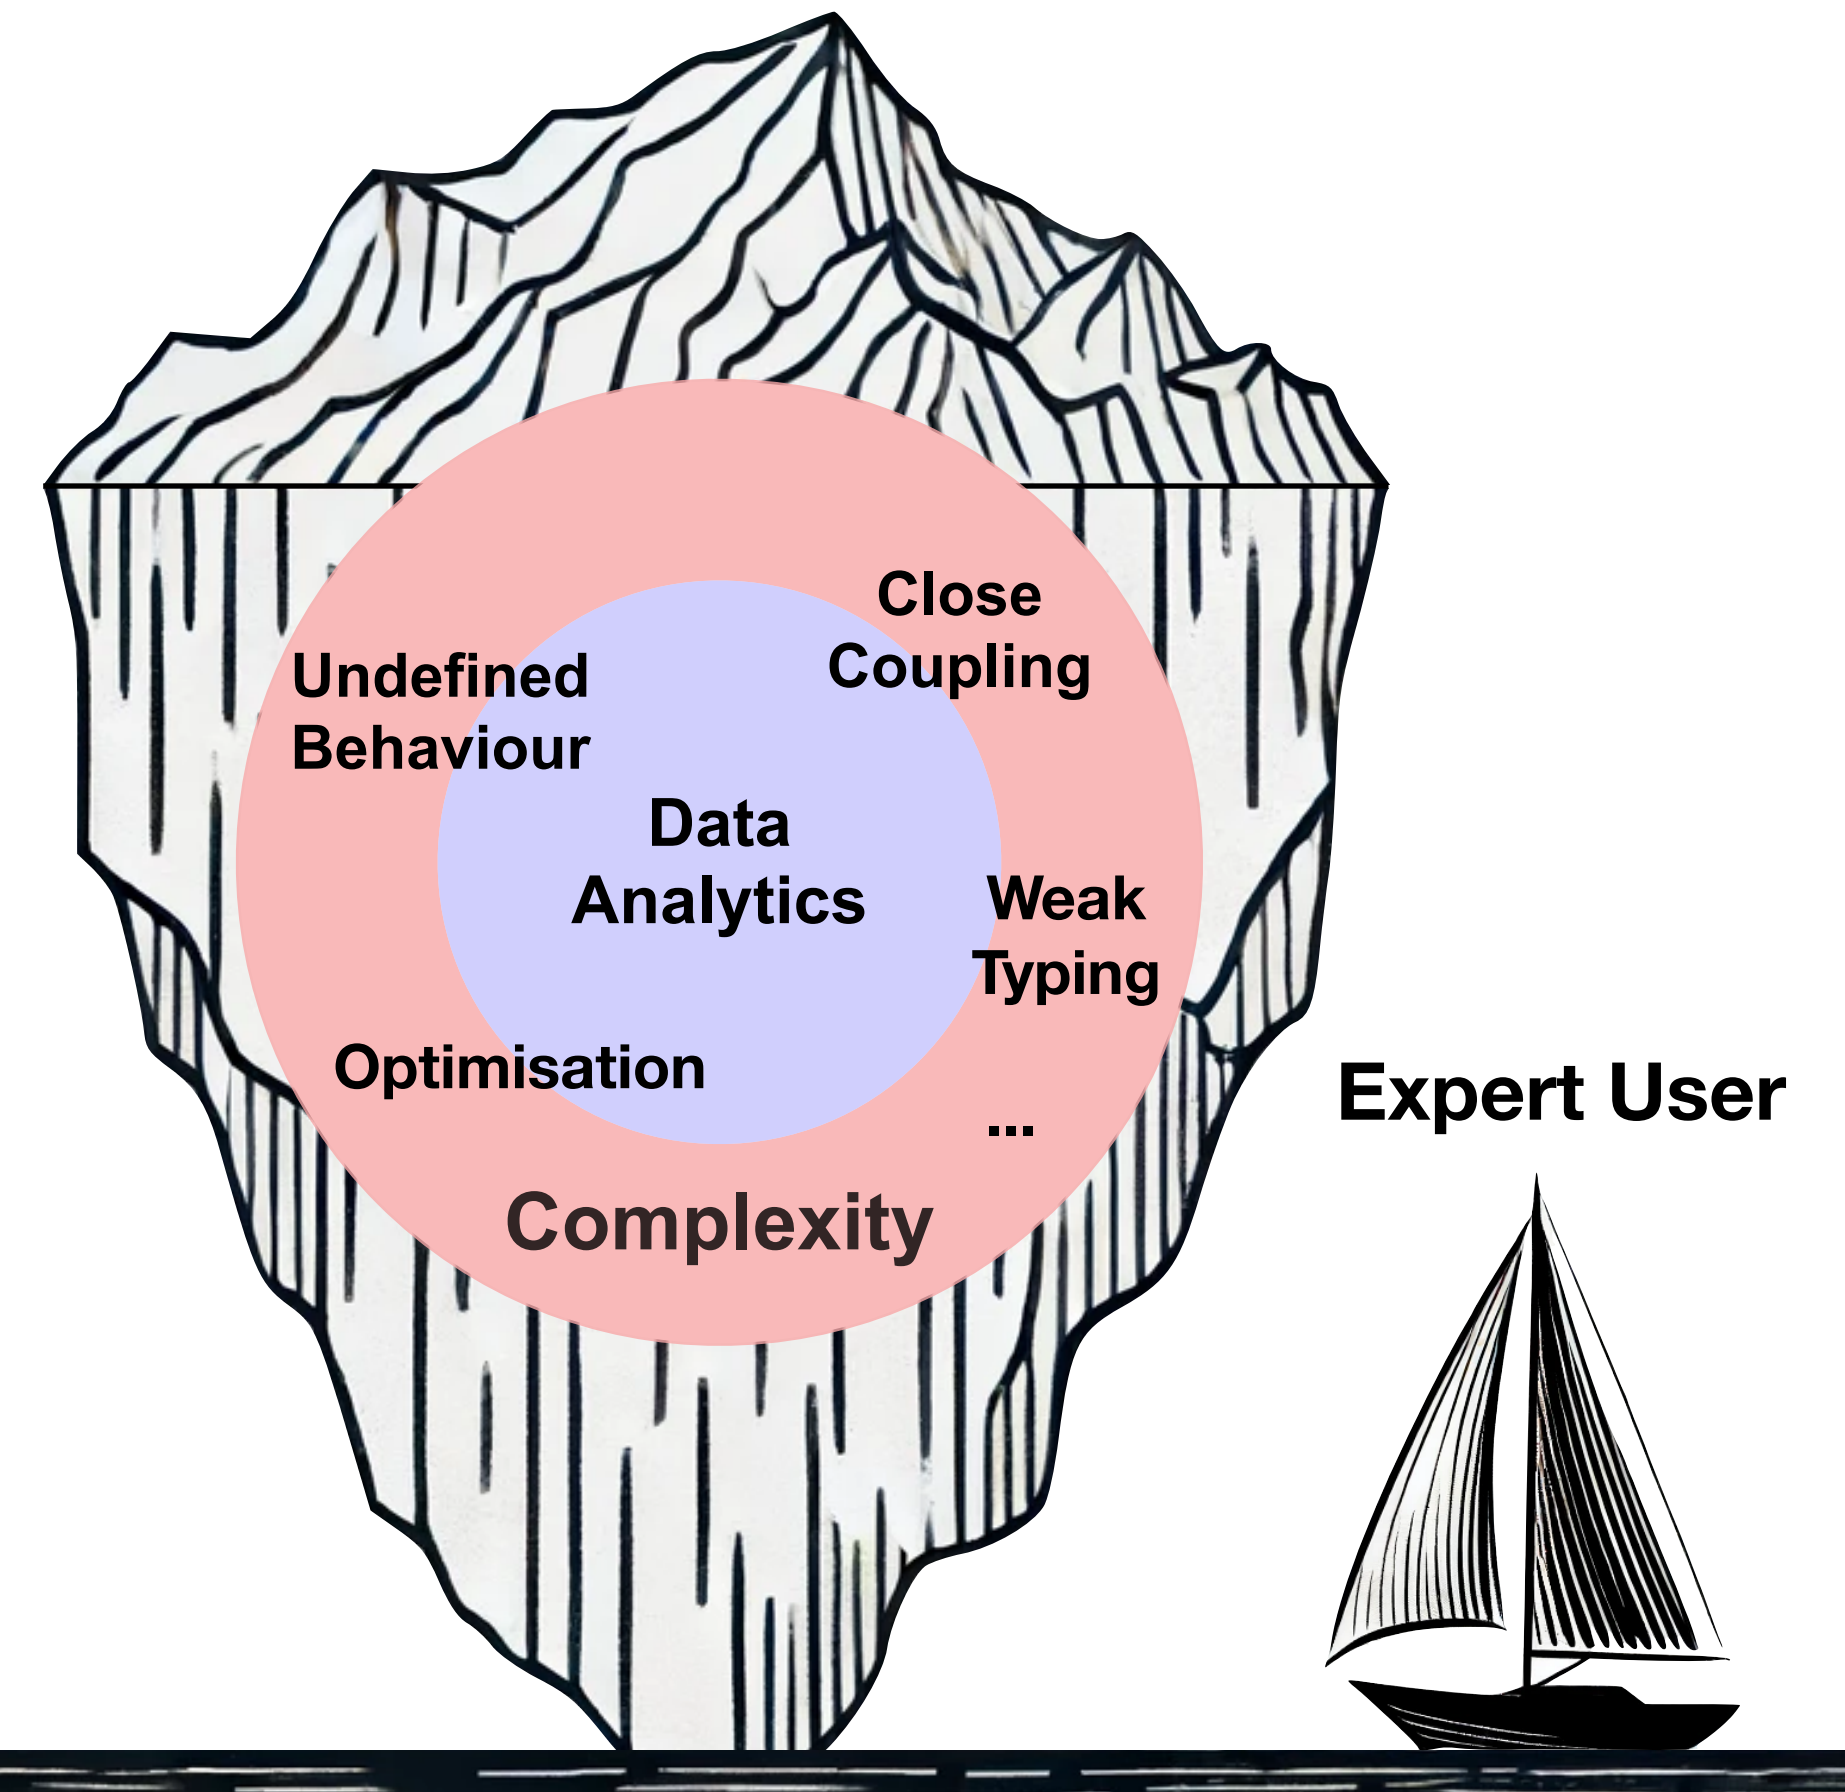
The Reality of Dataflow Systems

Expectation

User



Reality



Expert User

Solutions?



Solutions?

**Extend Flink SQL
with support for
User Defined Functions**




Solutions?

**Extend Flink SQL
with support for
User Defined Functions**



**Create a new
language**



AquaLang: Dataflow Programming

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);  
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
```


AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[ErrorEvent] = {
5     val parts = line.split(" ");
6     val eventType = parts[0];
7     if eventType == "ERROR" {
8         val code = parts[1].parse();
9         val severity = parts[2].parse();
10        val zone = parts[3];
11        val server = parts[4];
12        val message = parts[5];
13        Some(ErrorEvent(code, severity, zone, server, message))
14    } else {
15        None
16    }
17 }
```

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[ErrorEvent] = { ... }
5
6 Stream[RawEvent] :: source(Reader :: kafka("eventlog", ...), Format :: csv('|'), _.time)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window :: sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(Writer :: kafka("alerts", ...), Format :: json())
16  .run(Backend :: flink(...));
```

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, severity:u8);
3
4 def parse(line: String): Option[ErrorEvent] = { ... }
5
6 Stream[RawEvent]::source(Reader::kafka("eventlog", ...), Format::csv('|'), _.time)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(Writer::kafka("alerts", ...), Format::json())
16  .run(Backend::flink(...));
```

No undefined behaviour:
UDFs are verified to be
side-effect free.

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[ErrorEvent] = { ... }
5
6 Stream[RawEvent]::source(Reader::kafka("eventlog", ...), Format::csv('|'), _.time)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(Writer::kafka("alerts", ...), Format::json())
16  .run(Backend::flink(...));
```

No close coupling:
Dataflow graphs can be run on
different dataflow systems

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event ⇒ parse(event.line))
8   .filter(error ⇒ error.code == 500)
9   .keyBy(error ⇒ error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) ⇒ {
11    val avg_severity = errors.map(error ⇒ error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data ⇒ data.avg_severity > 2)
15  .sink(...).run(...);
```

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
```

Functional Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
```

Functional Syntax

```
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[Error] = { ... }
5
6 Stream[RawEvent]::source(...)
7   .flatMap(event => parse(event.line))
8   .filter(error => error.code == 500)
9   .keyBy(error => error.zone)
10  .window(Window::sliding(1m, 10m), (errors, zone, win) => {
11    val avg_severity = errors.map(error => error.severity).average();
12    record(zone, avg_severity, win.start, win.end)
13  })
14  .filter(data => data.avg_severity > 2)
15  .sink(...).run(...);
16
17 from event in Stream[RawEvent]::source(...)
18 from error in parse(event.line)
19 where error.code == 500
20 group error.zone
21   over win = Window::sliding(1m, 10m)
22   compute avg_severity = average of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink(...).run(...);
```

