



Stream All the Things!

Architectures for Data that Never Ends

Dean Wampler, Ph.D.
dean@deanwampler.com
[@deanwampler](https://twitter.com/deanwampler)

Free as in 🍺

go.lightbend.com/fast-data-architectures-for-streaming-applications-oreilly-2nd-edition

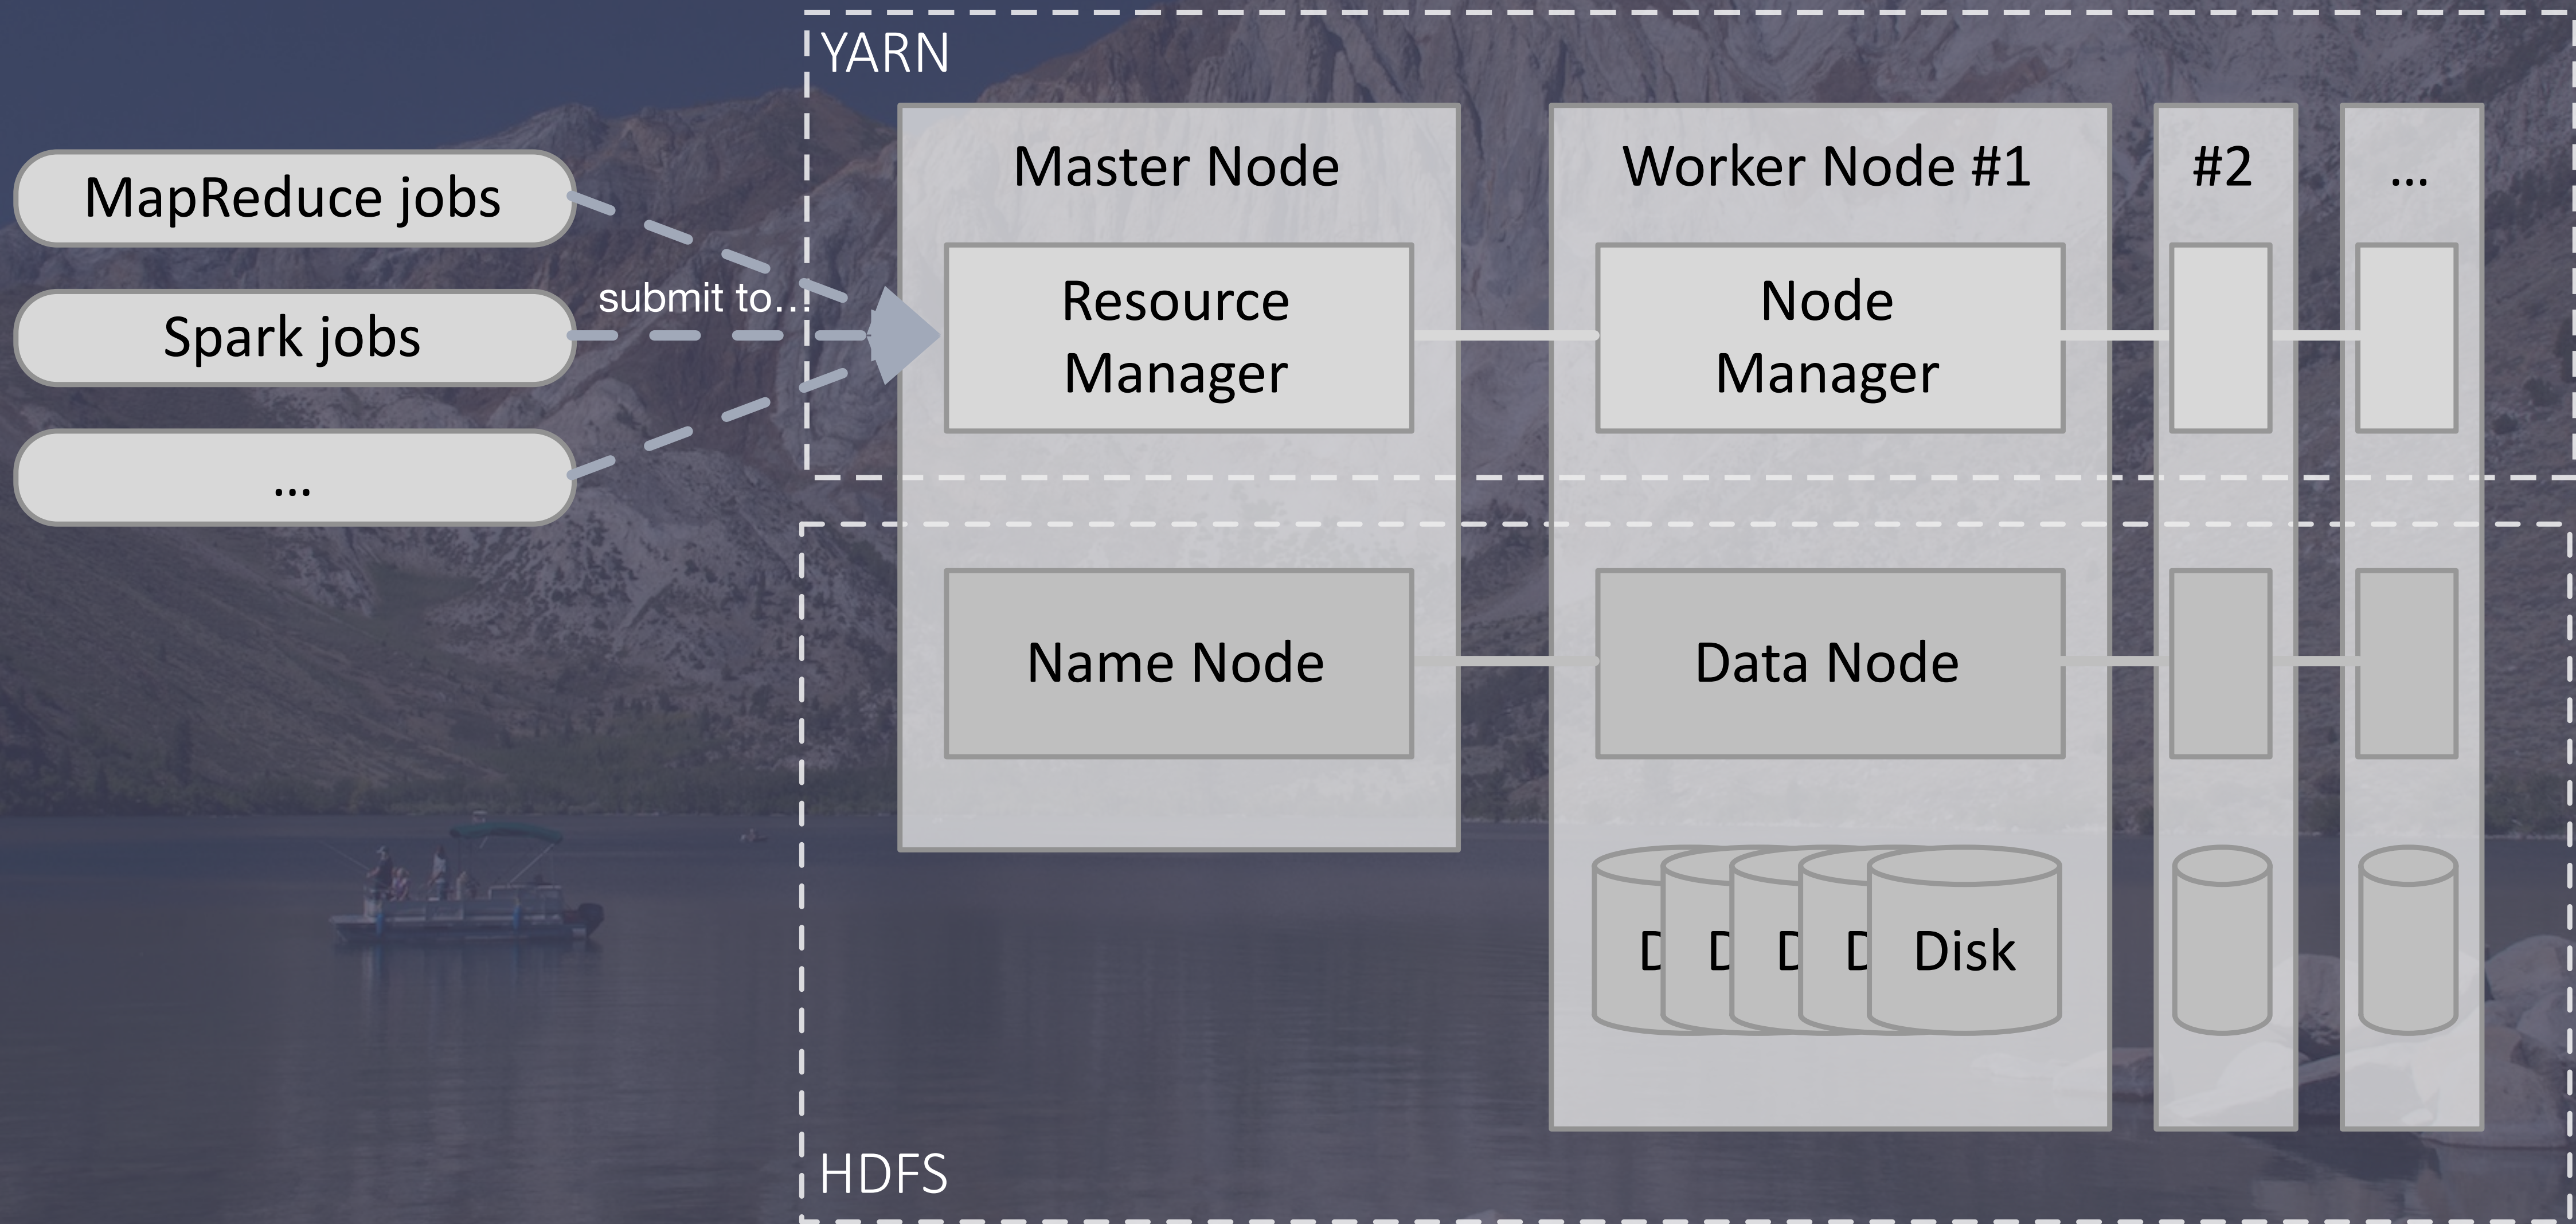


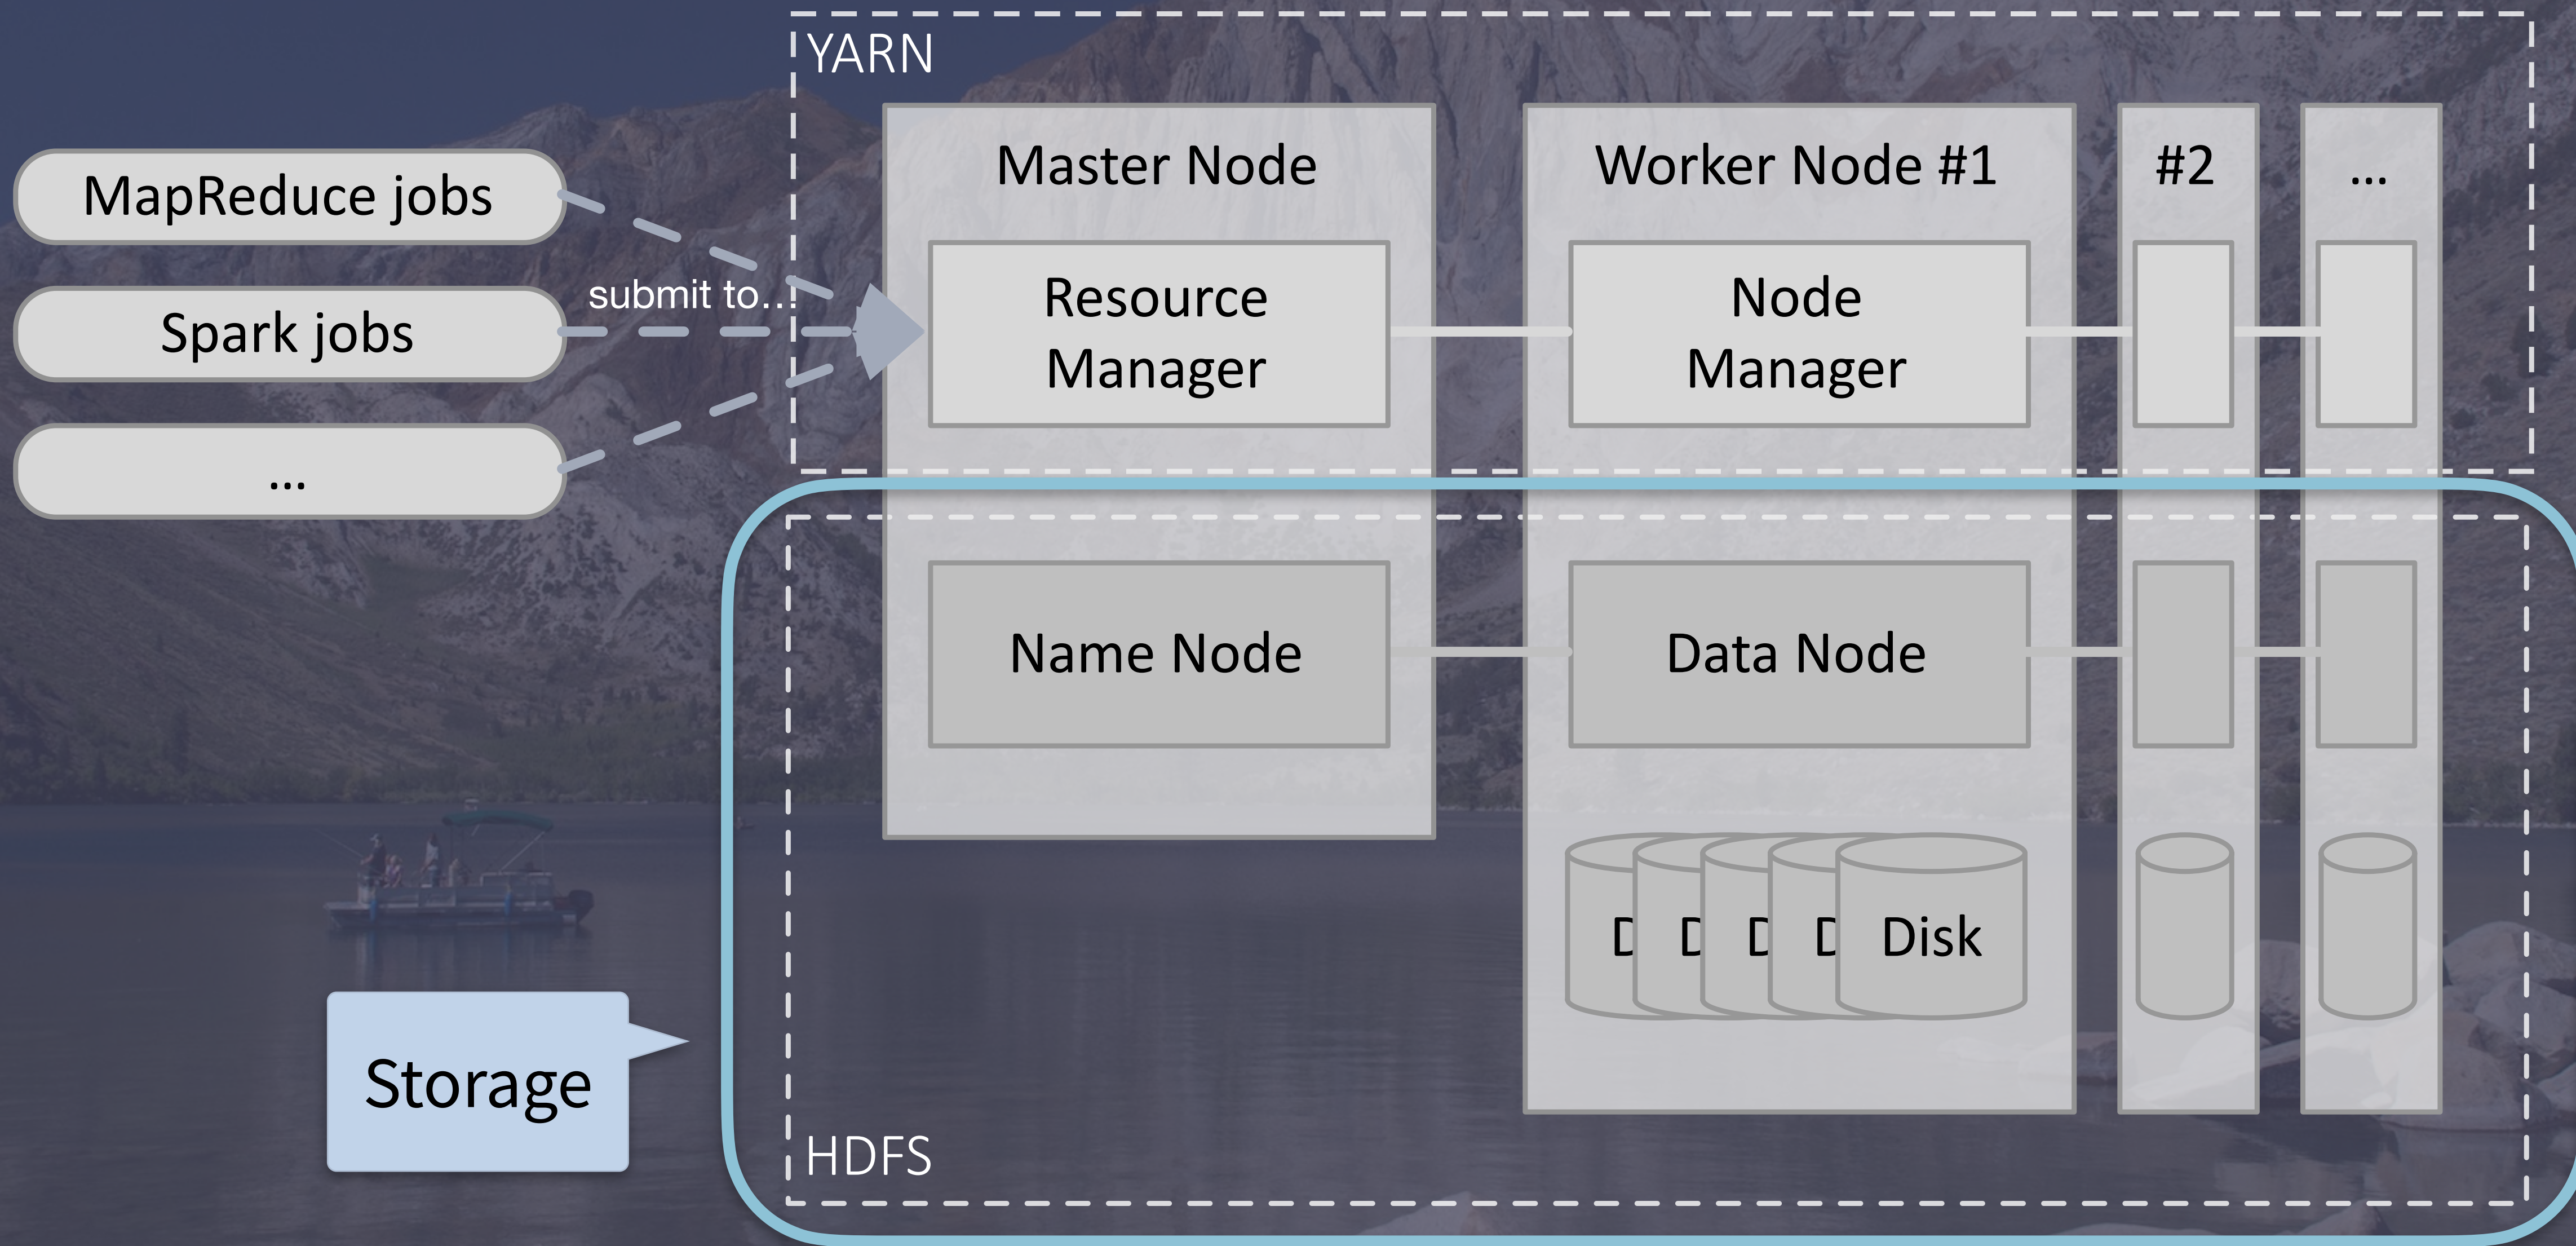


Streaming
in Context...