Functional Syntax

**Semantically
Equivalent**

Relational Syntax

Optimisations


AquaLang: Optimisation via Equality Saturation

AquaLang: Optimisation via Equality Saturation

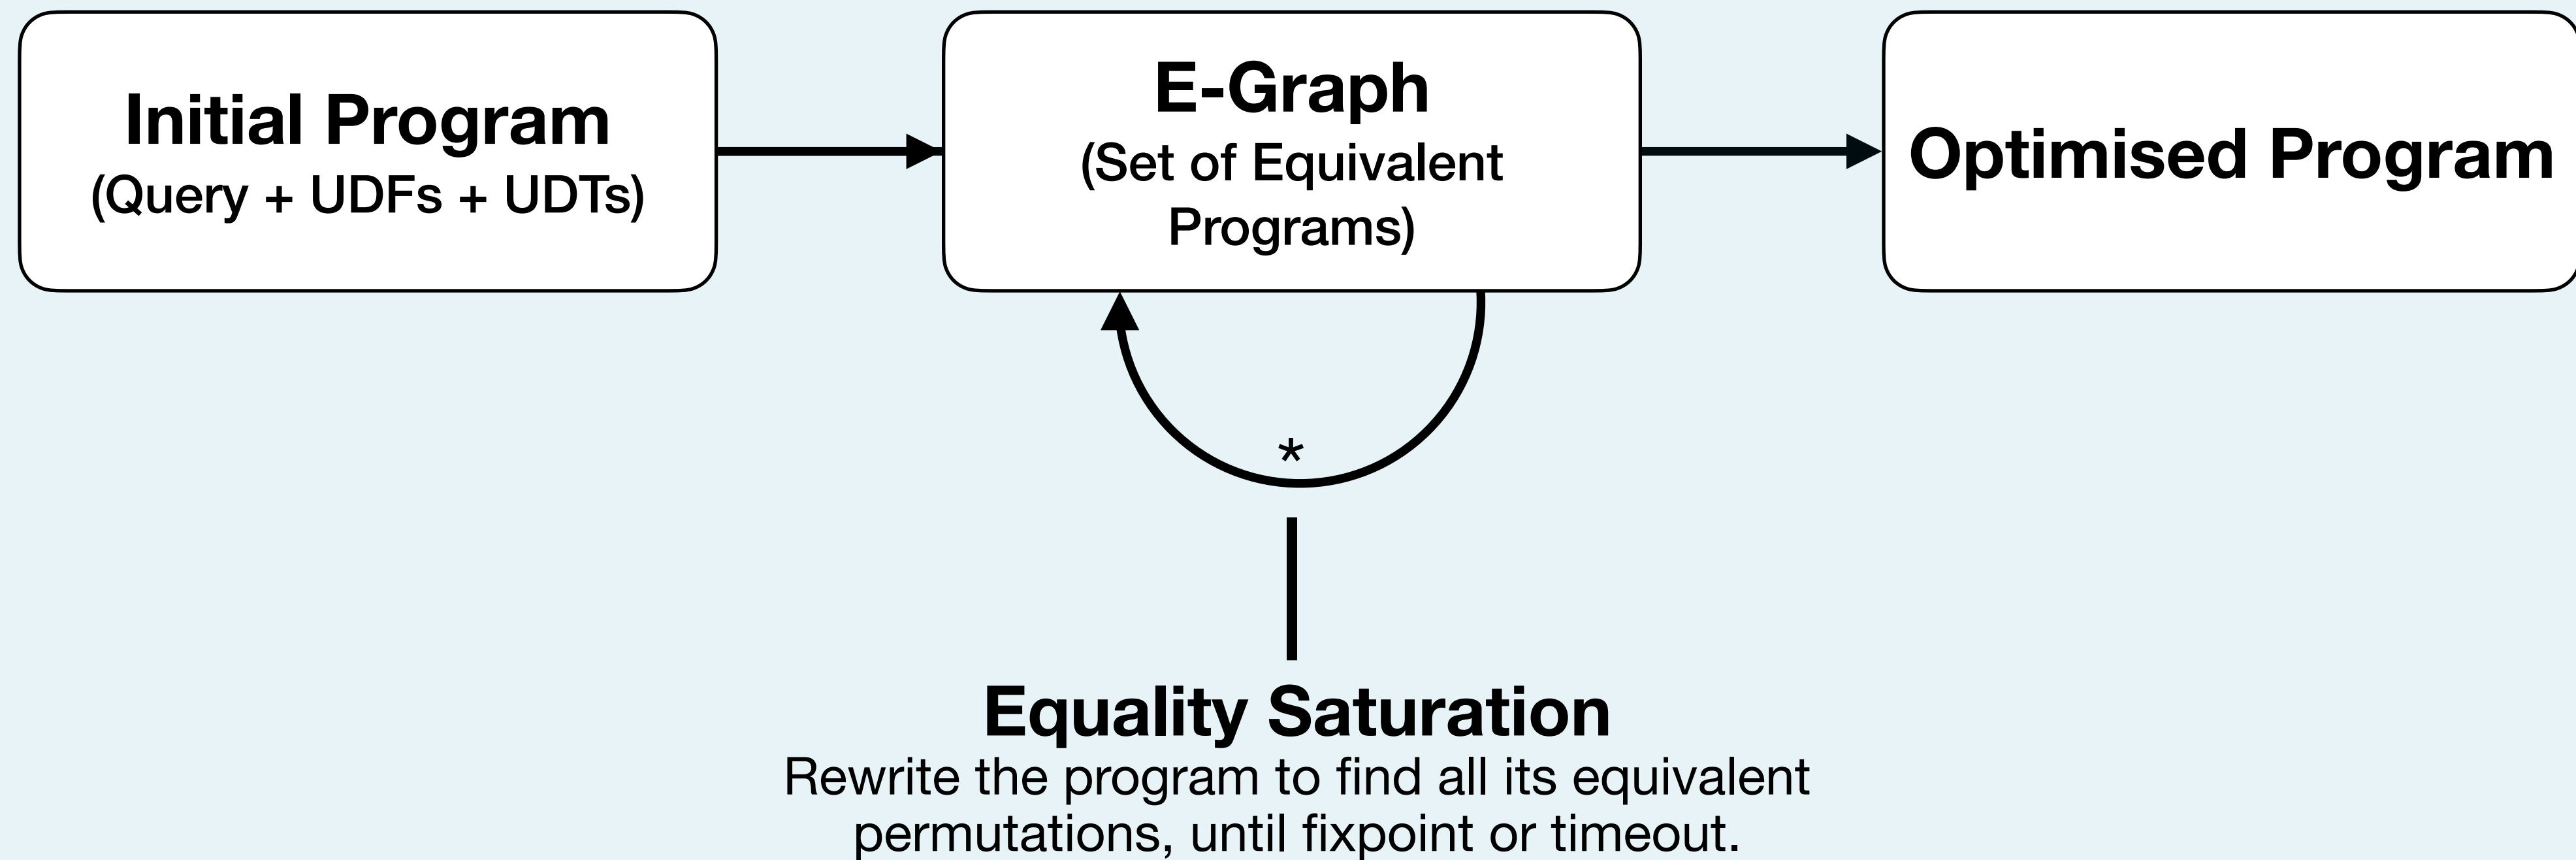
Initial Program
(Query + UDFs + UDTs)

Optimised Program

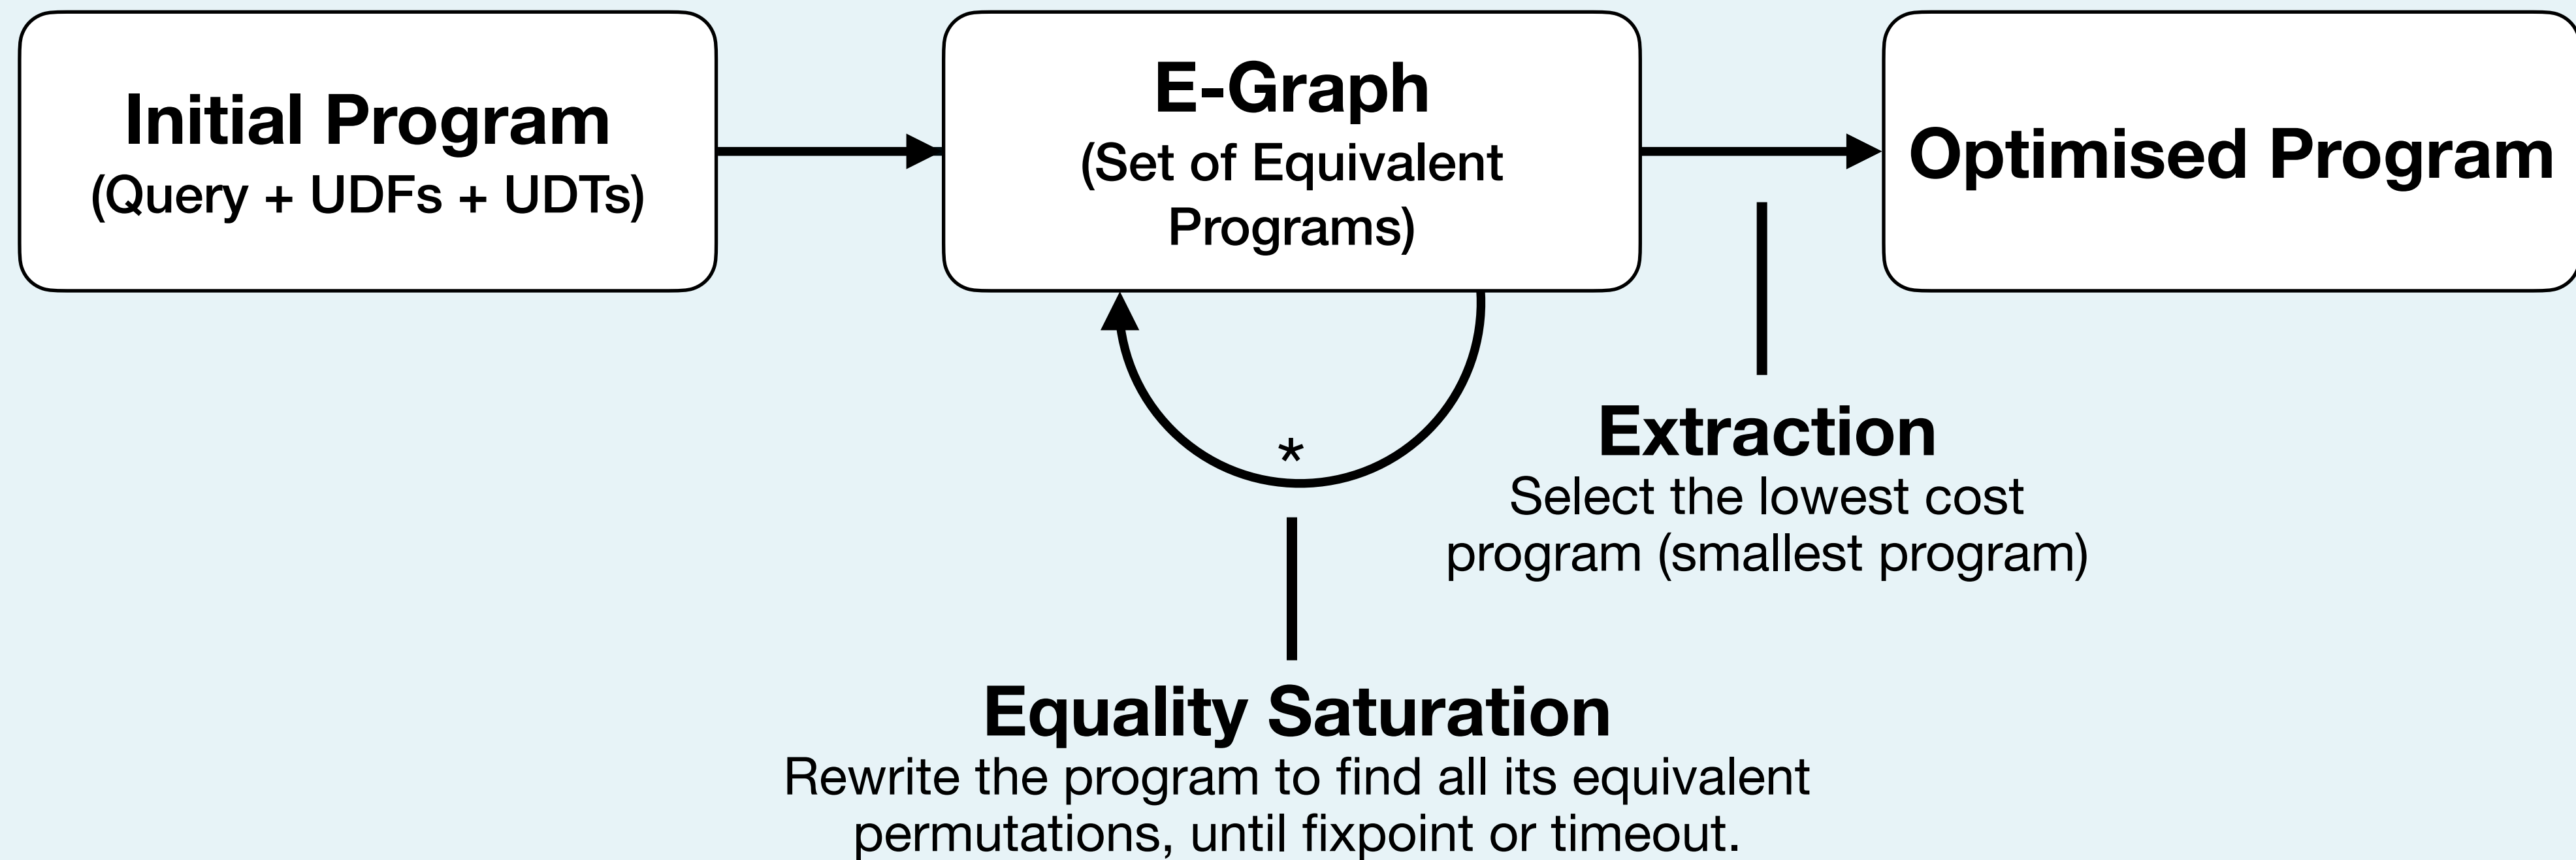
AquaLang: Optimisation via Equality Saturation