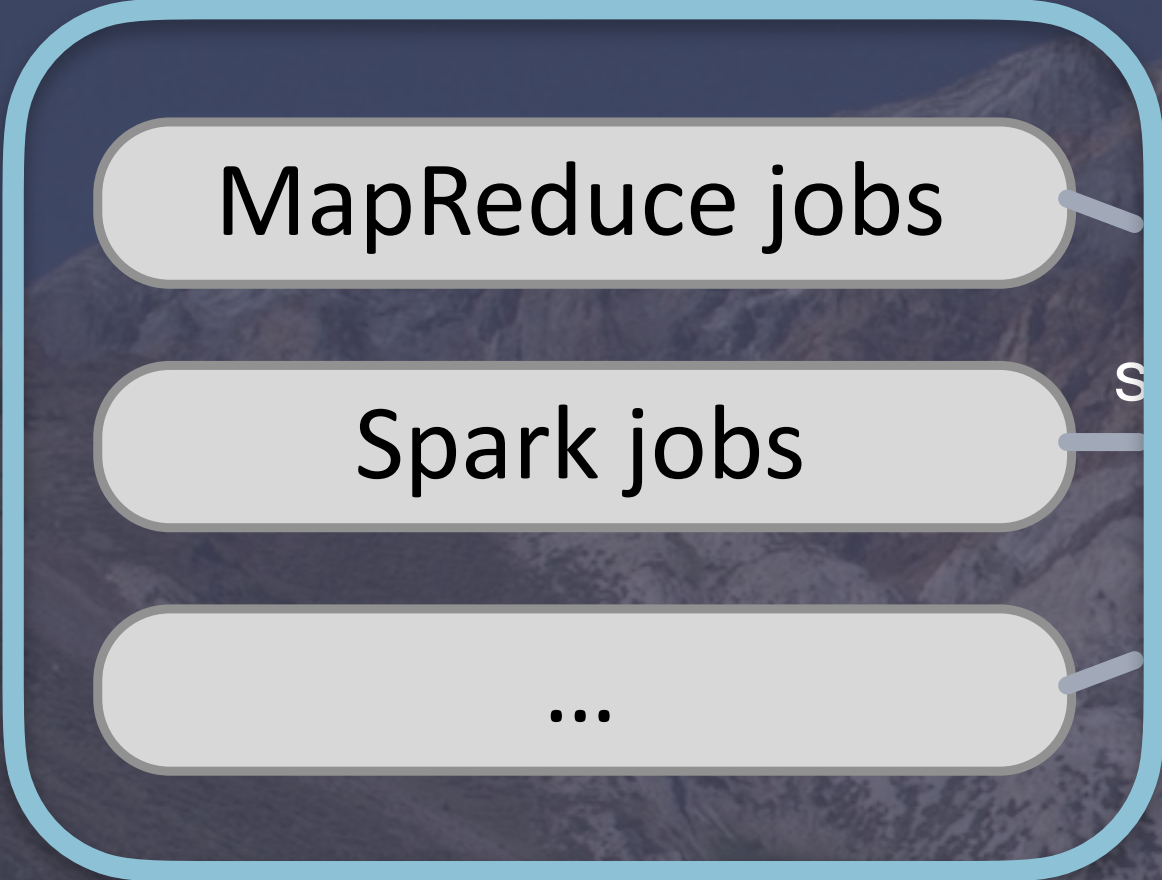


Hadoop: Classic Batch Architecture





Compute



Resource Management

MapReduce jobs

Spark jobs

...

submit to...

YARN

Master Node

Resource Manager

Name Node

Worker Node #1

Node Manager

Data Node

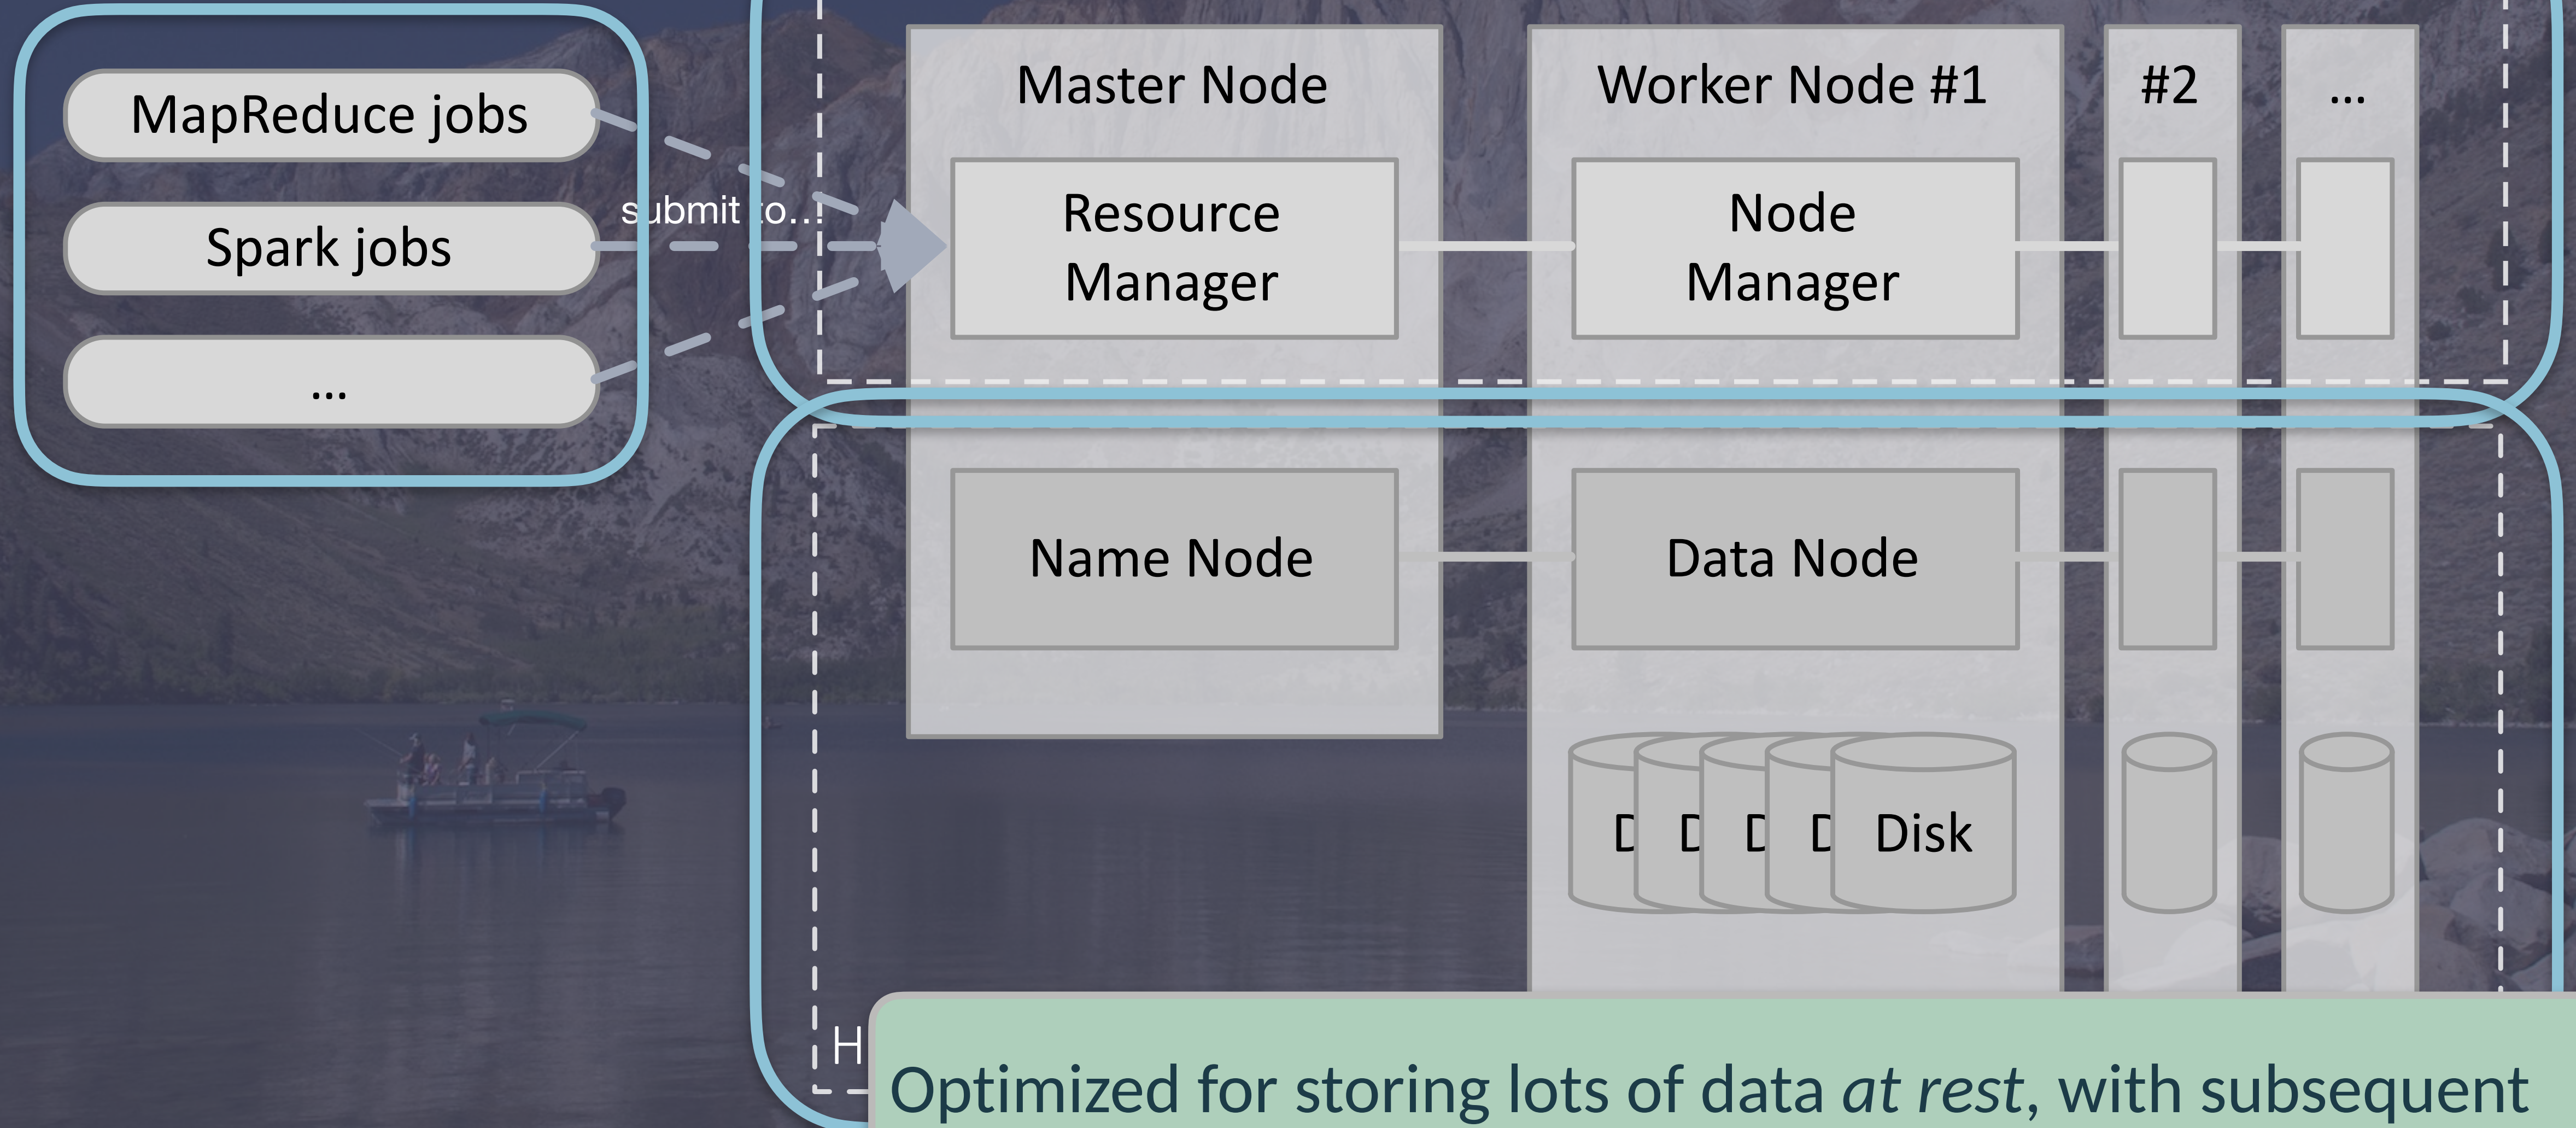
Disk

#2

...

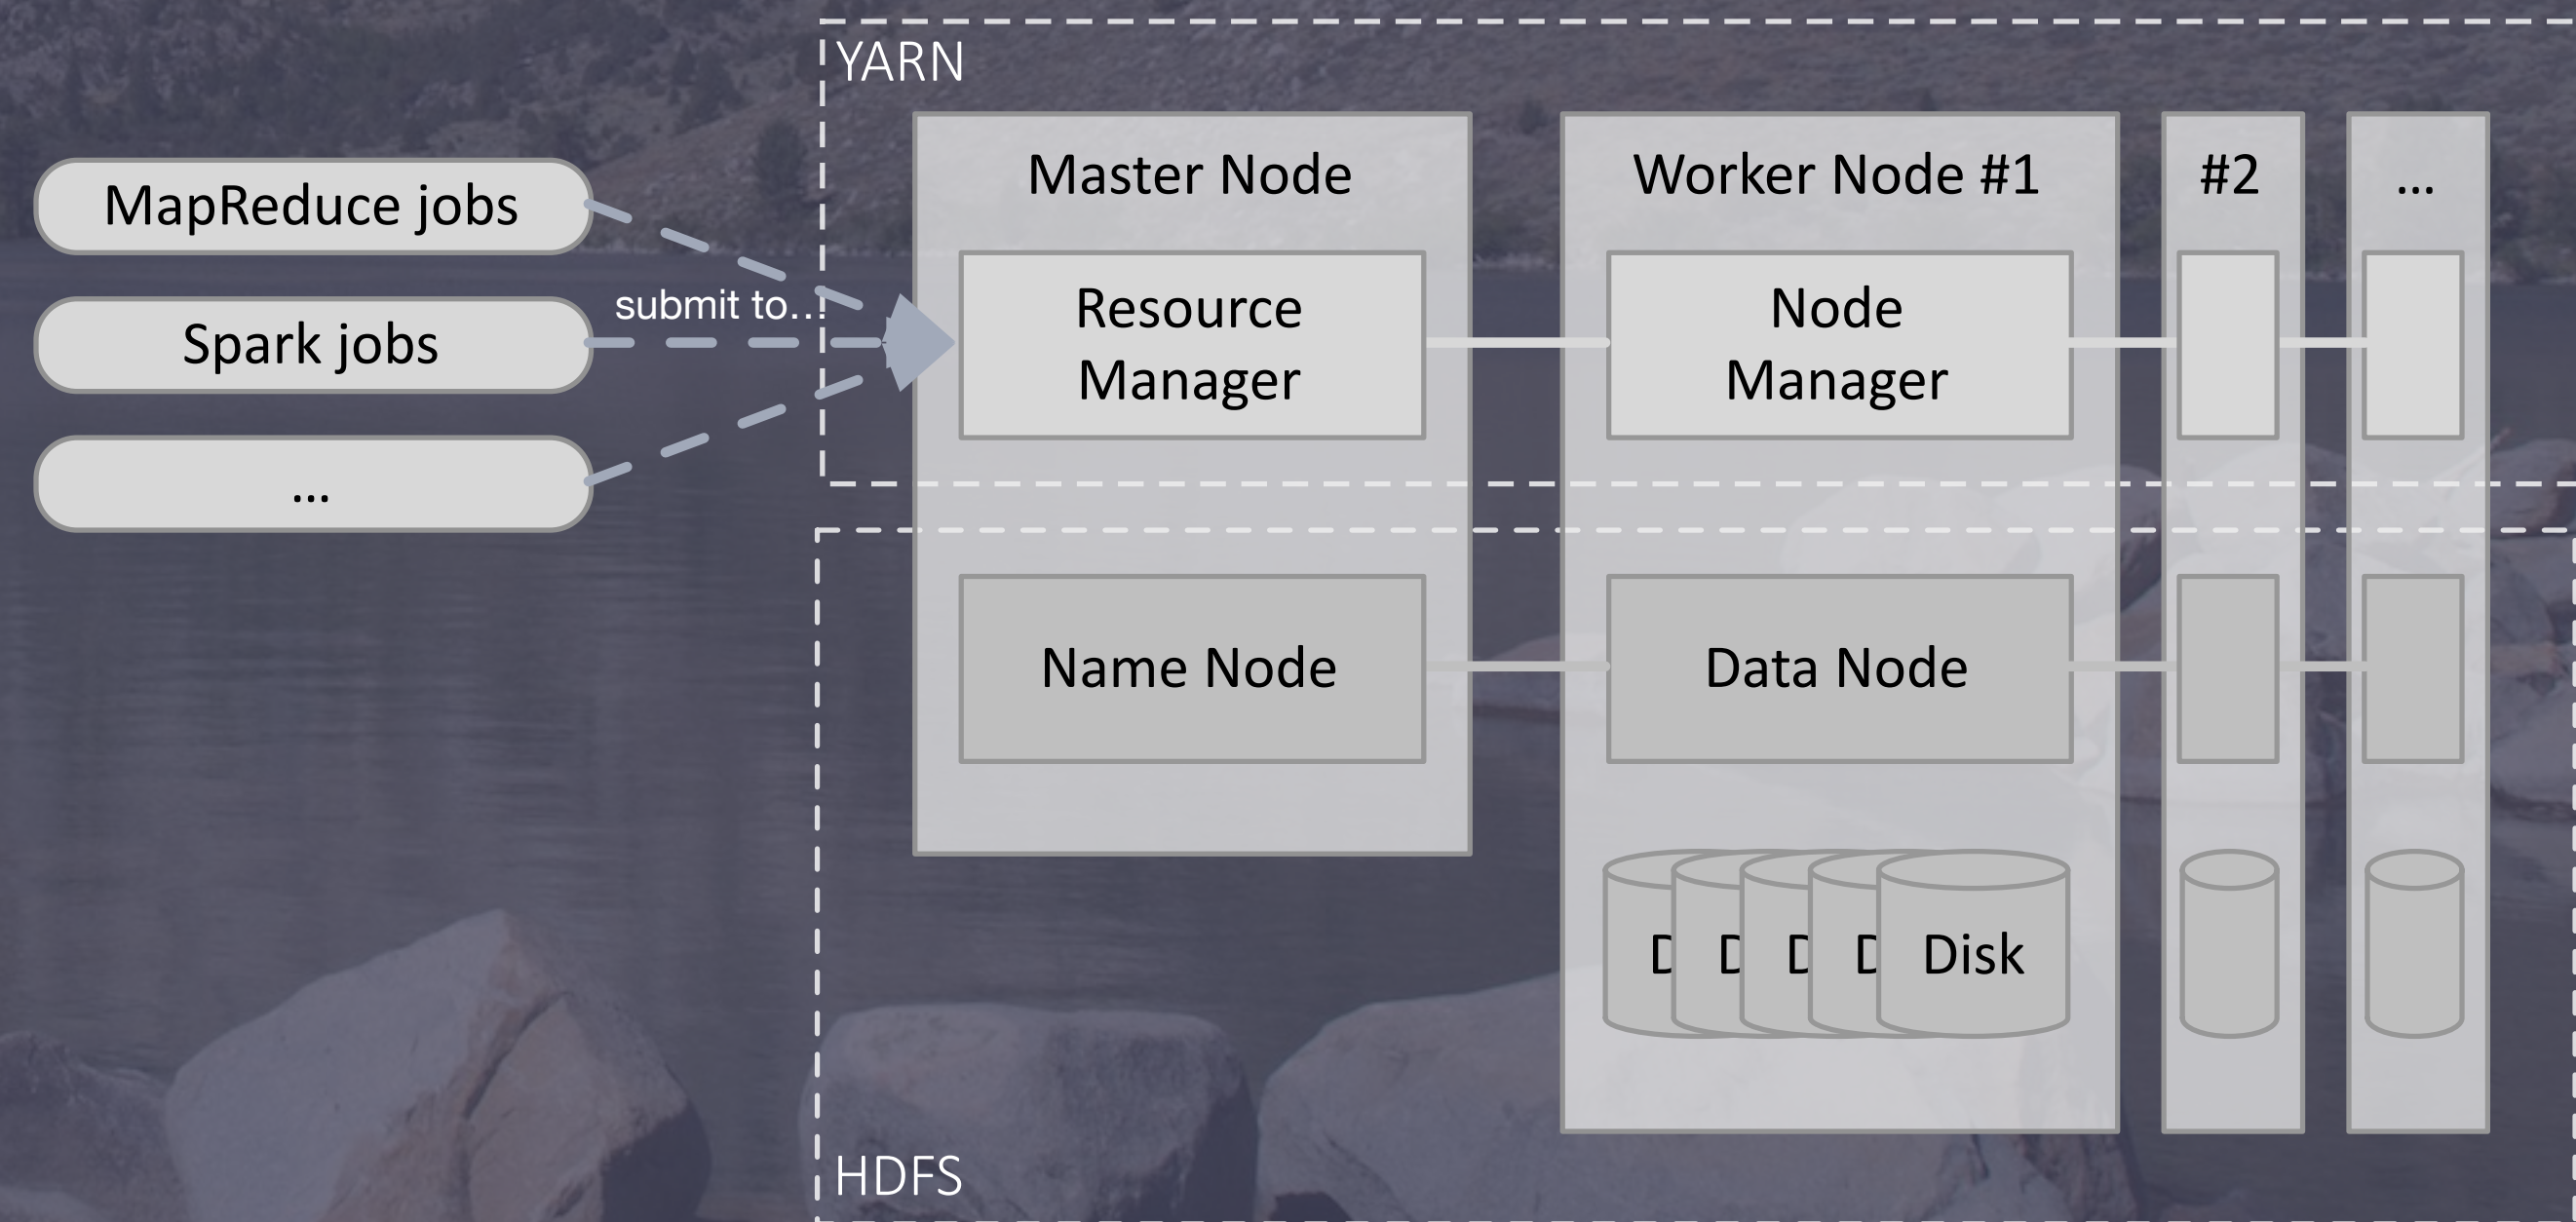
HDFS

Database
Deconstructed!

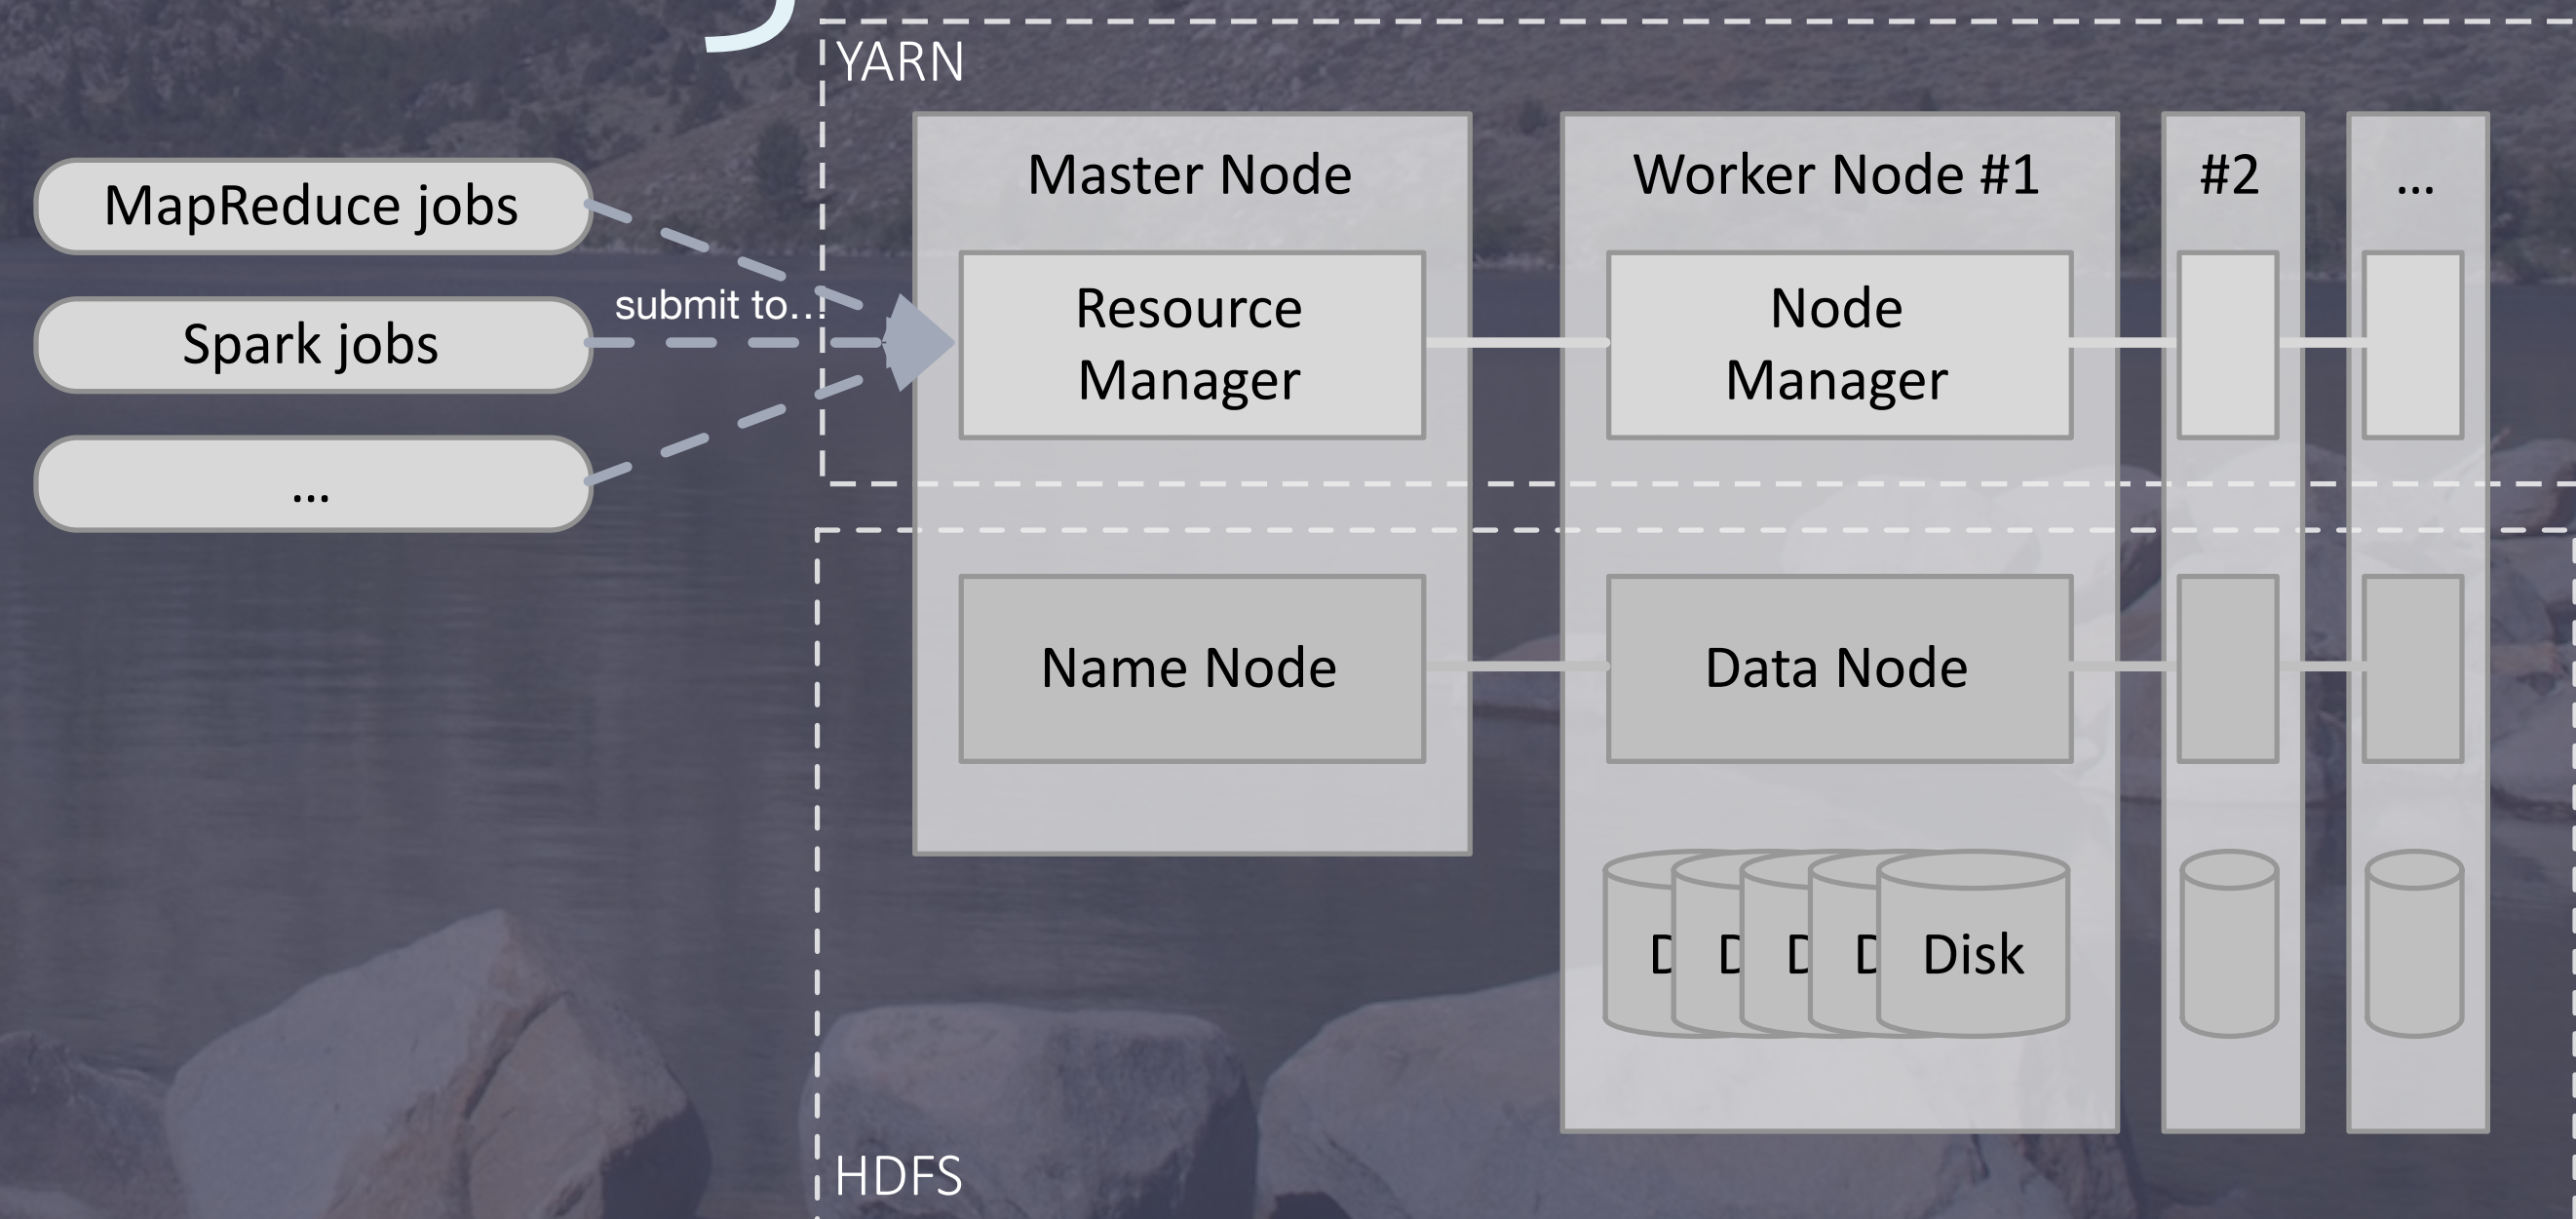


Optimized for storing lots of data *at rest*, with subsequent processing, but not optimized for data *in motion*.

- Characteristics
 - Batch and interactive queries
 - Massive storage - HDFS is the data “backplane”
- Integrate jobs through HDFS
- Multiuser jobs