AquaLang: Optimisation via Equality Saturation



AquaLang: Optimisation via Equality Saturation



AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[ErrorEvent] = { ... }
5
6 from event in Stream[RawEvent]::source(...)
7 from error in parse(event.line)
8 where error.code = 500
9 group zone = error.zone
10   over win = Window::sliding(1min, 10min)
11   compute avg_severity = avg of error.severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[ErrorEvent] = { ... }
5
6 from event in Stream[RawEvent]::source(...)
7 from error in parse(event.line)
8 where error.code = 500
9 group zone = error.zone
10   over win = Window::sliding(1min, 10min)
11   compute avg_severity = avg of error.severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[RawError] = { ... }
5
6 from event in Stream[RawEvent]::source(...)
7 from error in {
8     val parts = event.line.split(" ");
9     val eventType = parts[0];
10    if eventType == "ERROR" {
11        val code = parts[1].parse();
12        val severity = parts[2].parse();
13        val zone = parts[3];
14        val server = parts[4];
15        val message = parts[5];
16        Some(ErrorEvent(code, severity, zone, server, message))
17    } else { None }
18 }
19 where error.code = 500
20 group zone = error.zone
21     over win = Window::sliding(1min, 10min)
22     compute avg_severity = avg of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def parse(line: String): Option[RawError] = { ... }
5
6 from event in Stream[RawEvent]::source(...)
7 from error in {
8     val parts = event.line.split(" ");
9     val eventType = parts[0];
10    if eventType == "ERROR" {
11        val code = parts[1].parse();
12        val severity = parts[2].parse();
13        val zone = parts[3];
14        val server = parts[4];
15        val message = parts[5];
16        Some(ErrorEvent(code, severity, zone, server, message))
17    } else { None }
18 }
19 where error.code = 500
20 group zone = error.zone
21     over win = Window::sliding(1min, 10min)
22     compute avg_severity = avg of error.severity
23 where avg_severity > 2
24 select zone, avg_severity, win.start, win.end
25 into sink( ... );
```


AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 from error in {
6     val parts = event.line.split(" ");
7     val eventType = parts[0];
8     if eventType == "ERROR" {
9         val code = parts[1].parse();
10        val severity = parts[2].parse();
11        val zone = parts[3];
12        val server = parts[4];
13        val message = parts[5];
14        Some(ErrorEvent(code, severity, zone, server, message))
15    } else { None }
16 }
17 where error.code = 500
18 group zone = error.zone
19     over win = Window::sliding(1min, 10min)
20     compute avg_severity = avg of error.severity
21 where avg_severity > 2
22 select zone, avg_severity, win.start, win.end
23 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 from error in {
6     val parts = event.line.split(" ");
7     val eventType = parts[0];
8     if eventType == "ERROR" {
9         val code = parts[1].parse();
10        val severity = parts[2].parse();
11        val zone = parts[3];
12        val server = parts[4];
13        val message = parts[5];
14        Some(ErrorEvent(code, severity, zone, server, message))
15    } else { None }
16 }
17 where error.code = 500
18 group zone = error.zone
19     over win = Window::sliding(1min, 10min)
20     compute avg_severity = avg of error.severity
21 where avg_severity > 2
22 select zone, avg_severity, win.start, win.end
23 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 from error in {
8     if eventType == "ERROR" {
9         val code = parts[1].parse();
10        val severity = parts[2].parse();
11        val zone = parts[3];
12        val server = parts[4];
13        val message = parts[5];
14        Some(ErrorEvent(code, severity, zone, server, message))
15    } else { None }
16 }
17 where error.code == 500
18 group zone = error.zone
19     over win = Window::sliding(1min, 10min)
20     compute avg_severity = avg of error.severity
21 where avg_severity > 2
22 select zone, avg_severity, win.start, win.end
23 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 from error in {
8     if eventType == "ERROR" {
9         val code = parts[1].parse();
10        val severity = parts[2].parse();
11        val zone = parts[3];
12        val server = parts[4];
13        val message = parts[5];
14        Some(ErrorEvent(code, severity, zone, server, message))
15    } else { None }
16 }
17 where error.code == 500
18 group zone = error.zone
19     over win = Window::sliding(1min, 10min)
20     compute avg_severity = avg of error.severity
21 where avg_severity > 2
22 select zone, avg_severity, win.start, win.end
23 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 from error in {
9     val code = parts[1].parse();
10    val severity = parts[2].parse();
11    val zone = parts[3];
12    val server = parts[4];
13    val message = parts[5];
14    Some(ErrorEvent(code, severity, zone, server, message))
15 }
16 where error.code = 500
17 group zone = error.zone
18     over win = Window::sliding(1min, 10min)
19     compute avg_severity = avg of error.severity
20 where avg_severity > 2
21 select zone, avg_severity, win.start, win.end
22 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 from error in {
9     val code = parts[1].parse();
10    val severity = parts[2].parse();
11    val zone = parts[3];
12    val server = parts[4];
13    val message = parts[5];
14    Some(ErrorEvent(code, severity, zone, server, message))
15 }
16 where error.code = 500
17 group zone = error.zone
18     over win = Window::sliding(1min, 10min)
19     compute avg_severity = avg of error.severity
20 where avg_severity > 2
21 select zone, avg_severity, win.start, win.end
22 into sink( ... );
```


AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 from error in {
14     Some(ErrorEvent(code, severity, zone, server, message))
15 }
16 where error.code = 500
17 group zone = error.zone
18     over win = Window::sliding(1min, 10min)
19     compute avg_severity = avg of error.severity
20 where avg_severity > 2
21 select zone, avg_severity, win.start, win.end
22 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 from error in {
14     Some(ErrorEvent(code, severity, zone, server, message))
15 }
16 where error.code = 500
17 group zone = error.zone
18     over win = Window::sliding(1min, 10min)
19     compute avg_severity = avg of error.severity
20 where avg_severity > 2
21 select zone, avg_severity, win.start, win.end
22 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 from error in Some(ErrorEvent(code, severity, zone, server, message))
14 where error.code = 500
15 group zone = error.zone
16     over win = Window::sliding(1min, 10min)
17     compute avg_severity = avg of error.severity
18 where avg_severity > 2
19 select zone, avg_severity, win.start, win.end
20 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 from error in Some(ErrorEvent(code, severity, zone, server, message))
14 where error.code = 500
15 group zone = error.zone
16     over win = Window::sliding(1min, 10min)
17     compute avg_severity = avg of error.severity
18 where avg_severity > 2
19 select zone, avg_severity, win.start, win.end
20 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 select error = ErrorEvent(code, severity, zone, server, message)
14 where error.code = 500
15 group zone = error.zone
16     over win = Window::sliding(1min, 10min)
17     compute avg_severity = avg of error.severity
18 where avg_severity > 2
19 select zone, avg_severity, win.start, win.end
20 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 select error = ErrorEvent(code, severity, zone, server, message)
14 where error.code = 500
15 group zone = error.zone
16     over win = Window::sliding(1min, 10min)
17     compute avg_severity = avg of error.severity
18 where avg_severity > 2
19 select zone, avg_severity, win.start, win.end
20 into sink( ... );
```


AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 select code, severity, zone, server, message
14 where code = 500
15 group zone
16     over win = Window::sliding(1min, 10min)
17     compute avg_severity = avg of severity
18 where avg_severity > 2
19 select zone, avg_severity, win.start, win.end
20 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 val server = parts[4]
12 val message = parts[5]
13 select code, severity, zone, server, message
14 where code = 500
15 group zone
16     over win = Window::sliding(1min, 10min)
17     compute avg_severity = avg of severity
18 where avg_severity > 2
19 select zone, avg_severity, win.start, win.end
20 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11
12
13 select code, severity, zone
14 where code = 500
15 group zone
16     over win = Window::sliding(1min, 10min)
17     compute avg_severity = avg of severity
18 where avg_severity > 2
19 select zone, avg_severity, win.start, win.end
20 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 select code, severity, zone
12 where code = 500
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 val severity = parts[2].parse()
10 val zone = parts[3]
11 select code, severity, zone
12 where code = 500
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 where code = 500
10 val severity = parts[2].parse()
11 val zone = parts[3]
12 select code, severity, zone
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```


AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 where code = 500
10 val severity = parts[2].parse()
11 val zone = parts[3]
12 select code, severity, zone
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 where code = 500
10 val severity = parts[2].parse()
11 val zone = parts[3]
12 select severity, zone
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 where code = 500
10 val severity = parts[2].parse()
11 val zone = parts[3]
12 select severity, zone
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 val eventType = parts[0]
7 where eventType = "ERROR"
8 val code = parts[1].parse()
9 where code = 500
10 val severity = parts[2].parse()
11 val zone = parts[3]
12 select severity, zone
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6
7 where parts[0] = "ERROR"
8
9 where parts[1].parse() = 500
10
11
12 select severity = parts[2].parse(), zone = parts[3]
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6
7 where parts[0] = "ERROR"
8
9 where parts[1].parse() = 500
10
11
12 select severity = parts[2].parse(), zone = parts[3]
13 group zone
14     over win = Window::sliding(1min, 10min)
15     compute avg_severity = avg of severity
16 where avg_severity > 2
17 select zone, avg_severity, win.start, win.end
18 into sink( ... );
```


AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 where parts[0] = "ERROR"
7 where parts[1].parse() = 500
8 select severity = parts[2].parse(), zone = parts[3]
9 group zone
10     over win = Window::sliding(1min, 10min)
11     compute avg_severity = avg of severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 where parts[0] = "ERROR"
7 where parts[1].parse() = 500
8 select severity = parts[2].parse(), zone = parts[3]
9 group zone
10     over win = Window::sliding(1min, 10min)
11     compute avg_severity = avg of severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 where parts[0] = "ERROR"
7 and parts[1].parse() = 500
8 select severity = parts[2].parse(), zone = parts[3]
9 group zone
10 over win = Window::sliding(1min, 10min)
11 compute avg_severity = avg of severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```

AquaLang: Optimisation

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 where parts[0] = "ERROR"
7     and parts[1].parse() = 500
8 select severity = parts[2].parse(), zone = parts[3]
9 group zone
10     over win = Window::sliding(1min, 10min)
11     compute avg_severity = avg of severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```

More features

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 where parts[0] = "ERROR"
7     and parts[1].parse() = 500
8 select severity = parts[2].parse(), zone = parts[3]
9 group zone
10     over win = Window::sliding(1min, 10min)
11     compute avg_severity = avg of severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```


AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 from event in Stream[RawEvent]::source(...)
5 val parts = event.line.split(" ")
6 where parts[0] = "ERROR"
7     and parts[1].parse() = 500
8 select severity = parts[2].parse(), zone = parts[3]
9 group zone
10     over win = Window::sliding(1min, 10min)
11     compute avg_severity = avg of severity
12 where avg_severity > 2
13 select zone, avg_severity, win.start, win.end
14 into sink( ... );
```

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def queryAlerts(input: Stream[RawEvent], errorCode: u8): Stream[_] =
5     from event in input
6     val parts = event.line.split(" ")
7     where parts[0] = "ERROR"
8         and parts[1].parse() = errorCode
9     select severity = parts[2].parse(), zone = parts[3]
10    group zone
11        over win = Window::sliding(1min, 10min)
12        compute avg_severity = avg of severity
13    where avg_severity > 2
14    select zone, avg_severity, win.start, win.end
15    into sink( ... );
```

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def queryAlerts(input: Stream[RawEvent], errorCode: u8): Stream[_] =
5     from event in input
6     val parts = event.line.split(" ")
7     where parts[0] = "ERROR"
8         and parts[1].parse() = errorCode
9     select severity = parts[2].parse(), zone = parts[3]
10    group zone
11        over win = Window::sliding(1min, 10min)
12        compute avg_severity = avg of severity
13    where avg_severity > 2
14    select zone, avg_severity, win.start, win.end
15    into sink( ... );
```

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def queryAlerts(input: Stream[RawEvent], errorCode: u8): Stream[_] =
5   from event in input
6   val parts = event.line.split(" ")
7   where parts[0] = "ERROR"
8         and parts[1].parse() = errorCode
9   select severity = parts[2].parse(), zone = parts[3]
10  group zone
11    over win = Window::sliding(1min, 10min)
12    compute avg_severity = avg of severity
13  where avg_severity > 2
14  select zone, avg_severity, win.start, win.end
15  into sink( ... );
```