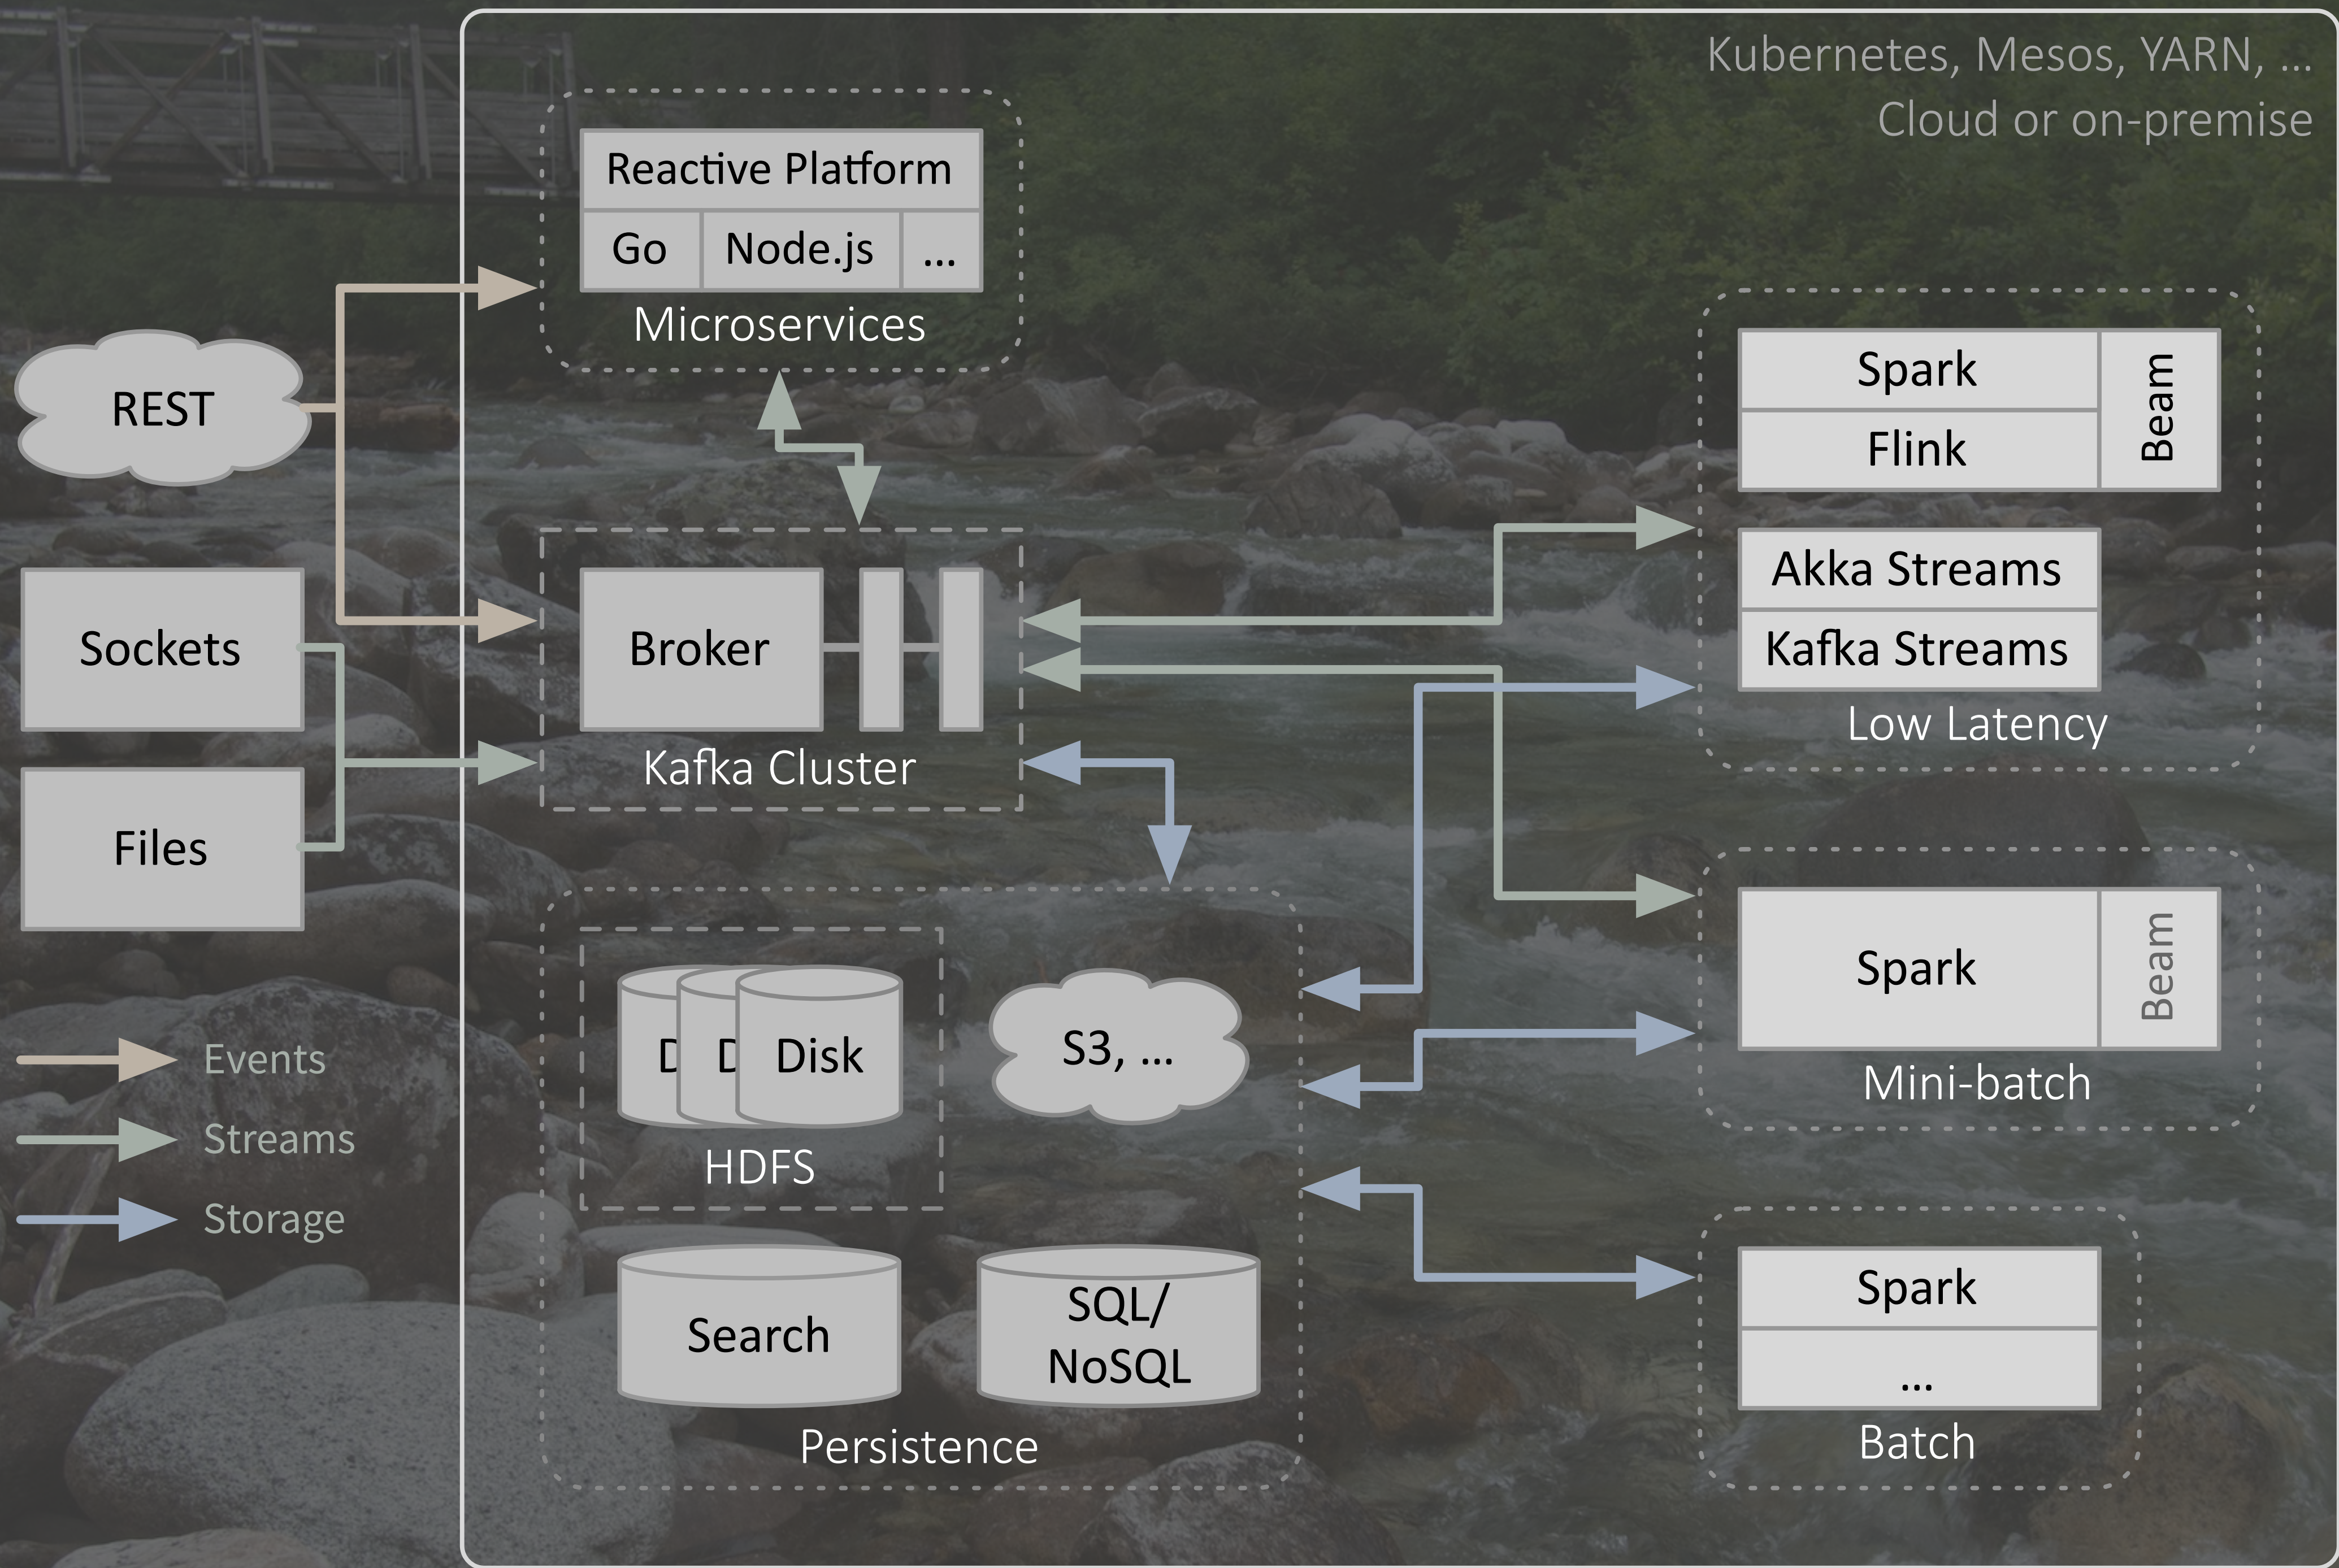


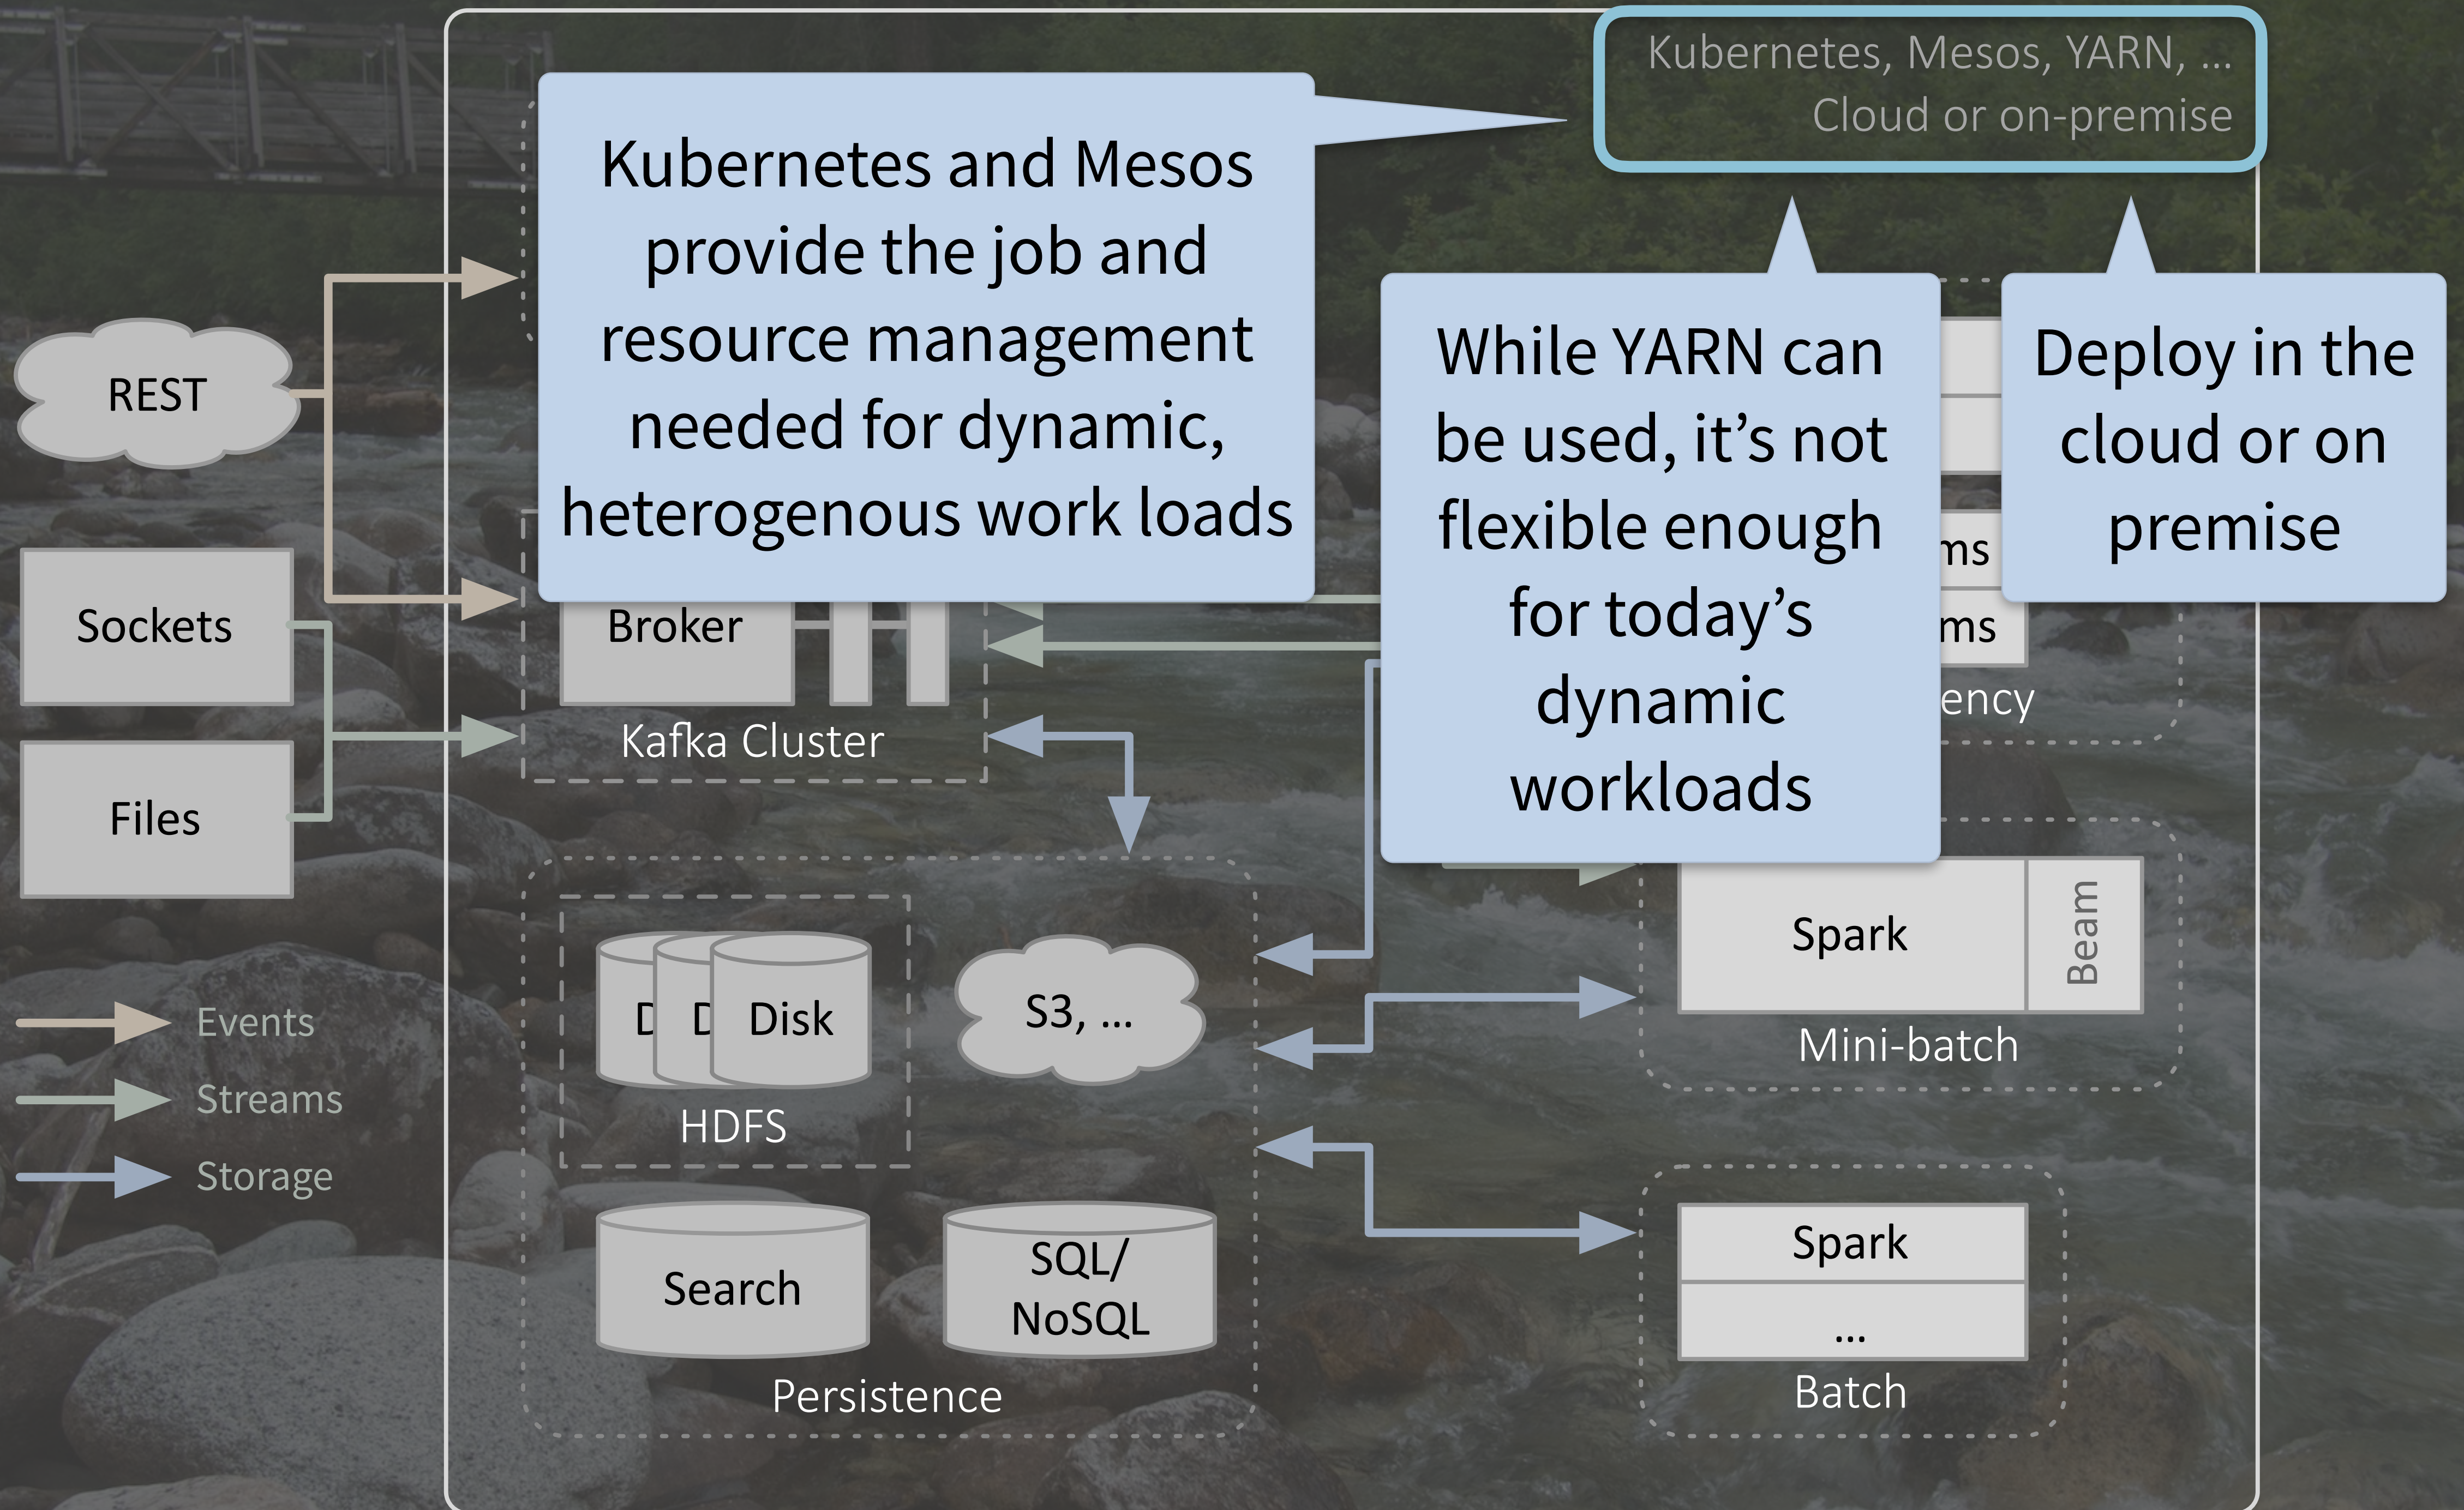
- Use Cases
 - Data warehouse replacement
 - Interactive exploration
 - Offline ML model training

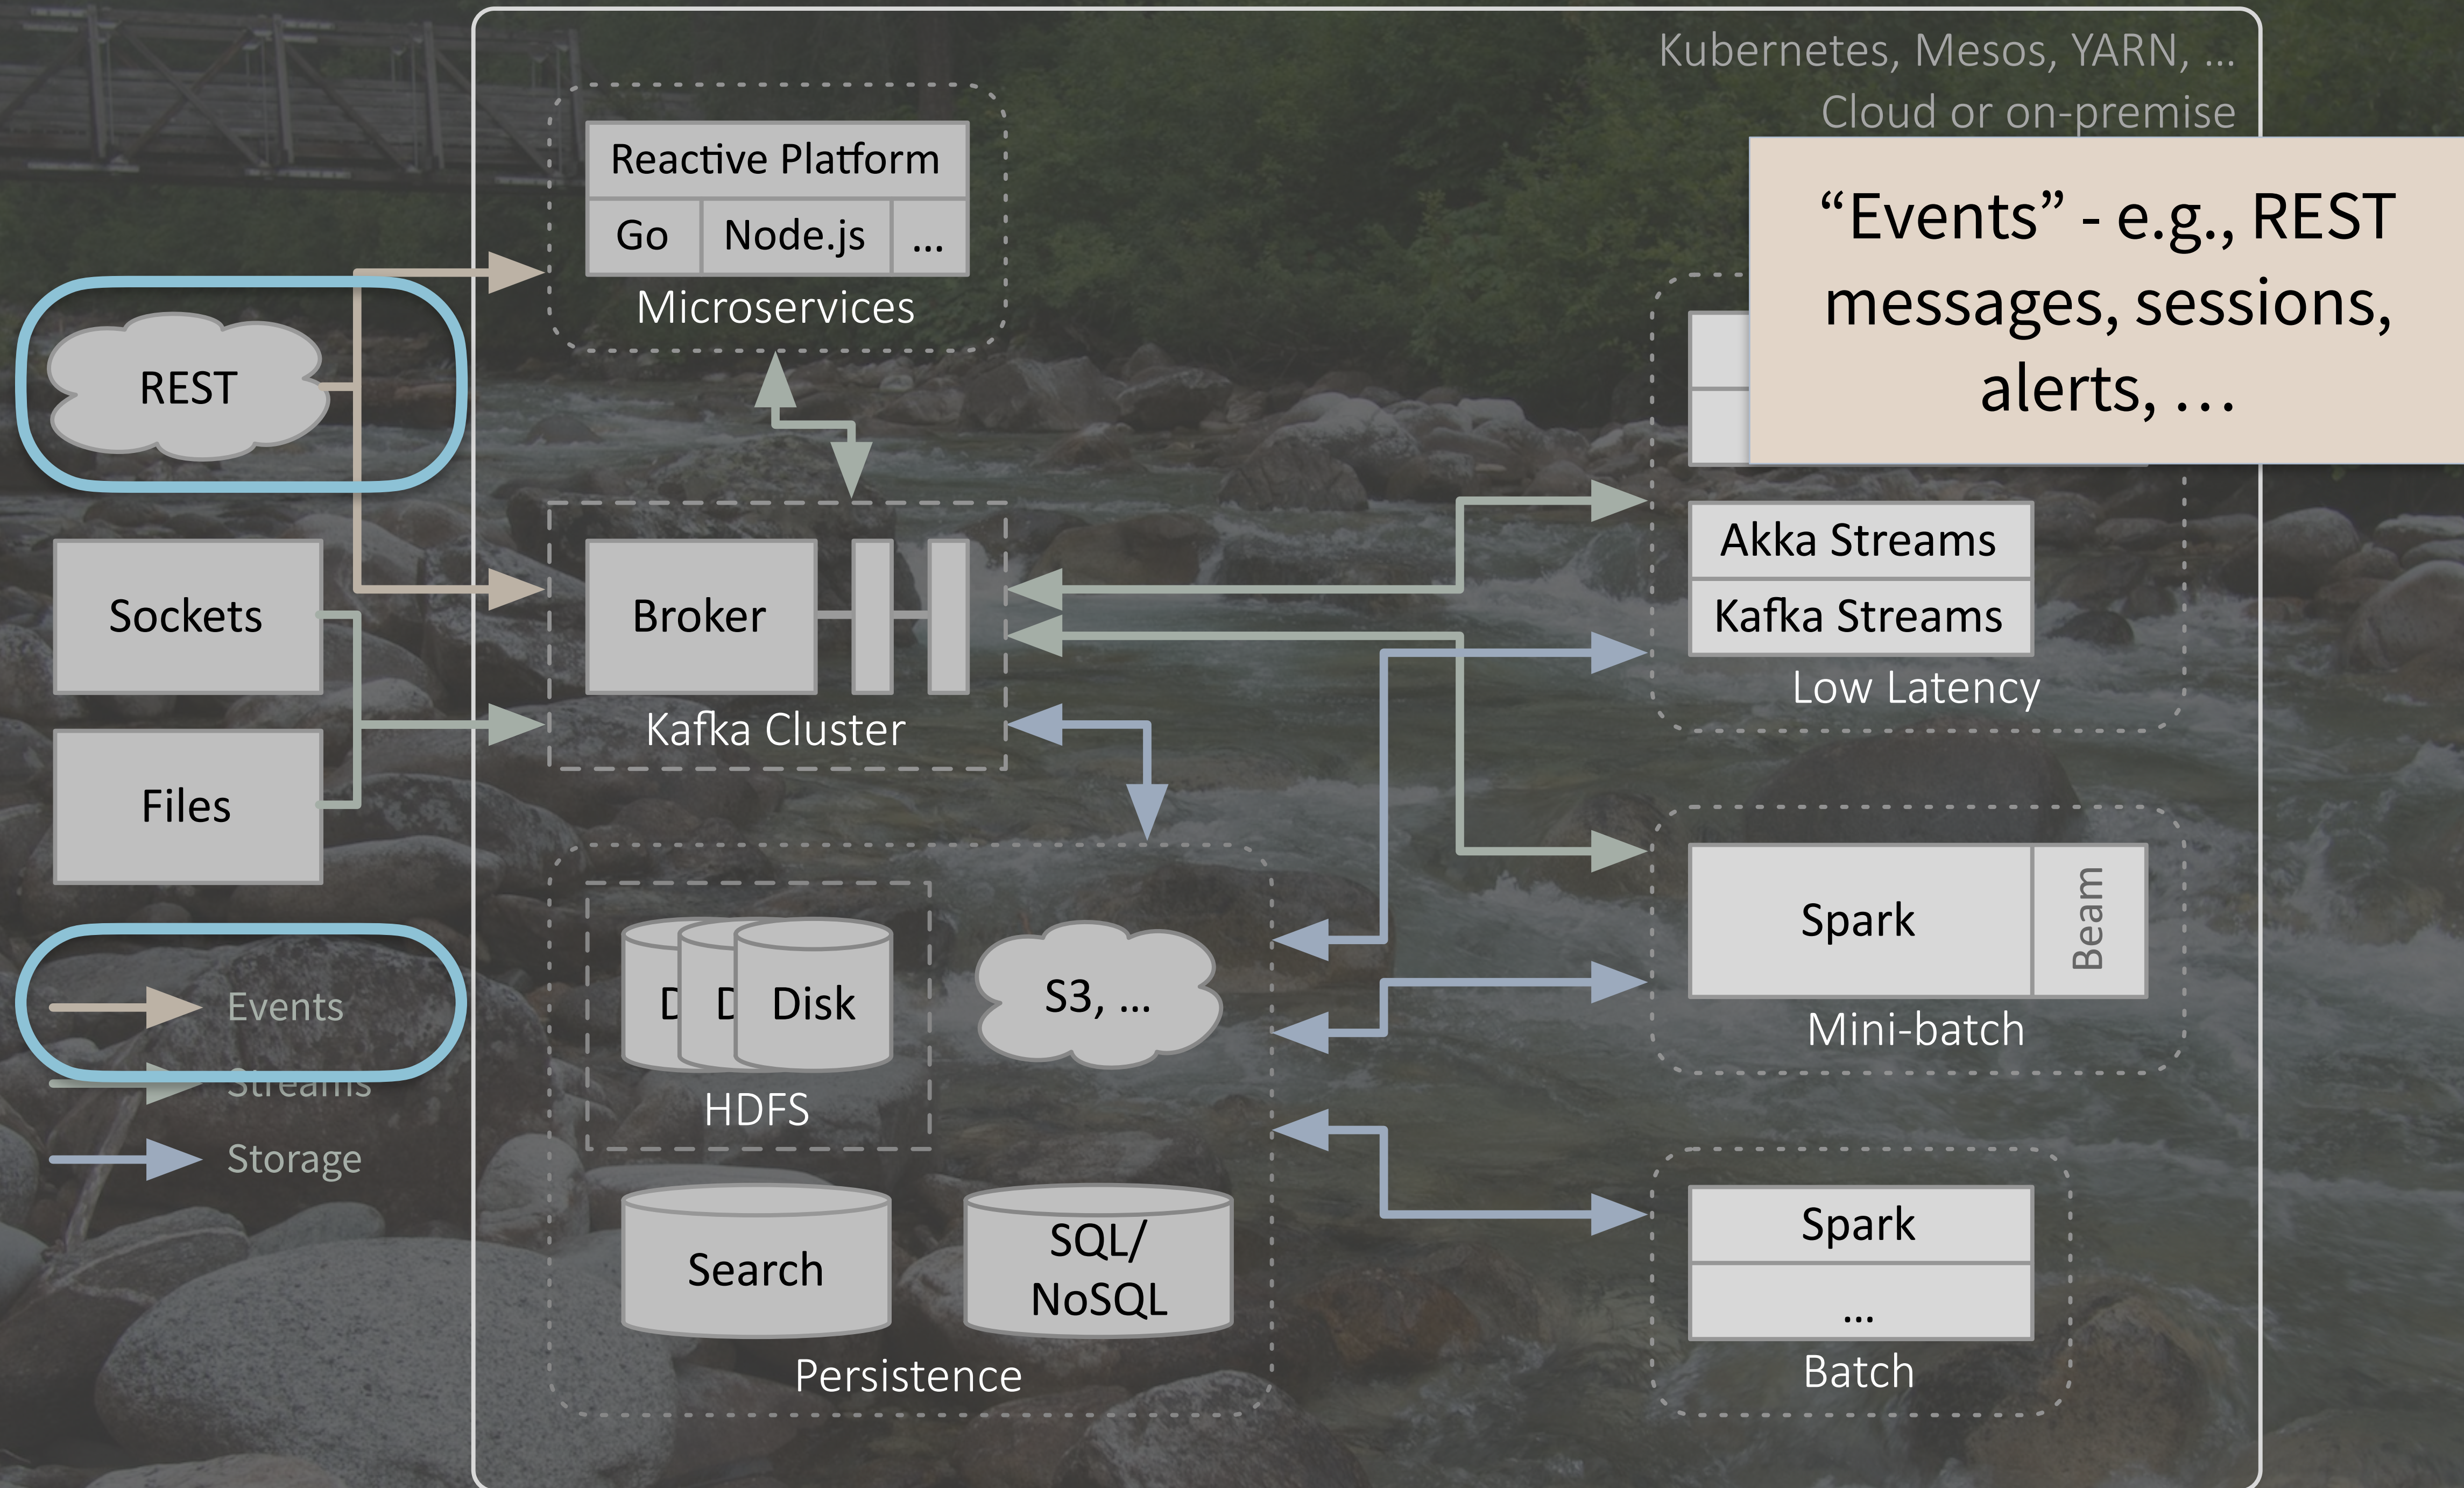


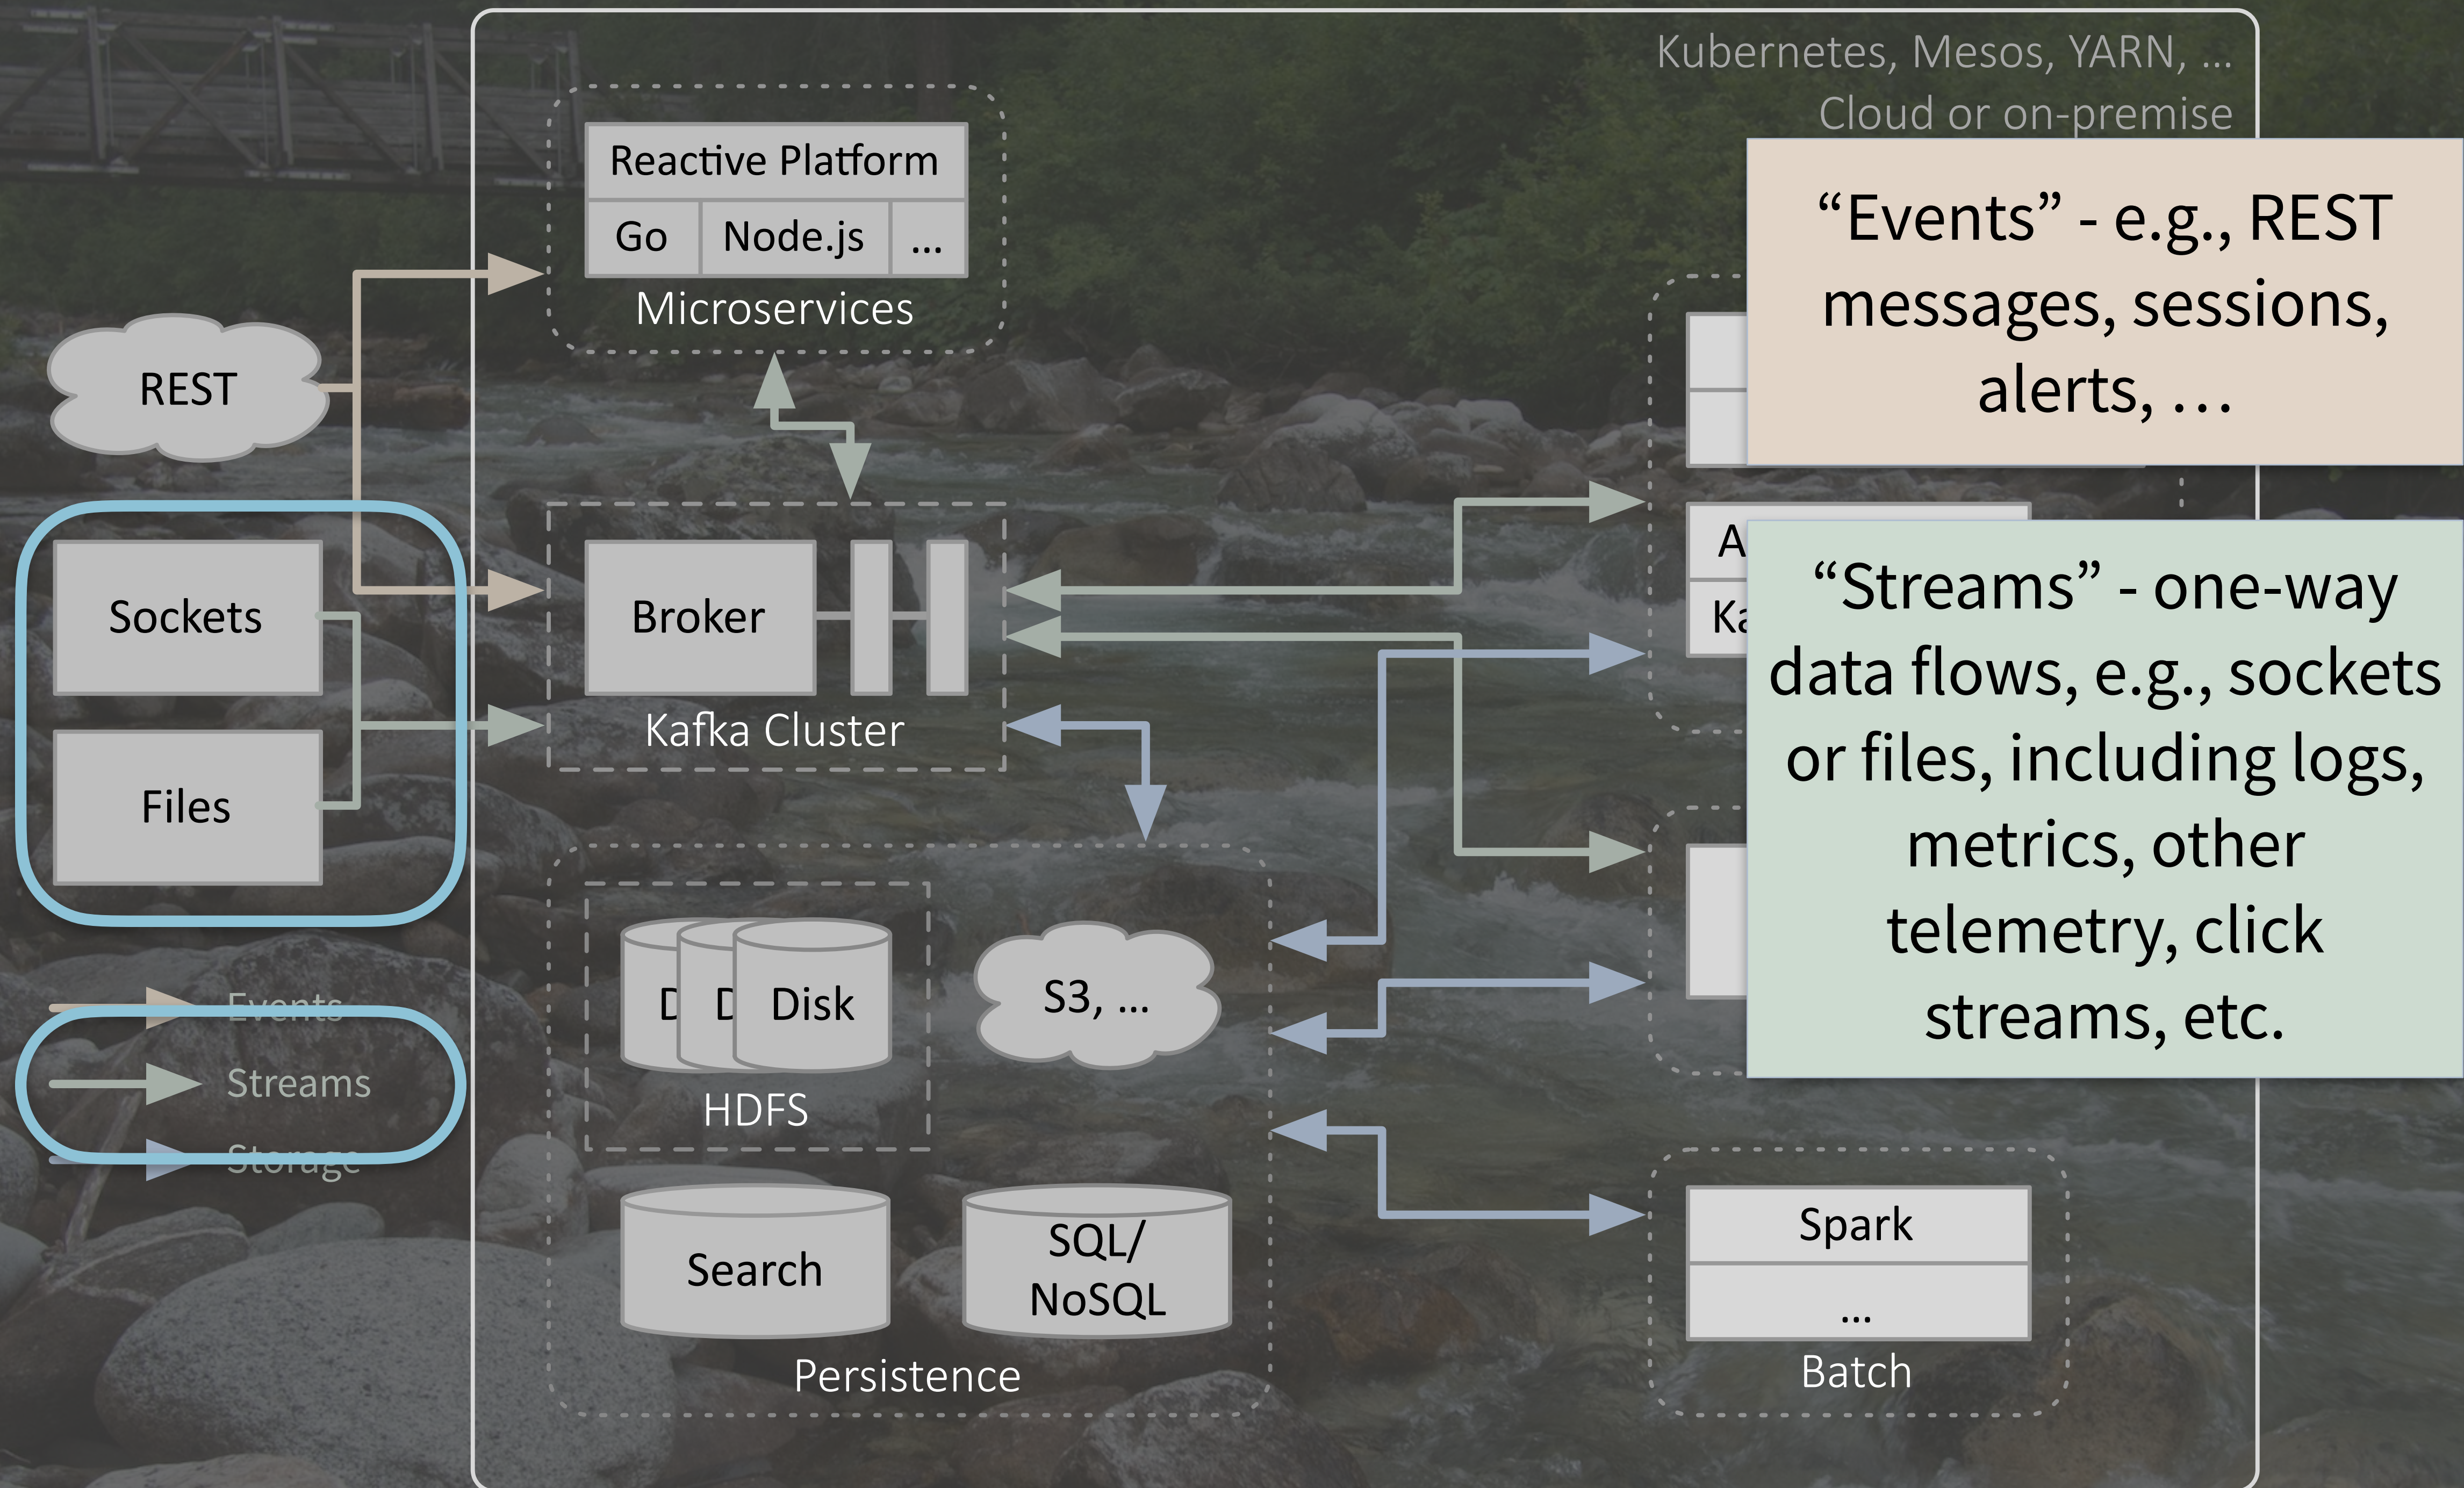