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def queryAlerts(input: Stream[RawEvent], errorCode: u8): Stream[_] =
5     from event in input
6     val parts = event.line.split(" ")
7     where parts[0] = "ERROR"
8         and parts[1].parse() = errorCode
9     select severity = parts[2].parse(), zone = parts[3]
10    group zone
11        over win = Window::sliding(1min, 10min)
12        compute avg_severity = avg of severity
13    where avg_severity > 2
14    select zone, avg_severity, win.start, win.end
15    into sink( ... );
```

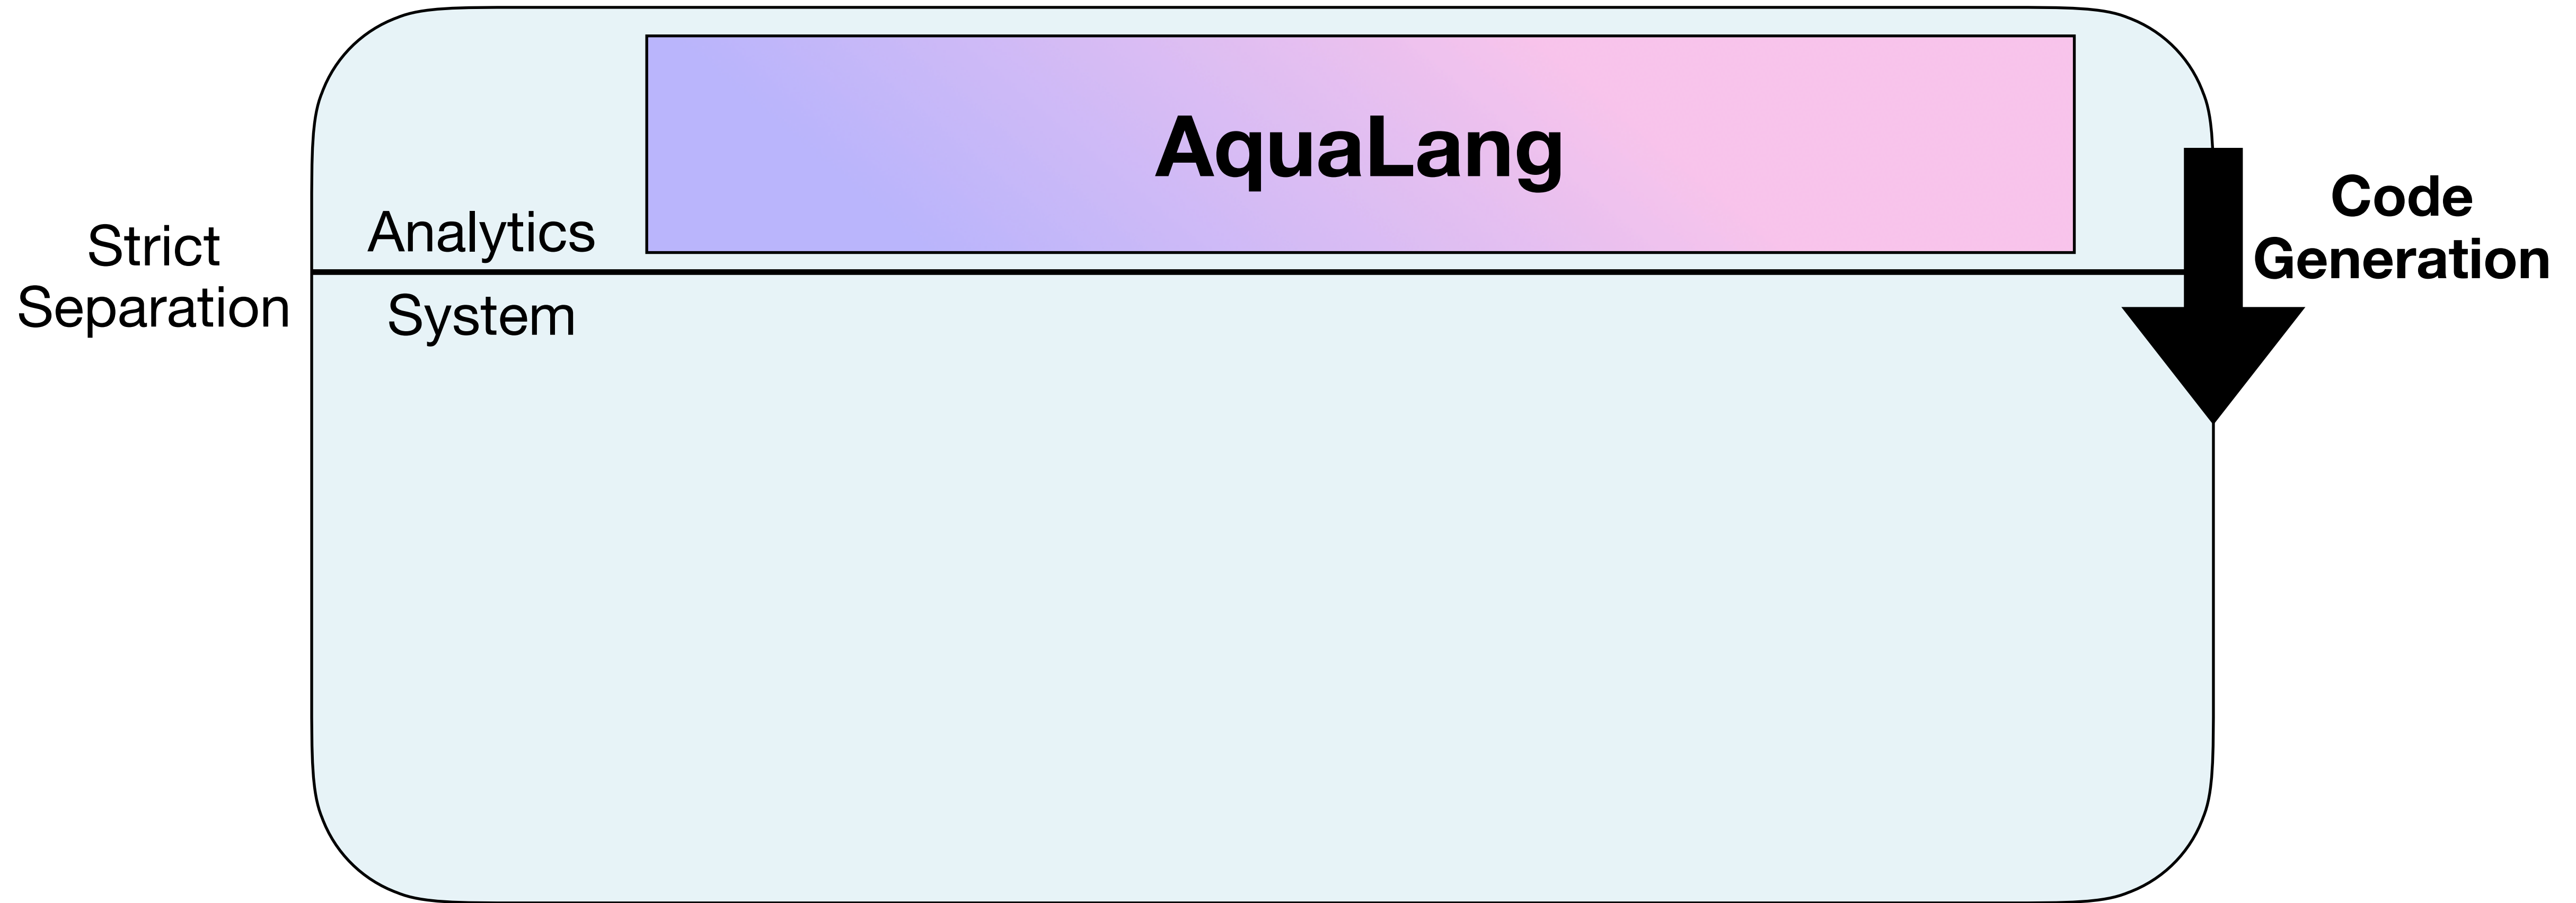
AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def queryAlerts(input: Stream[RawEvent], errorCode: u8): Stream[_] =
5     from event in input
6     val parts = event.line.split(" ")
7     where parts[0] = "ERROR"
8         and parts[1].parse() = errorCode
9     select severity = parts[2].parse(), zone = parts[3]
10    group zone
11        over win = Window::sliding(1min, 10min)
12        compute avg_severity = avg of severity
13    where avg_severity > 2
14    select zone, avg_severity, win.start, win.end
15    into sink( ... );
16
17 from alert in queryAlerts(Stream::source(...), 401)
18 where alert.zone = "Z2"
19 into sink(...).run(...);
```