New Streaming, “Fast Data” Architecture

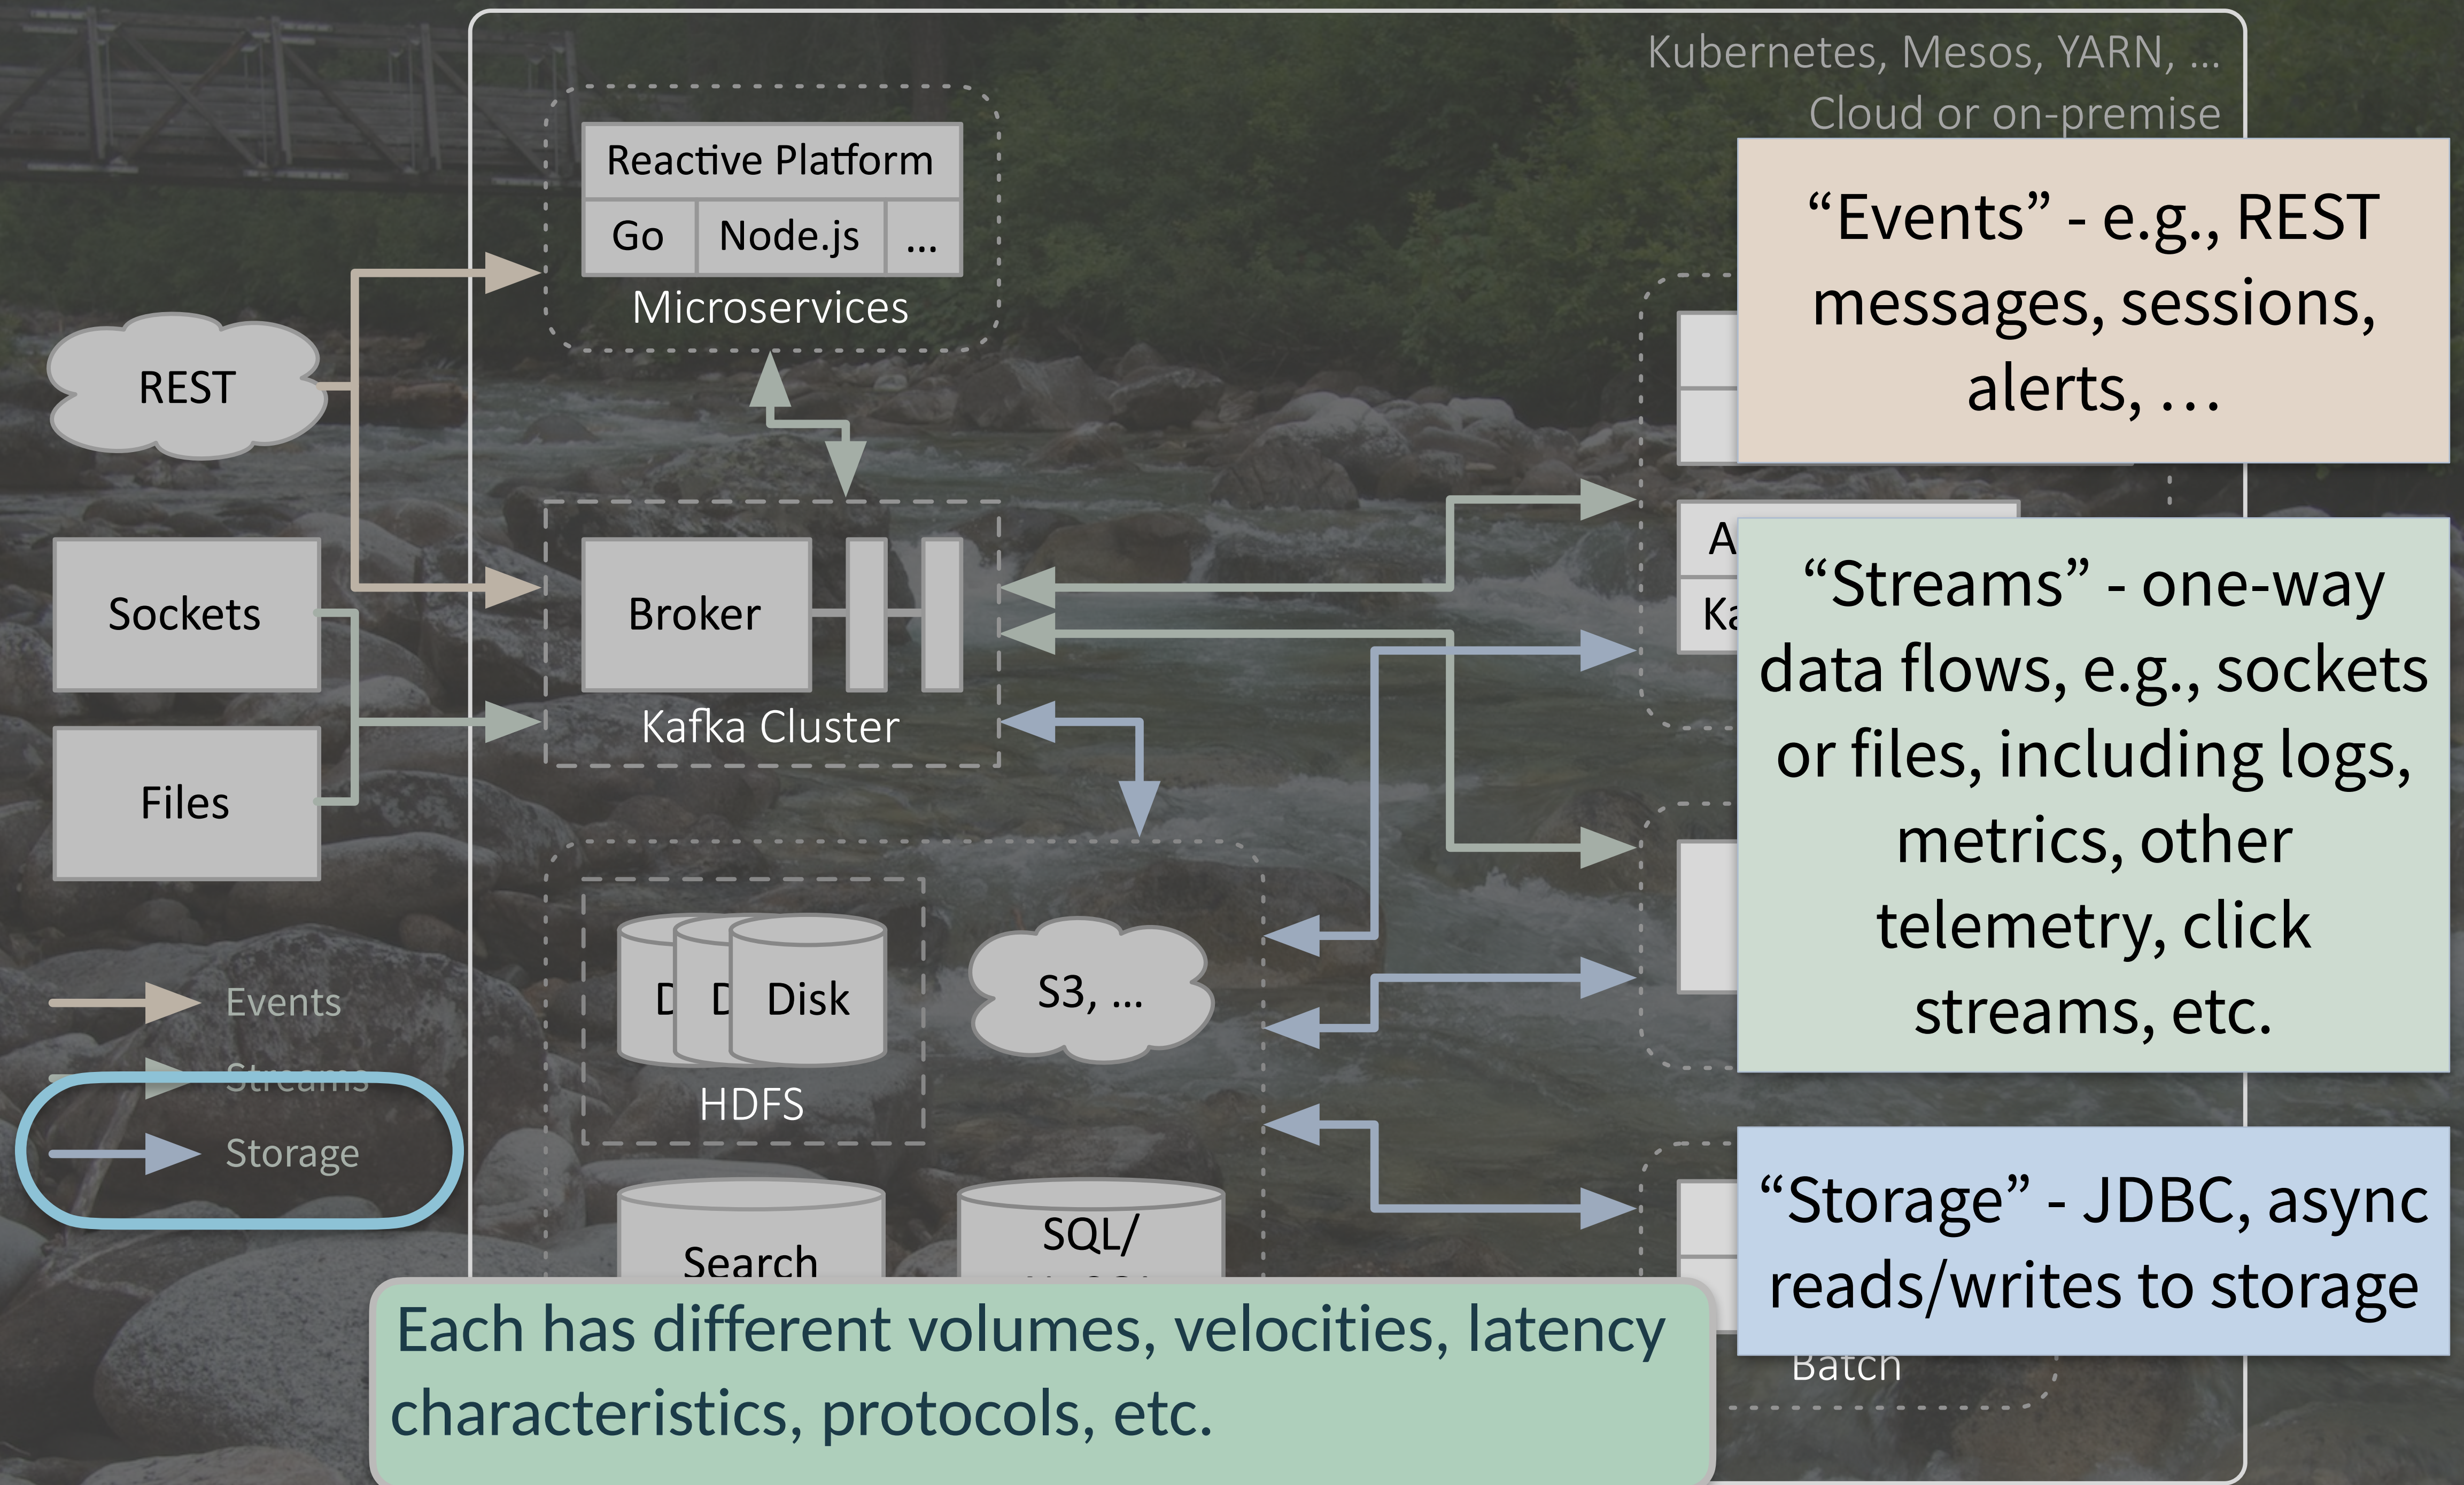


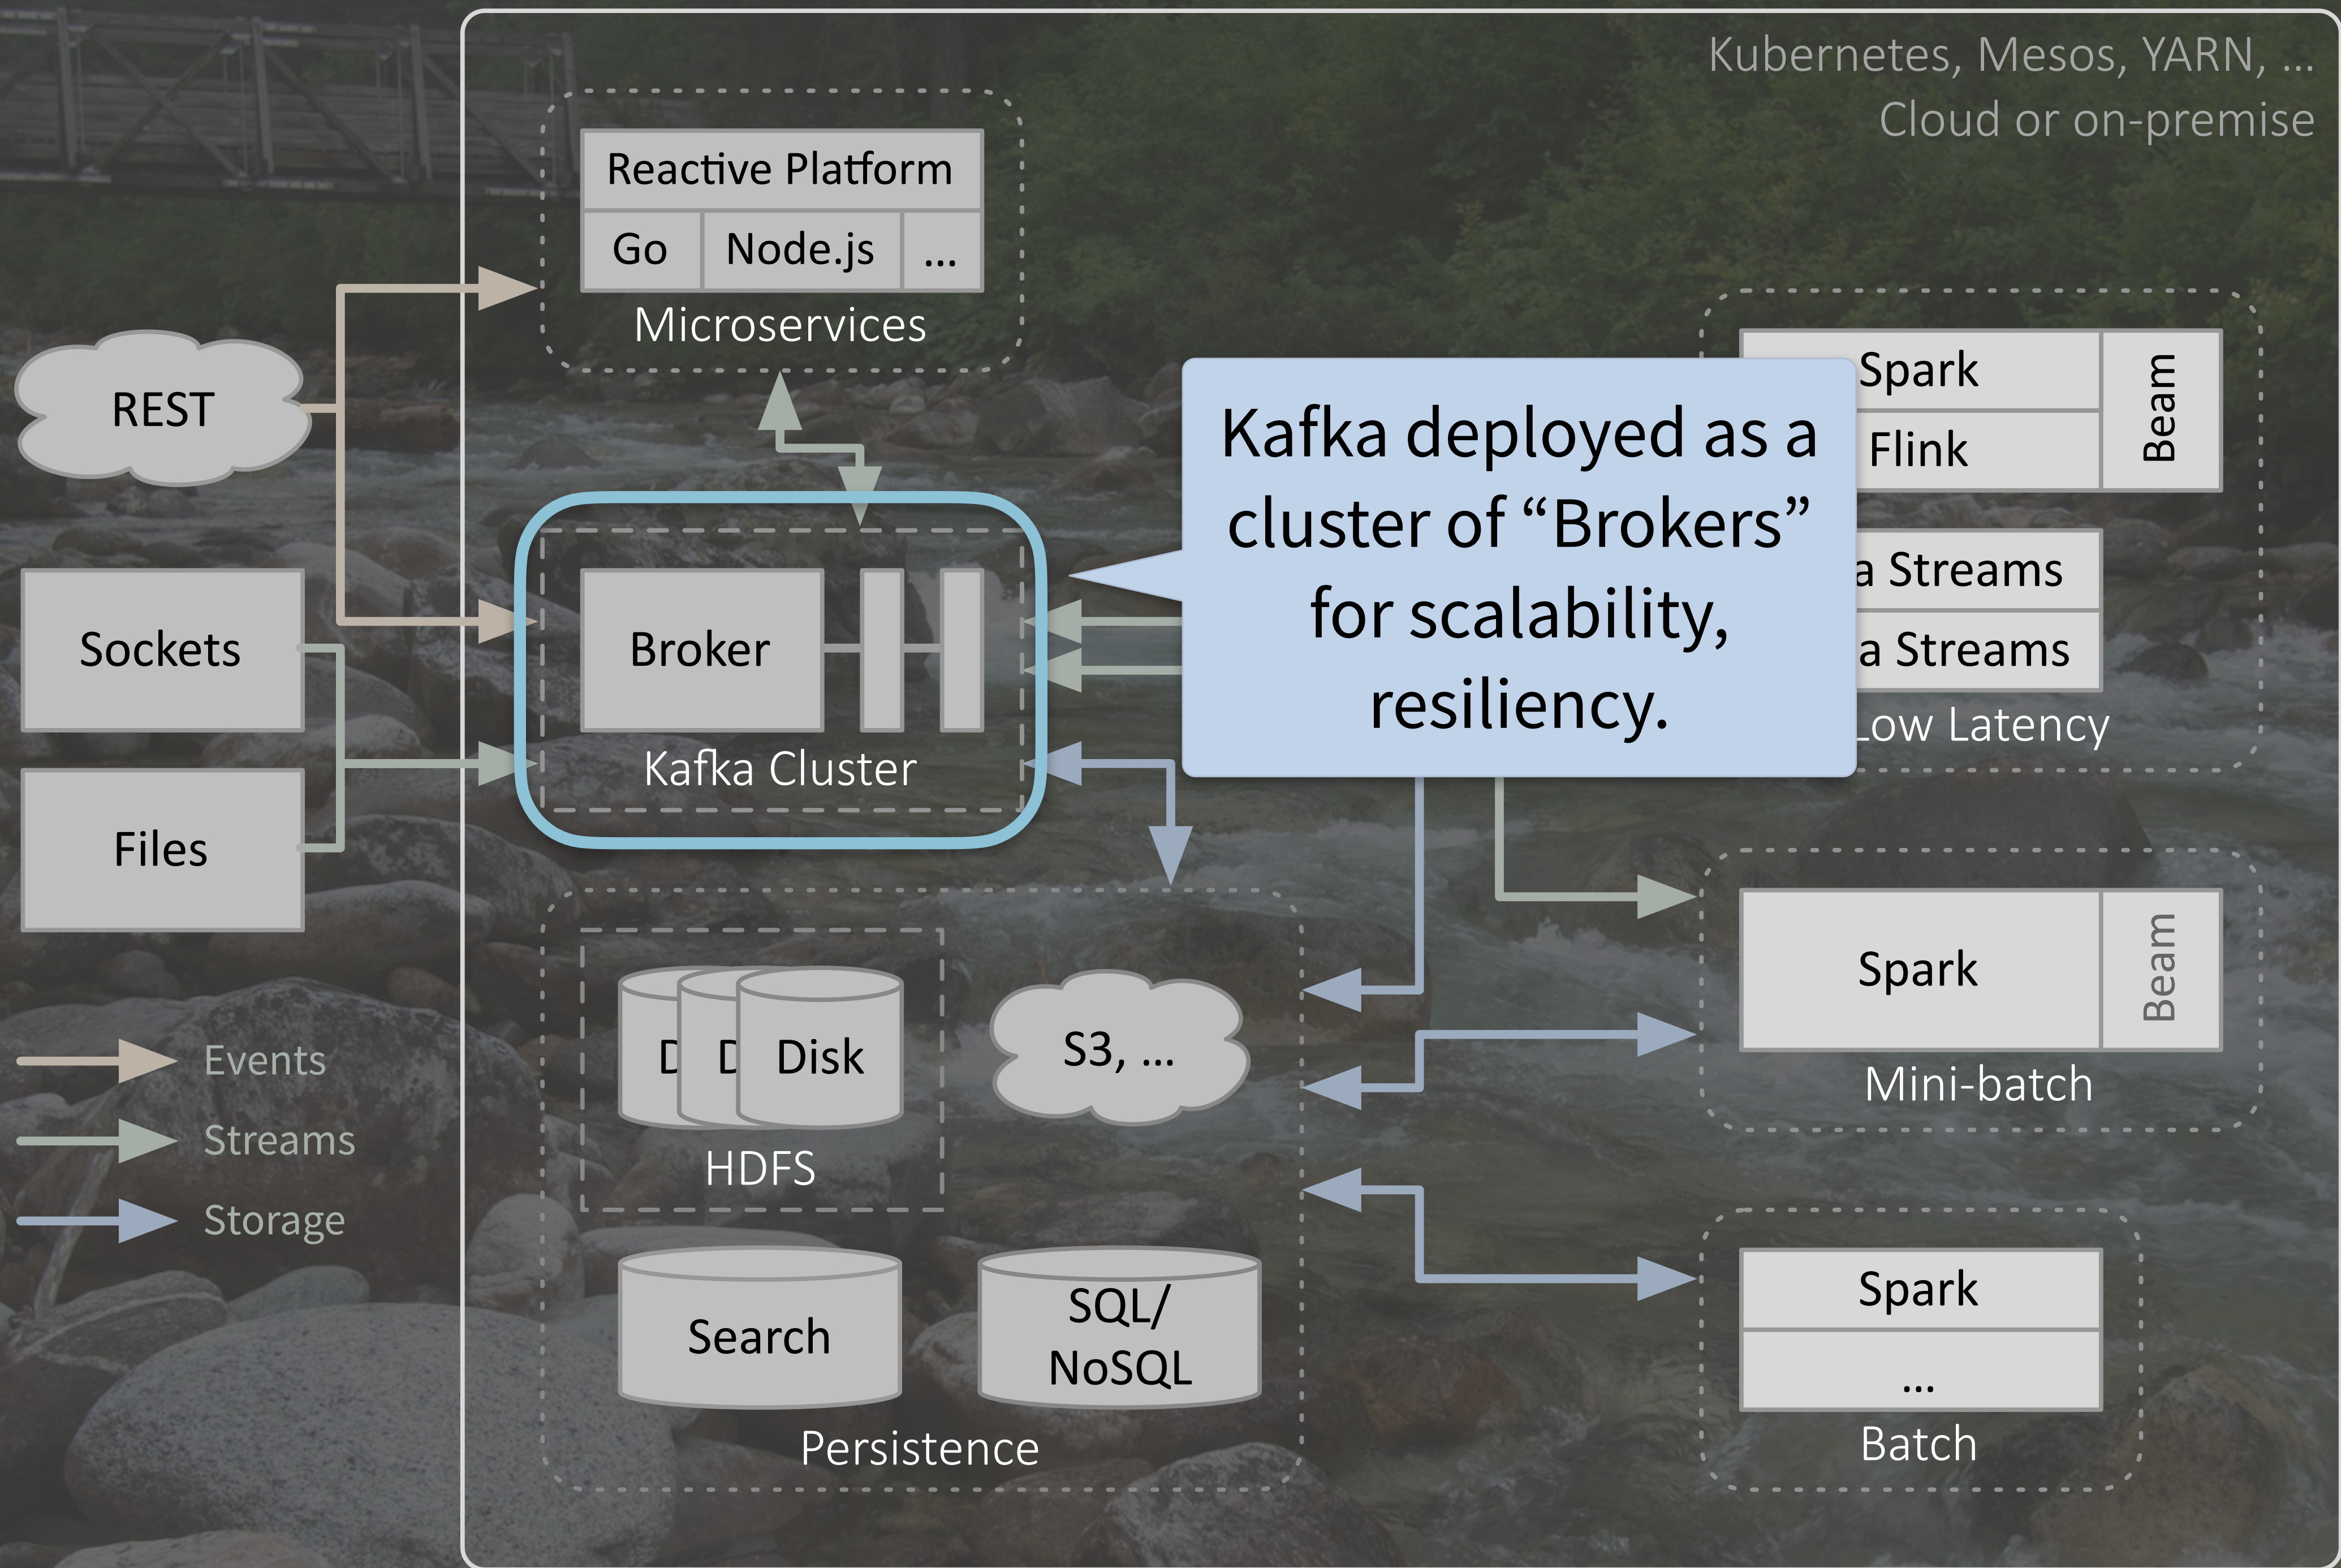


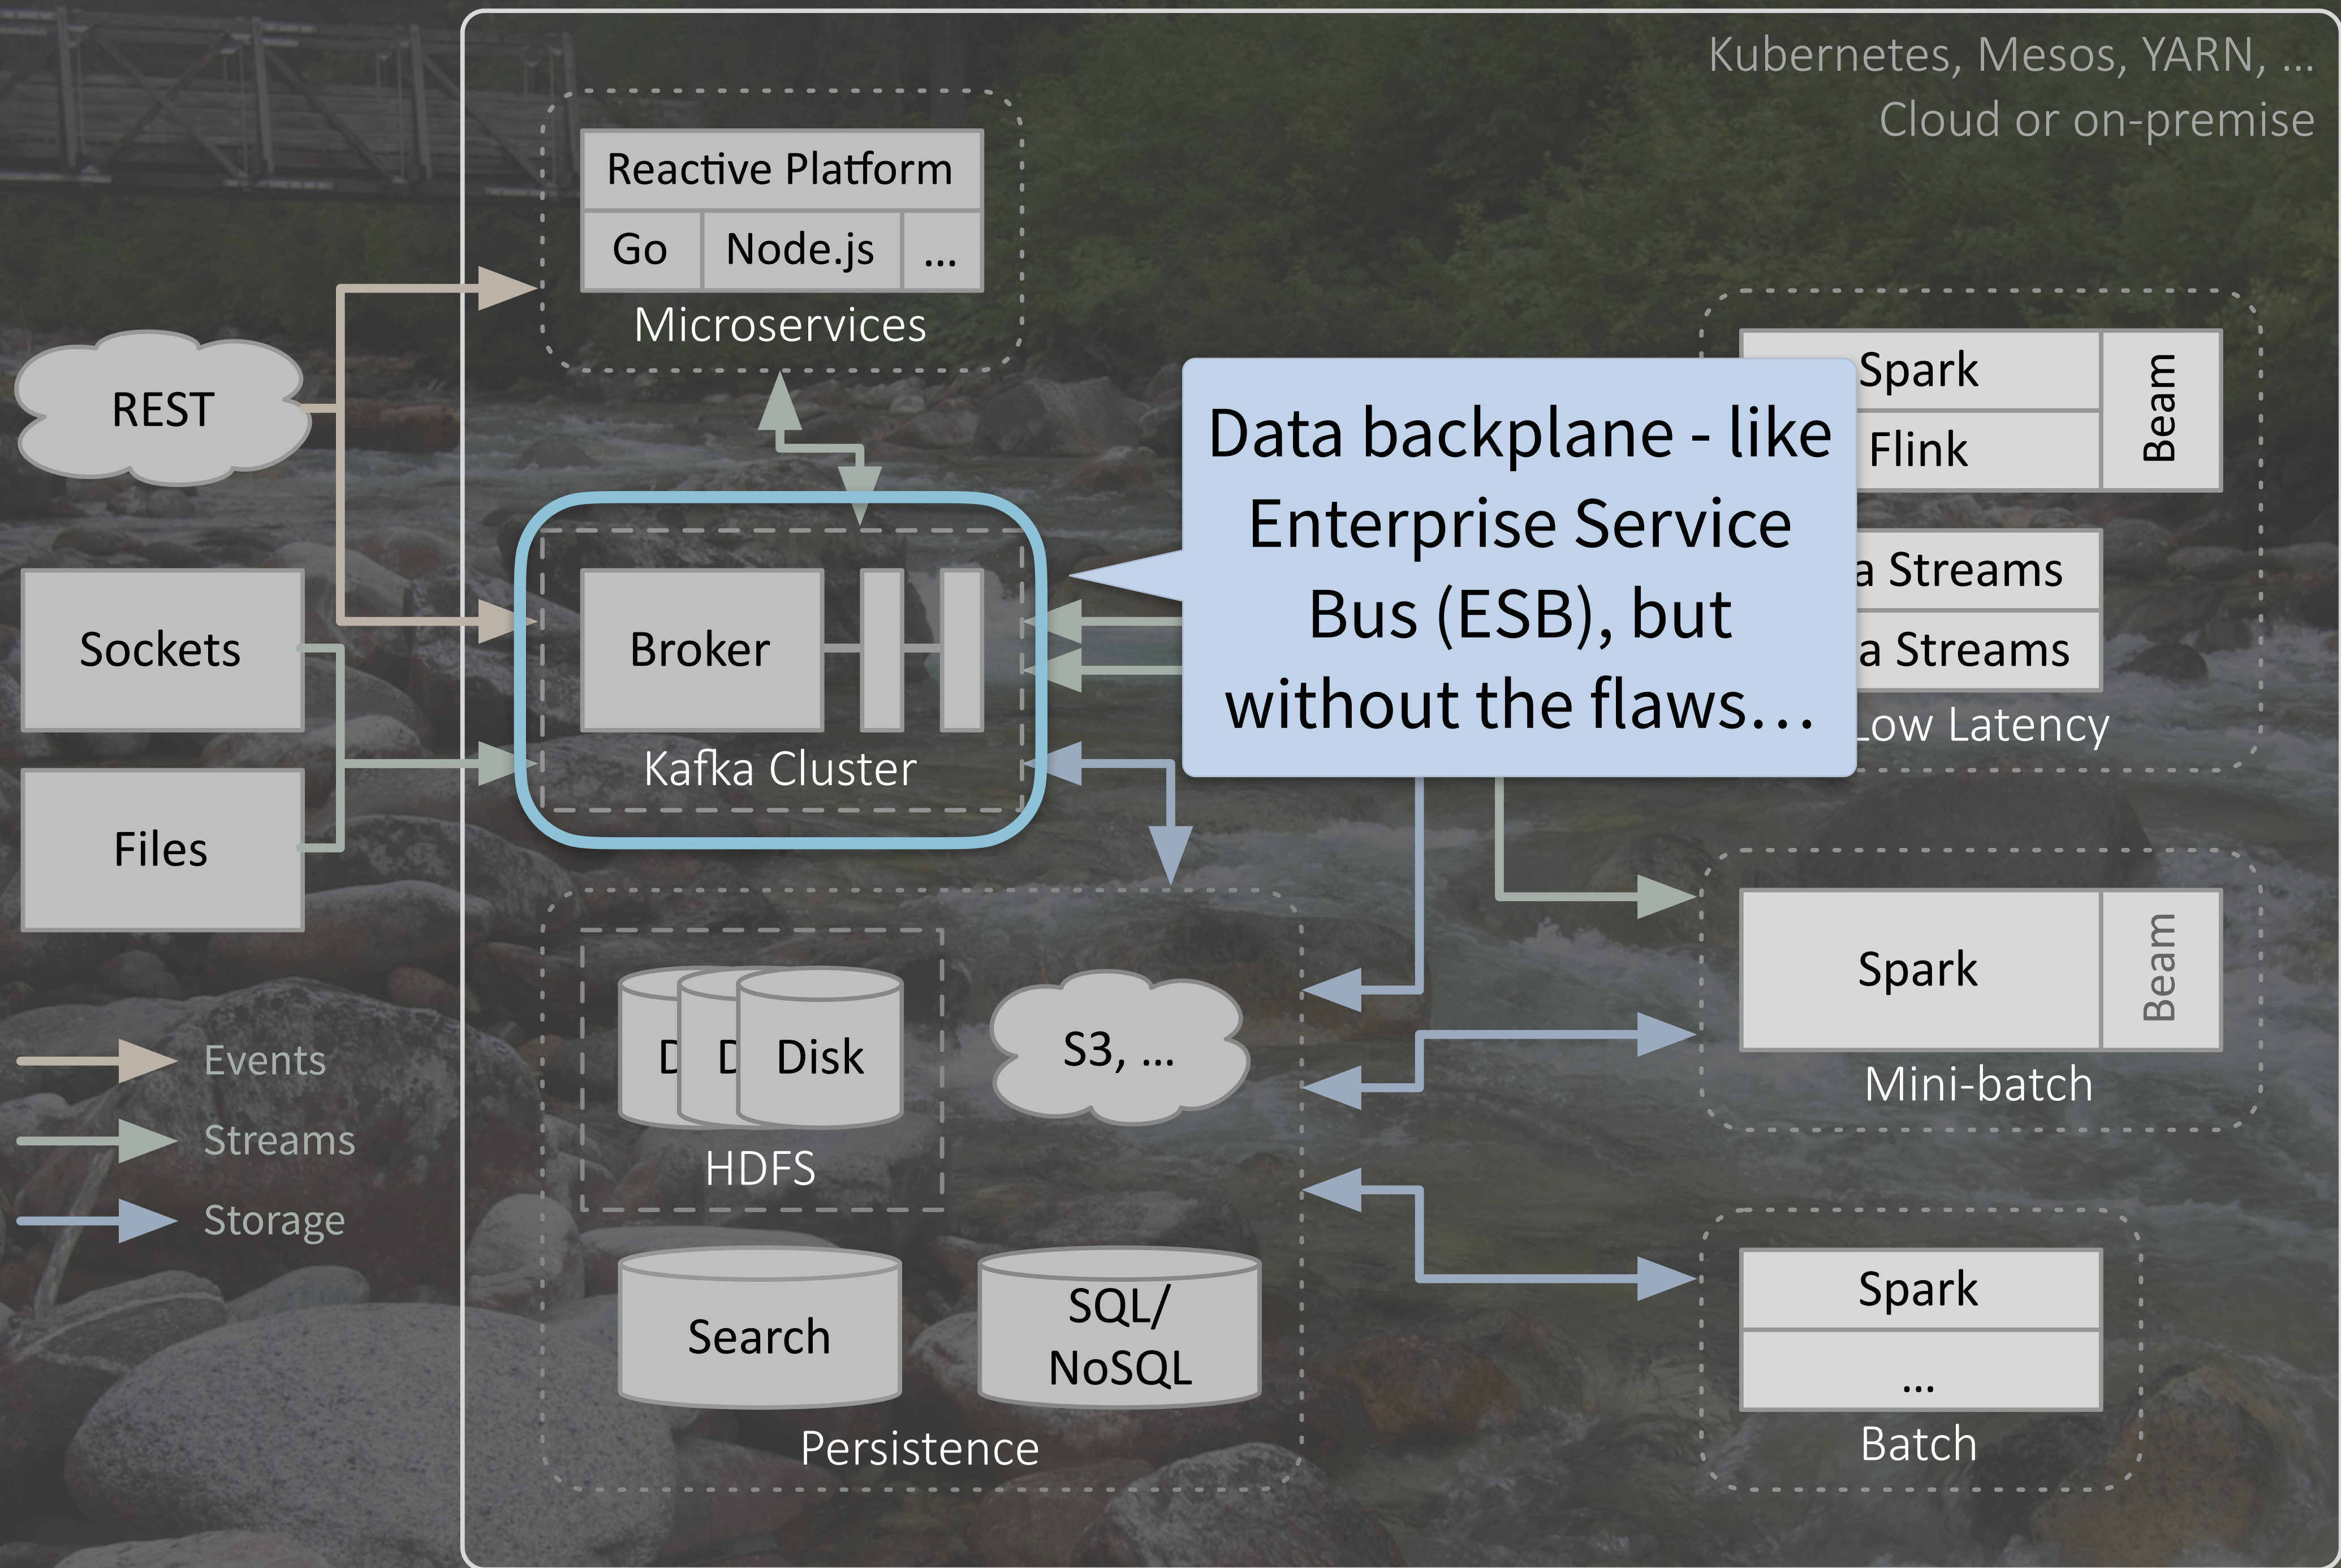












Why Kafka?