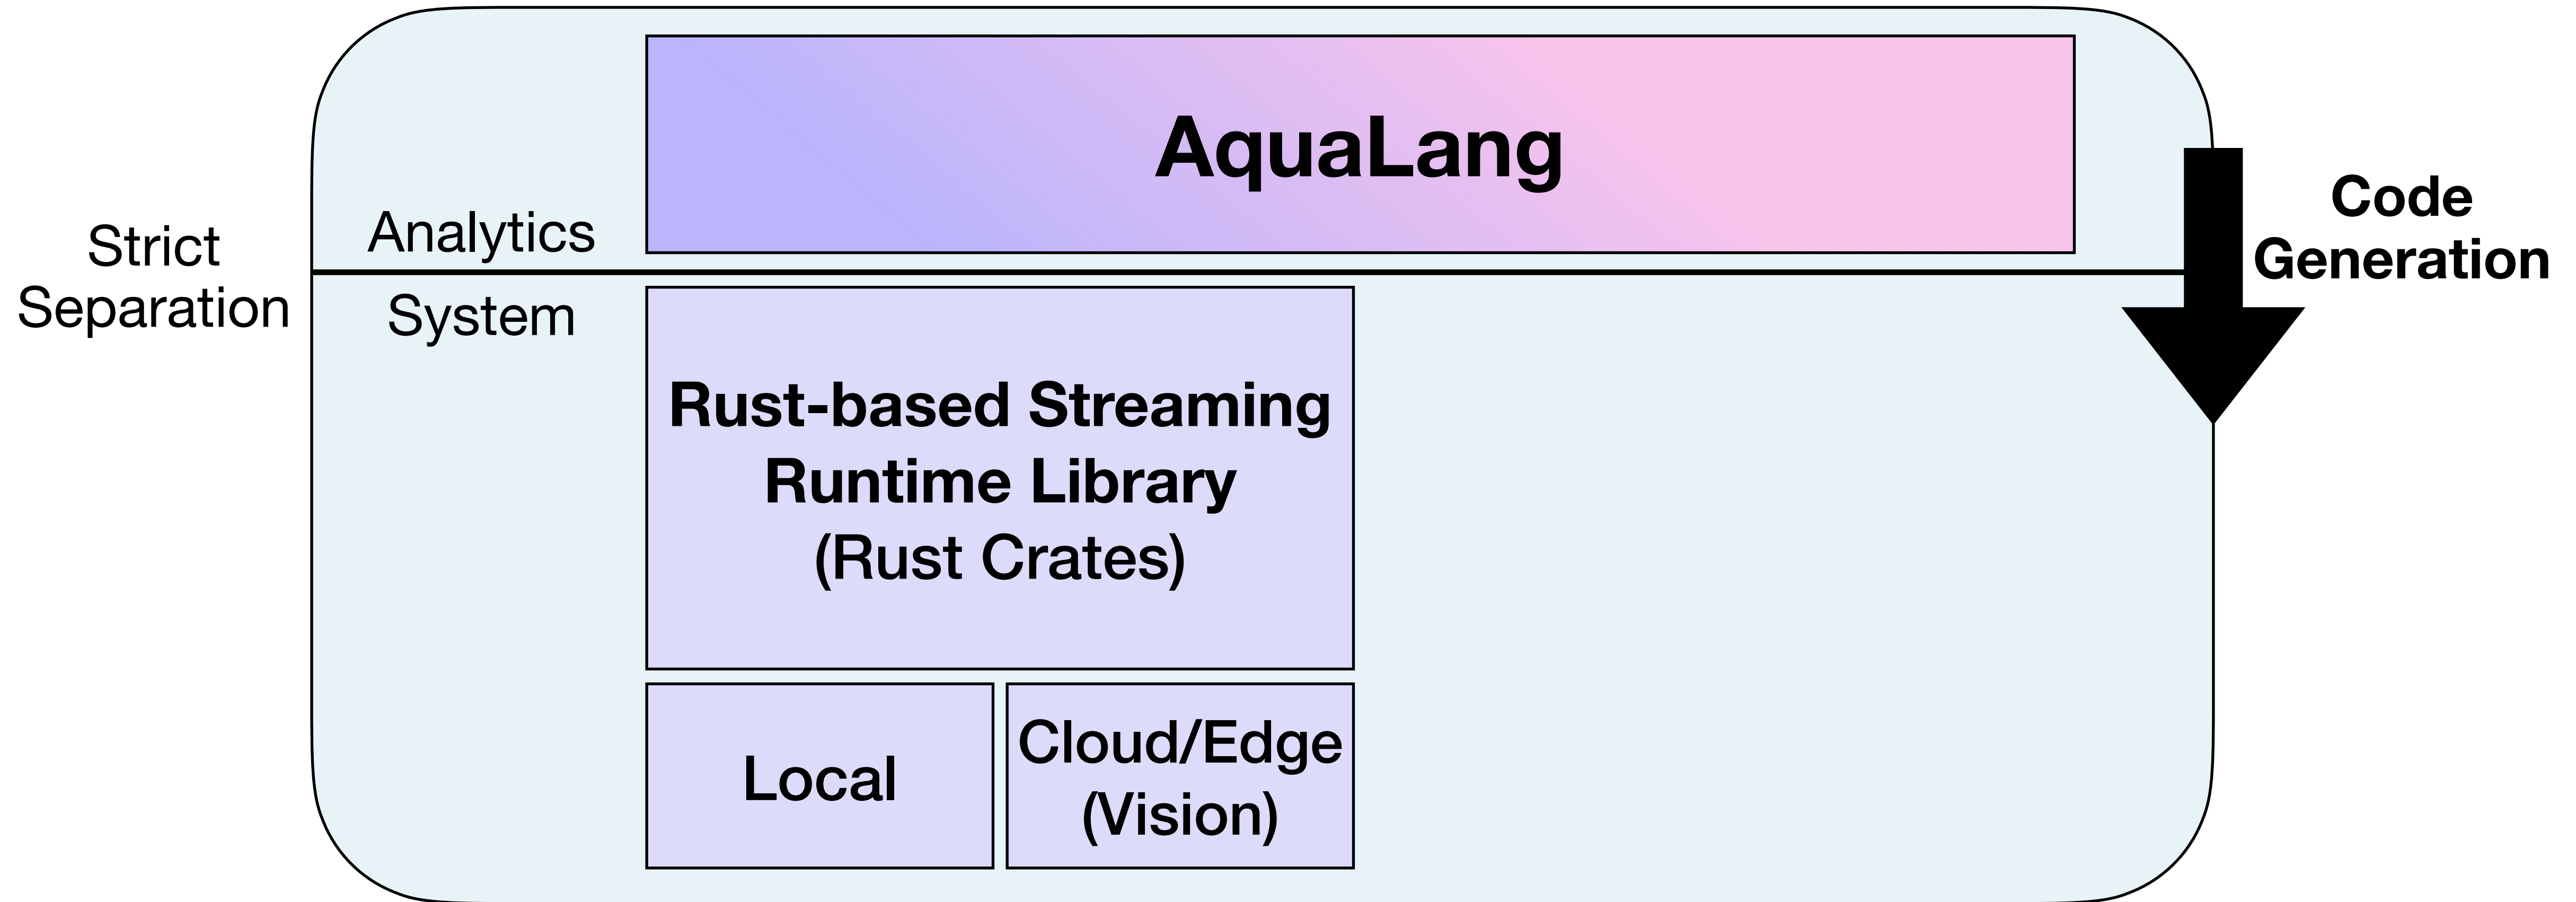

AquaLang: Dataflow Programming

```
1 struct RawEvent(ts:Time, line:String);
2 struct ErrorEvent(code:u8, severity:u8, zone:String, server:String, message:String);
3
4 def queryAlerts(input: Stream[RawEvent], errorCode: u8): Stream[_] =
5     from event in input
6     val parts = event.line.split(" ")
7     where parts[0] = "ERROR"
8         and parts[1].parse() = errorCode
9     select severity = parts[2].parse(), zone = parts[3]
10    group zone
11        over win = Window::sliding(1min, 10min)
12        compute avg_severity = avg of severity
13    where avg_severity > 2
14    select zone, avg_severity, win.start, win.end
15    into sink( ... );
16
17 from alert in queryAlerts(Stream::source(...), 401)
18 where alert.zone = "Z2"
19 into sink(...).run(...);
```

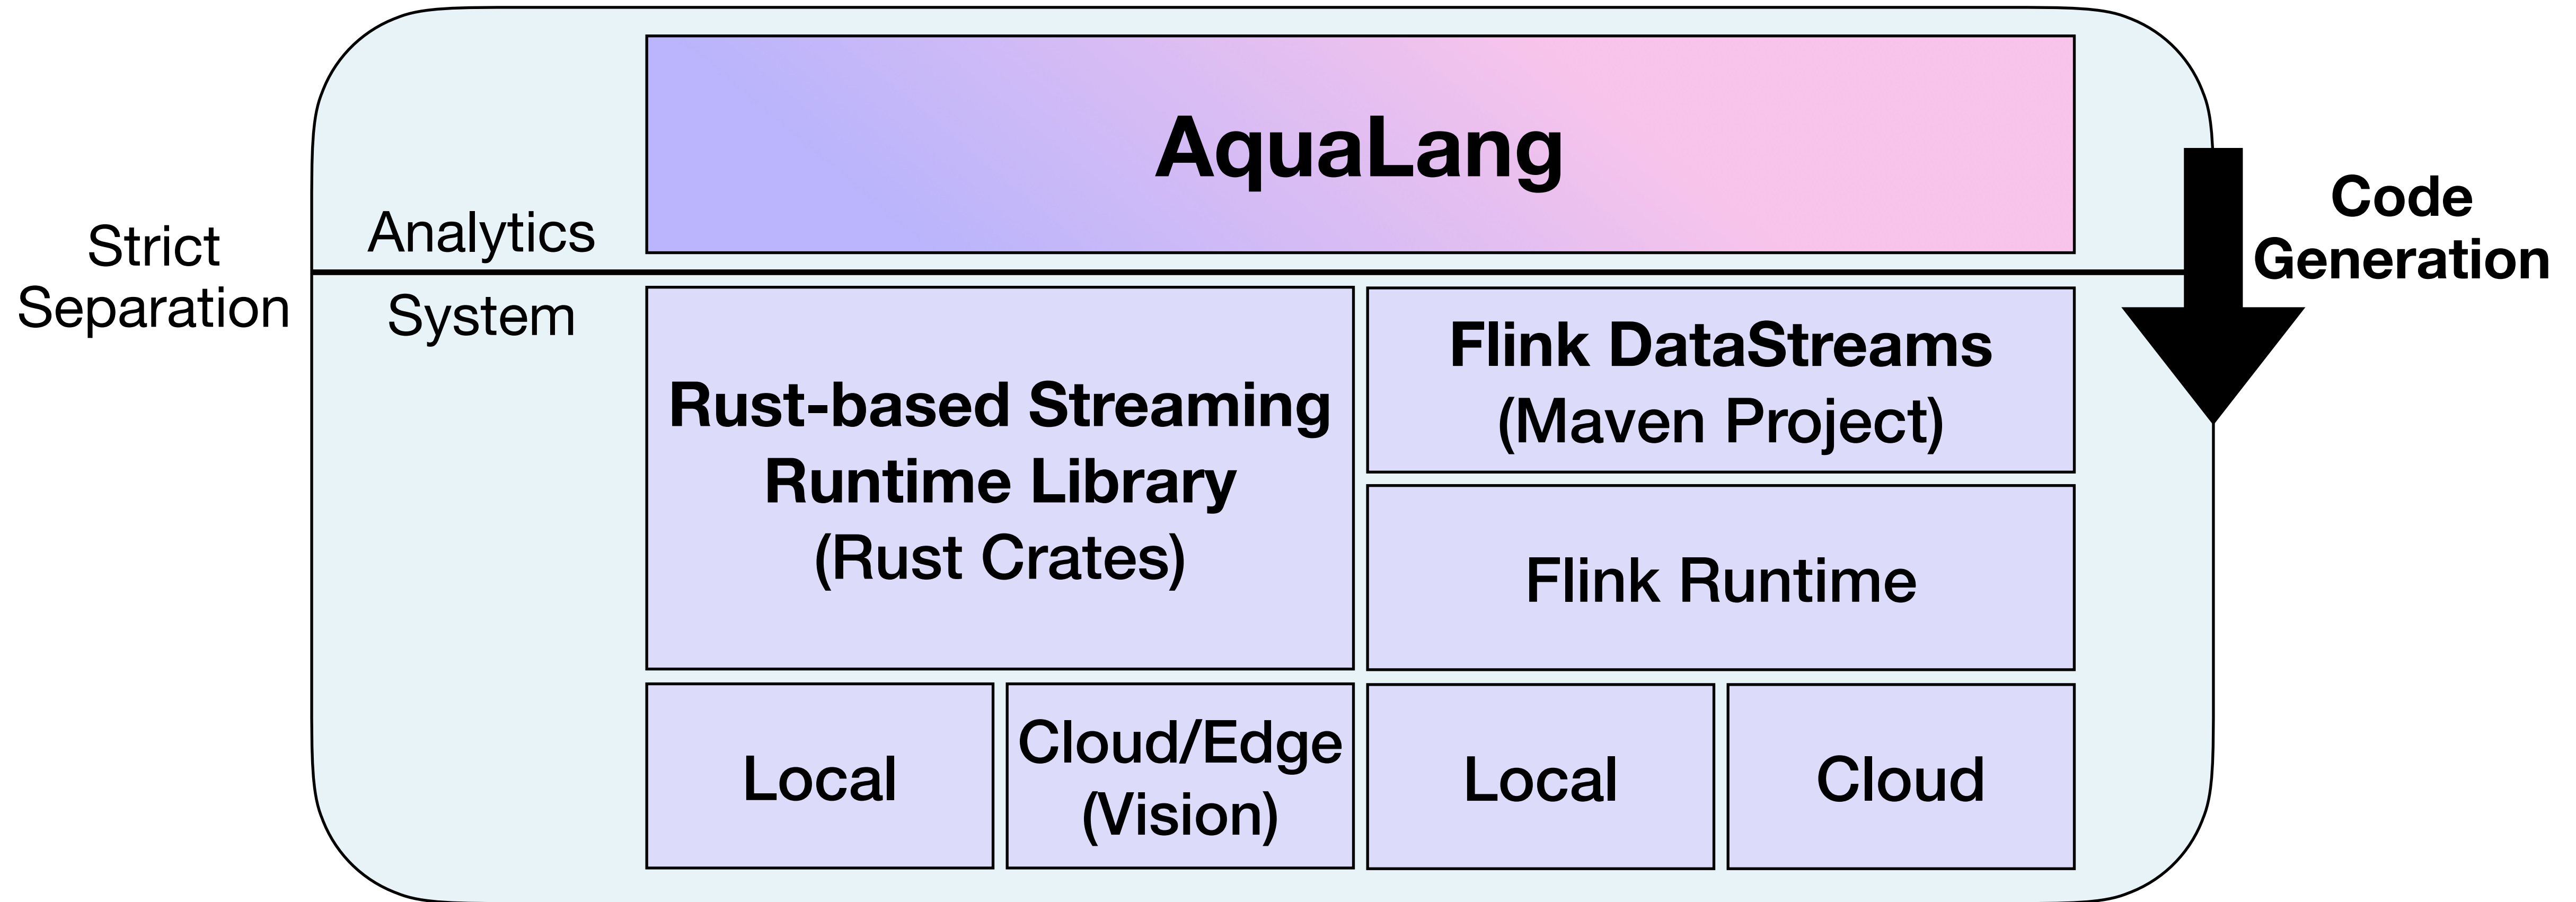
AquaLang: Code Generation



AquaLang: Code Generation



AquaLang: Code Generation



Conclusion

- **AquaLang** is a programming language that targets streaming dataflow systems
- **AquaLang** is designed to be compositional, safe, fast, and easy to use
- We are aiming to release a beta version later this year
- If you find **AquaLang** interesting, please get in touch! :)
 - **Website:** <https://aqua-language.github.io/>
 - **GitHub:** <https://github.com/aqua-language/aqua>
 - **Paper:** AquaLang - A Dataflow Programming Language (DEBS'24)