Organized into topics

Topics are partitioned, replicated, and distributed

Kafka

Partition 1

Topic A

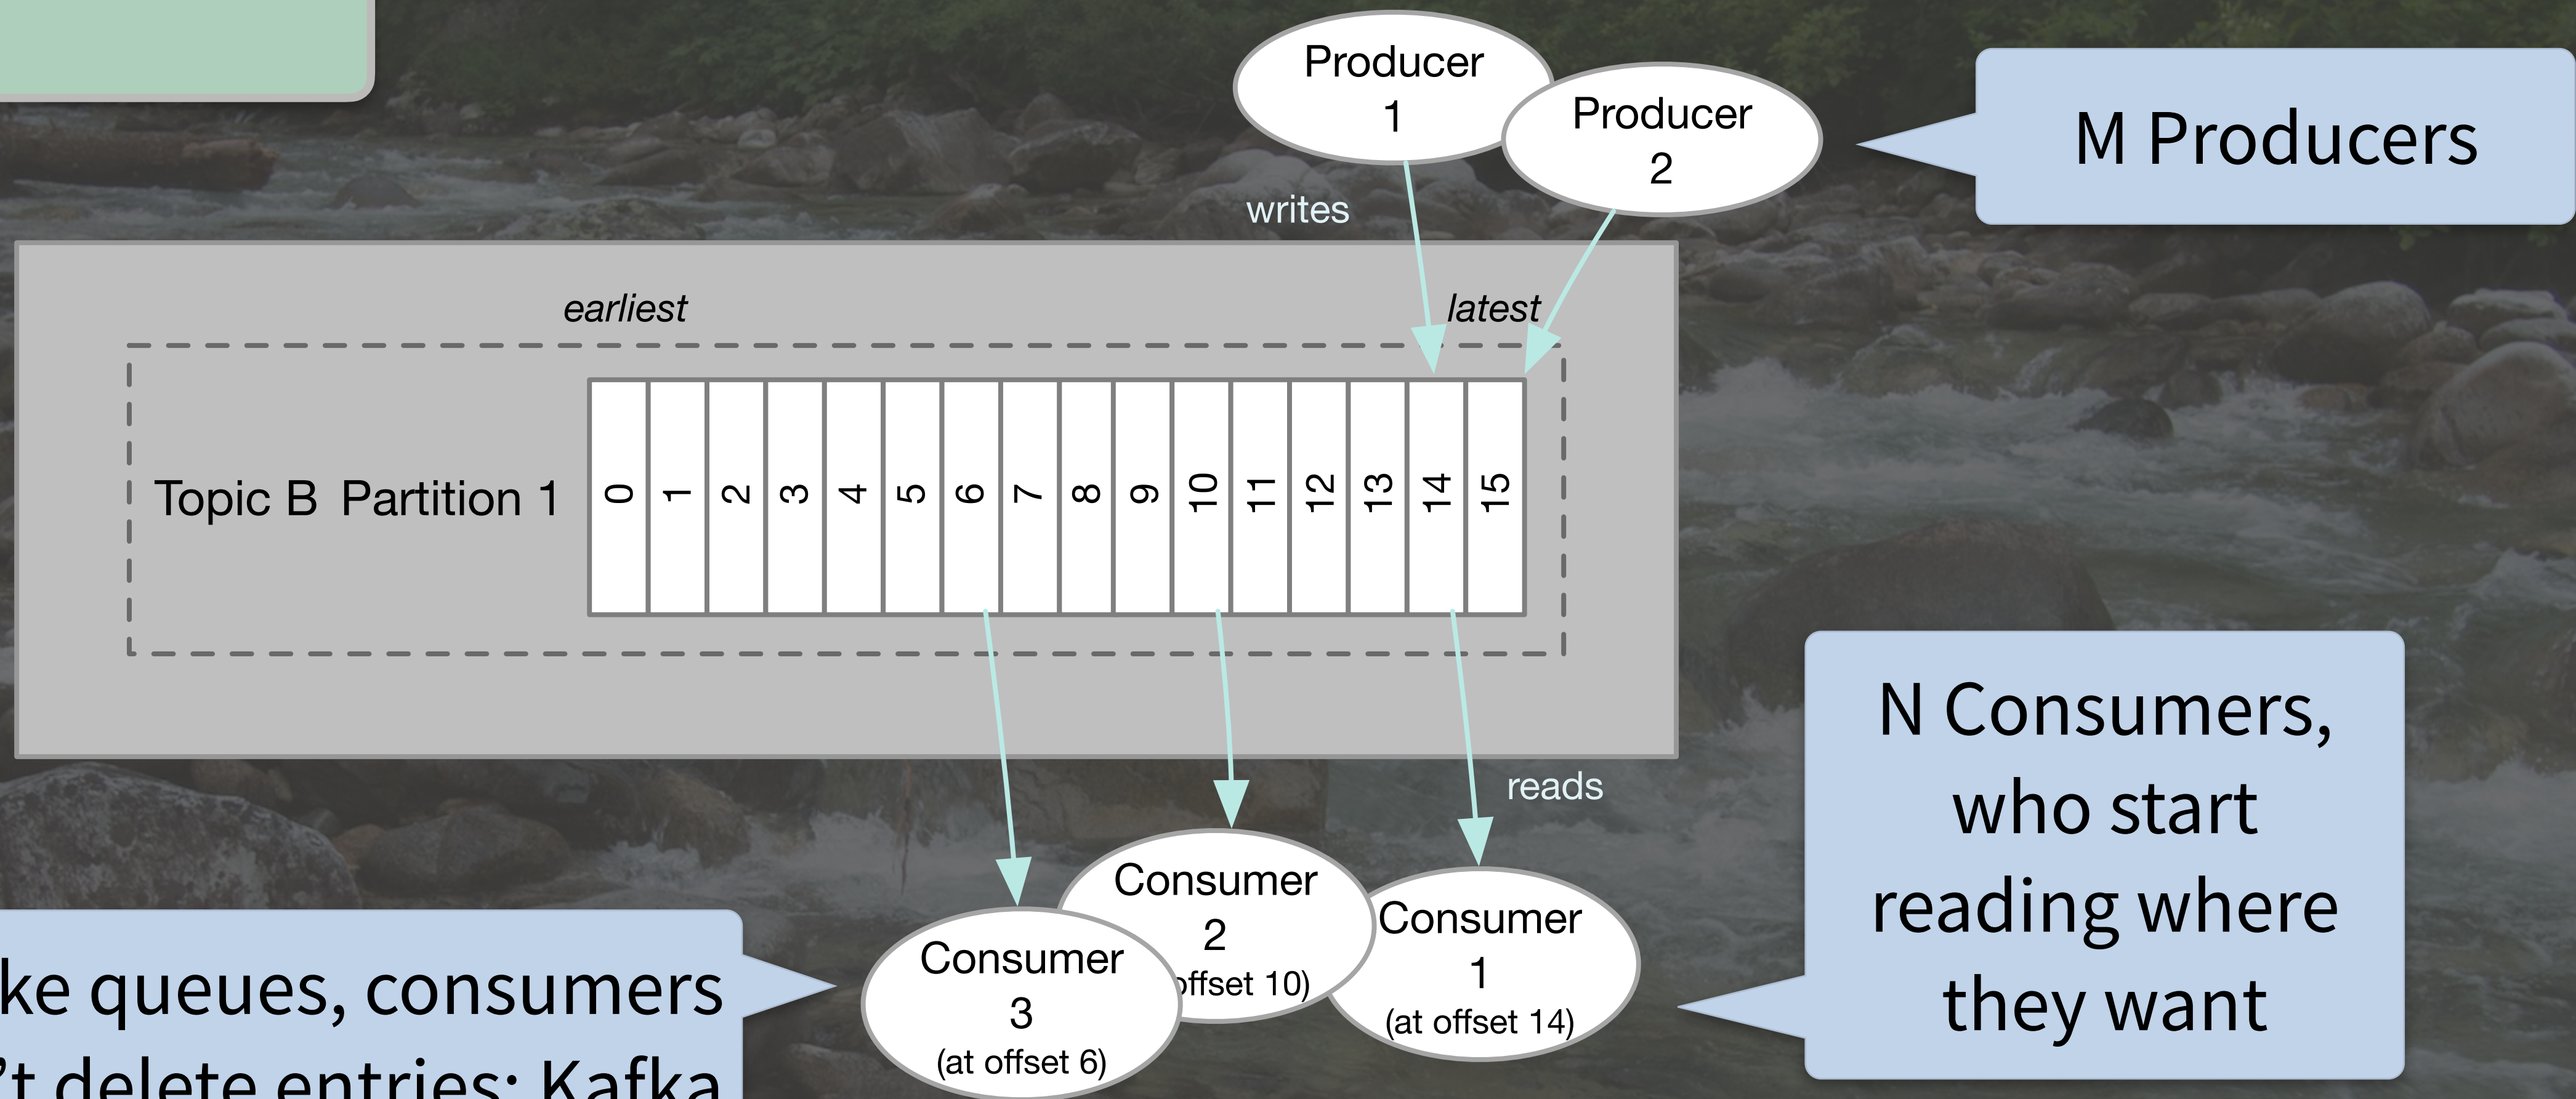
Partition 2

Topic B Partition 1



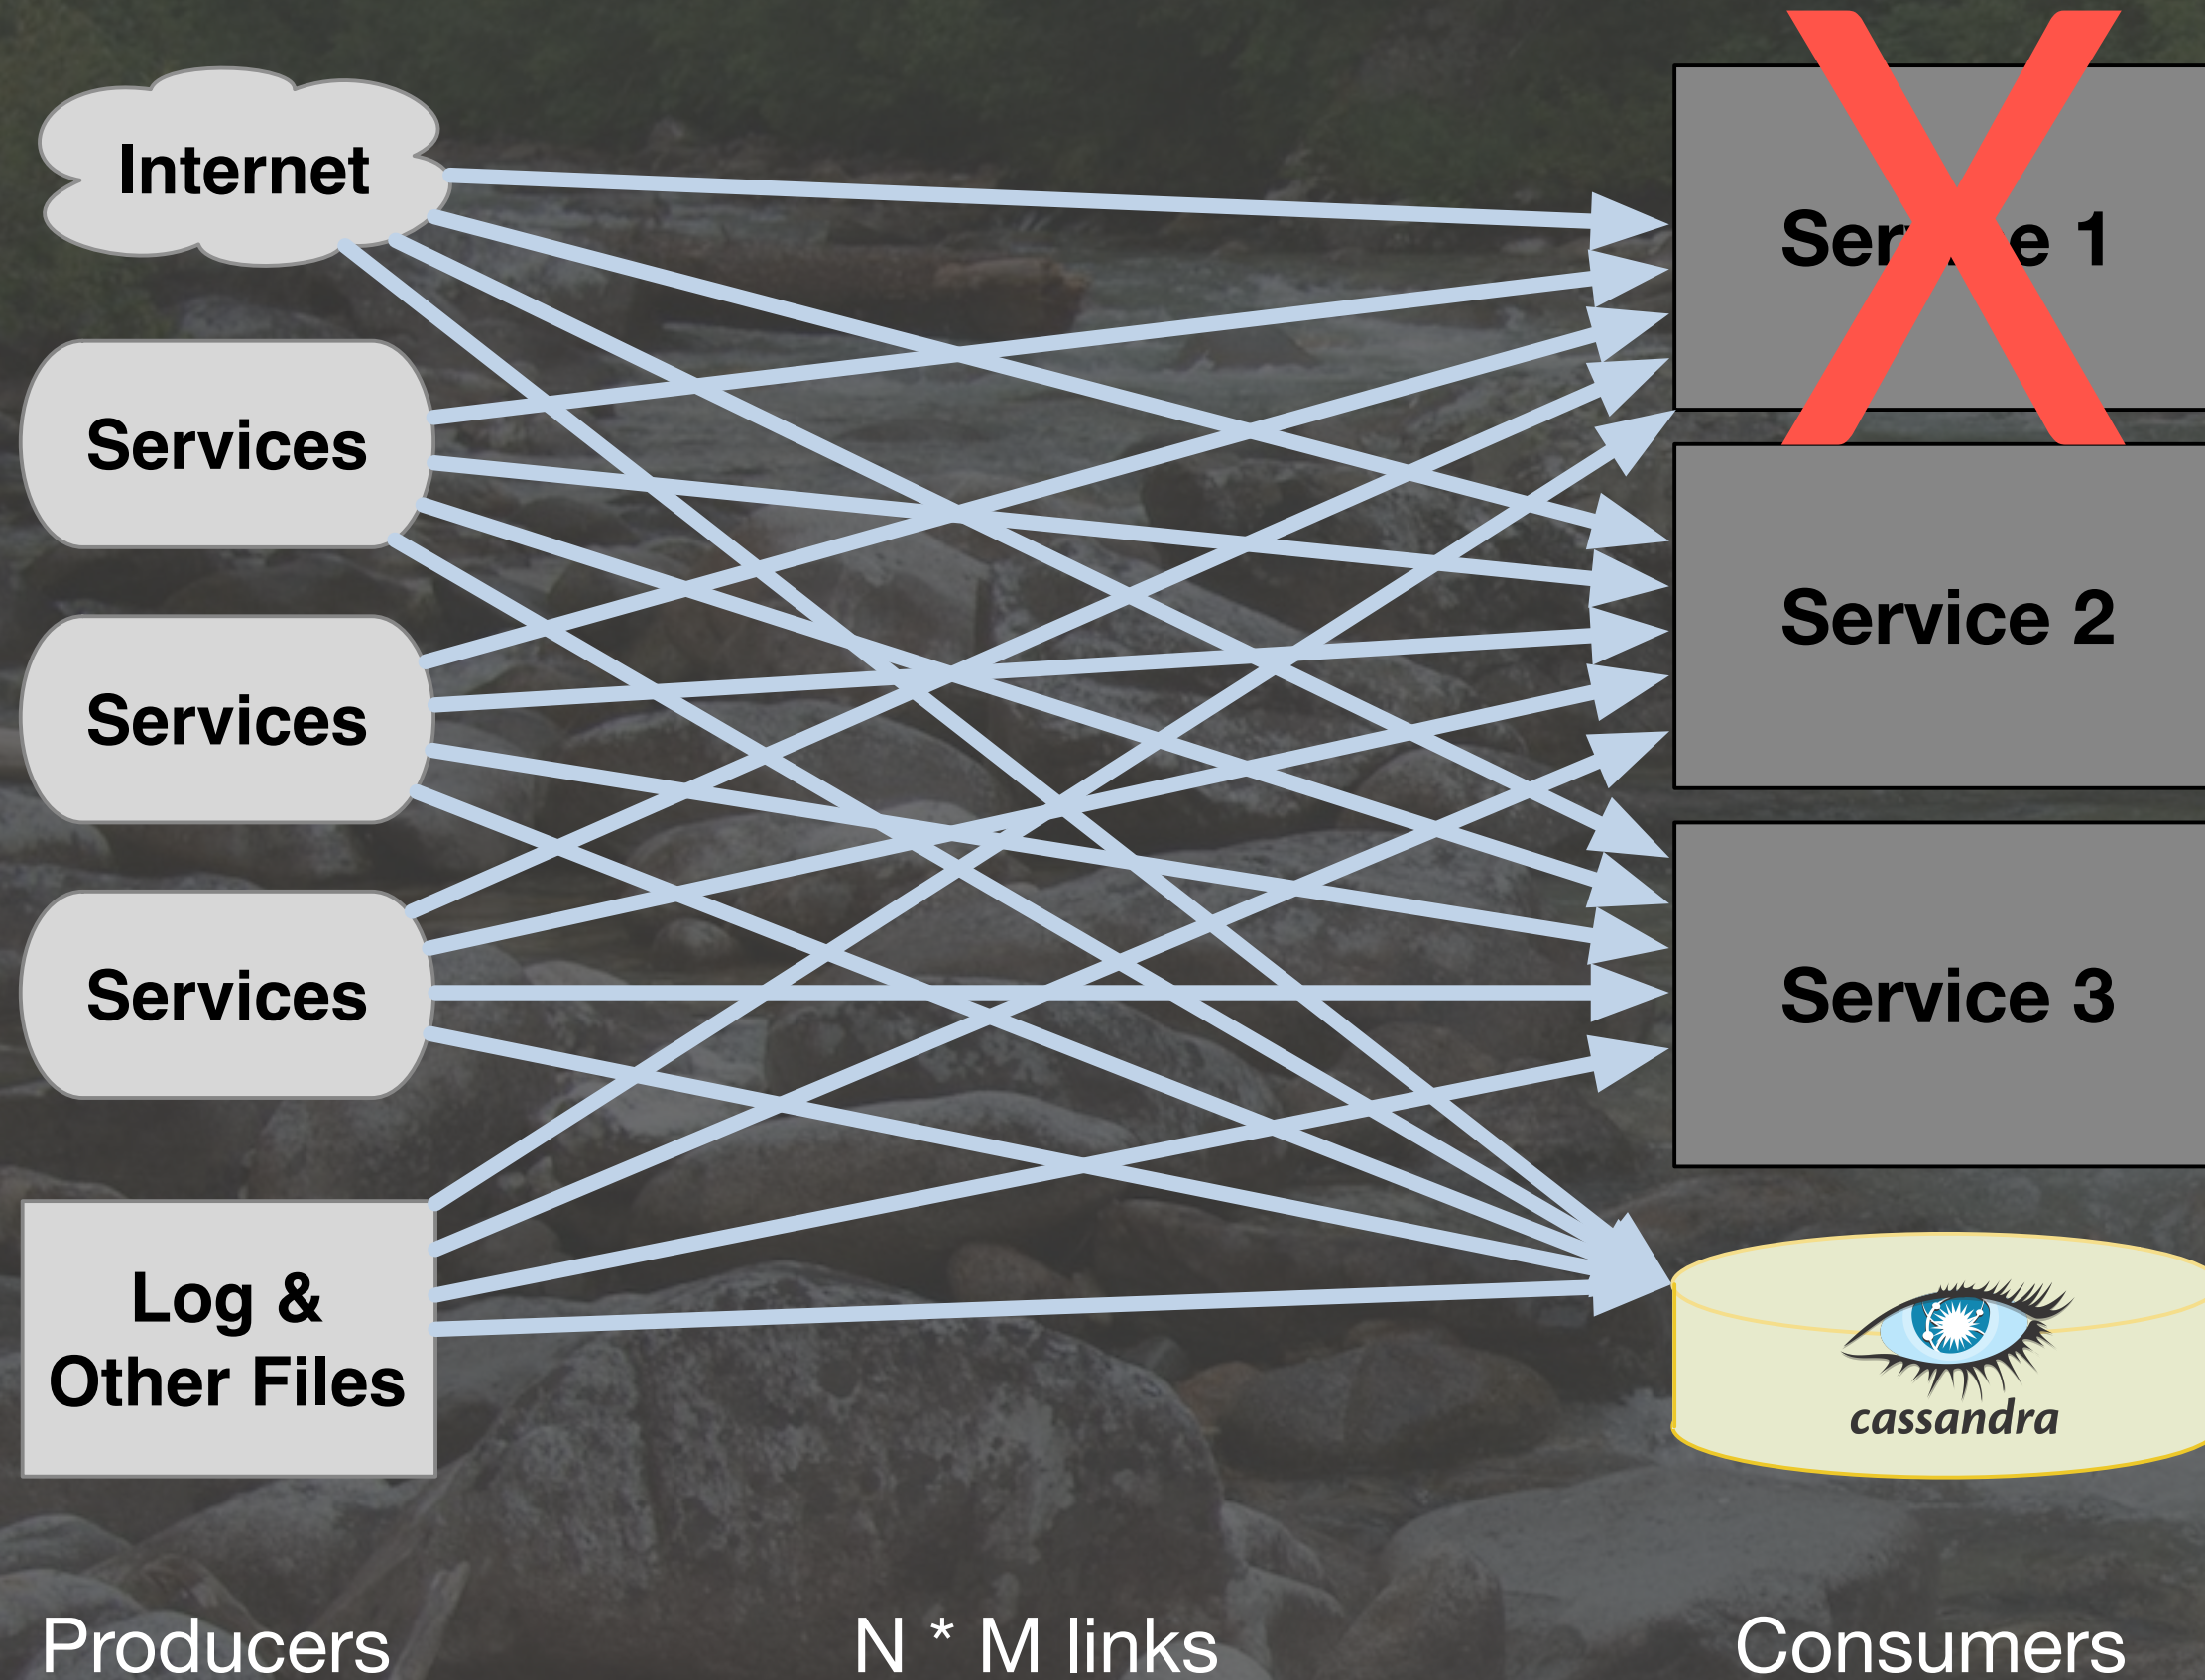
Why Kafka?

Logs, *not* queues!



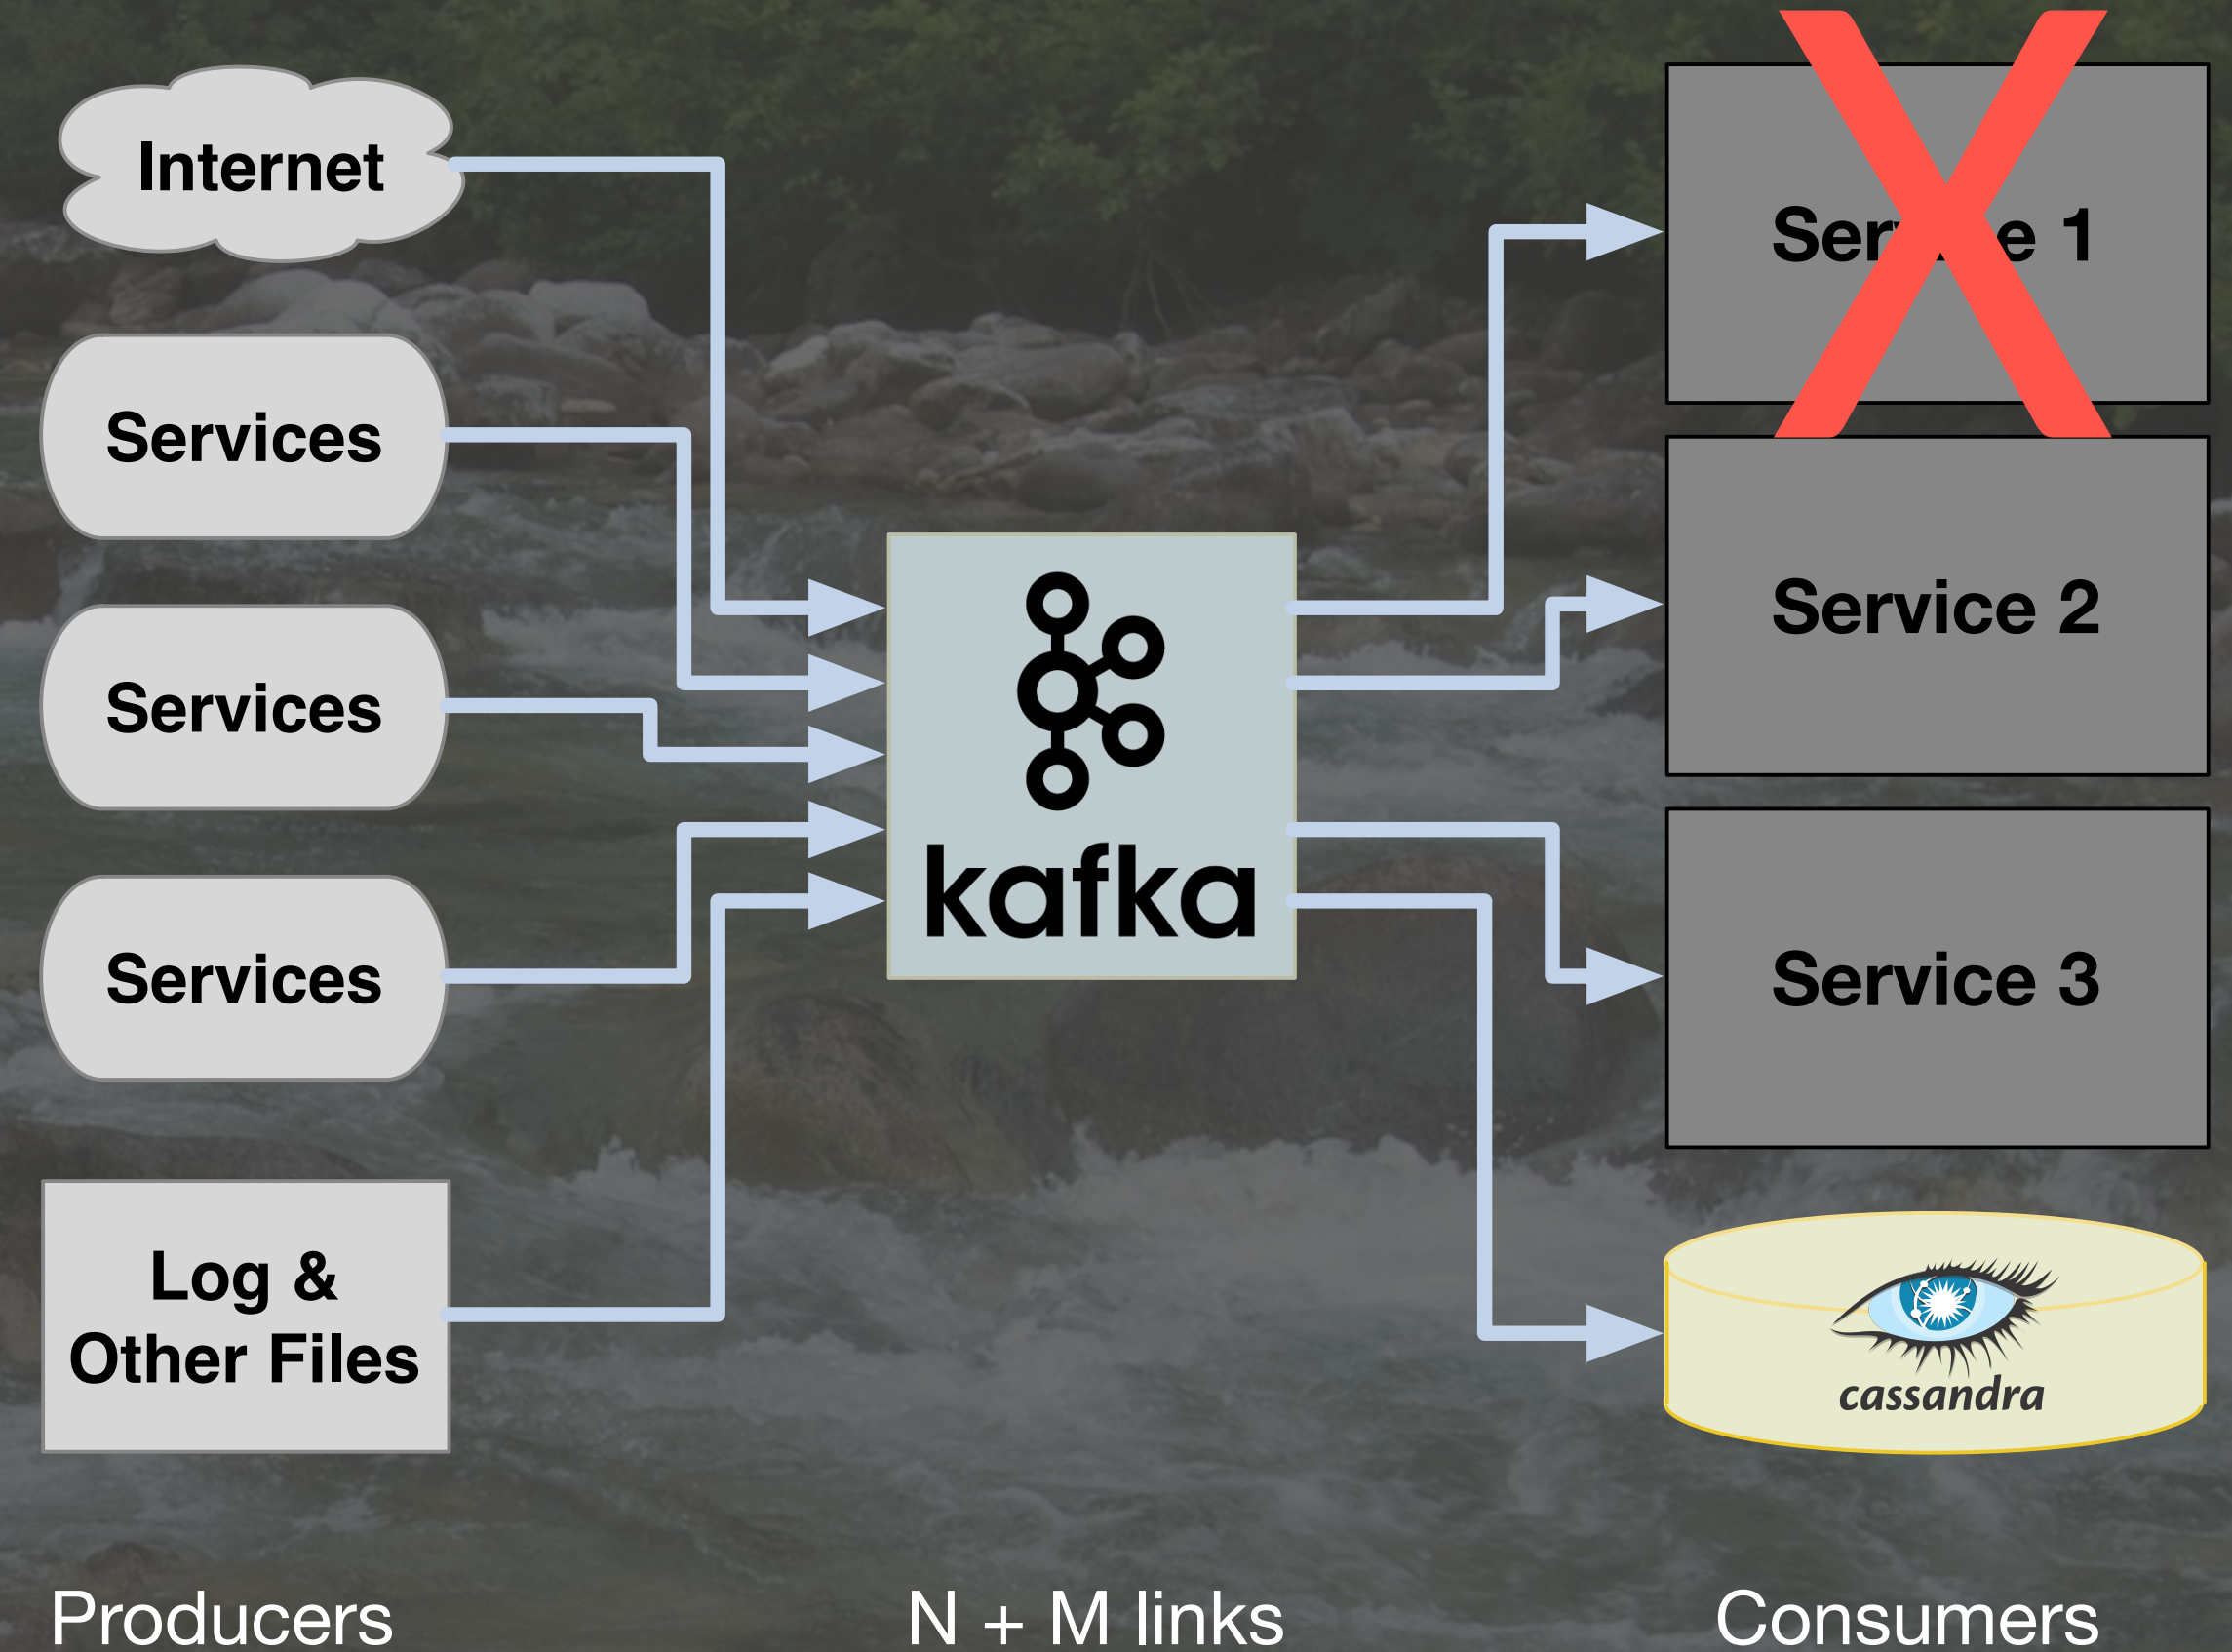
Using Kafka

Before:



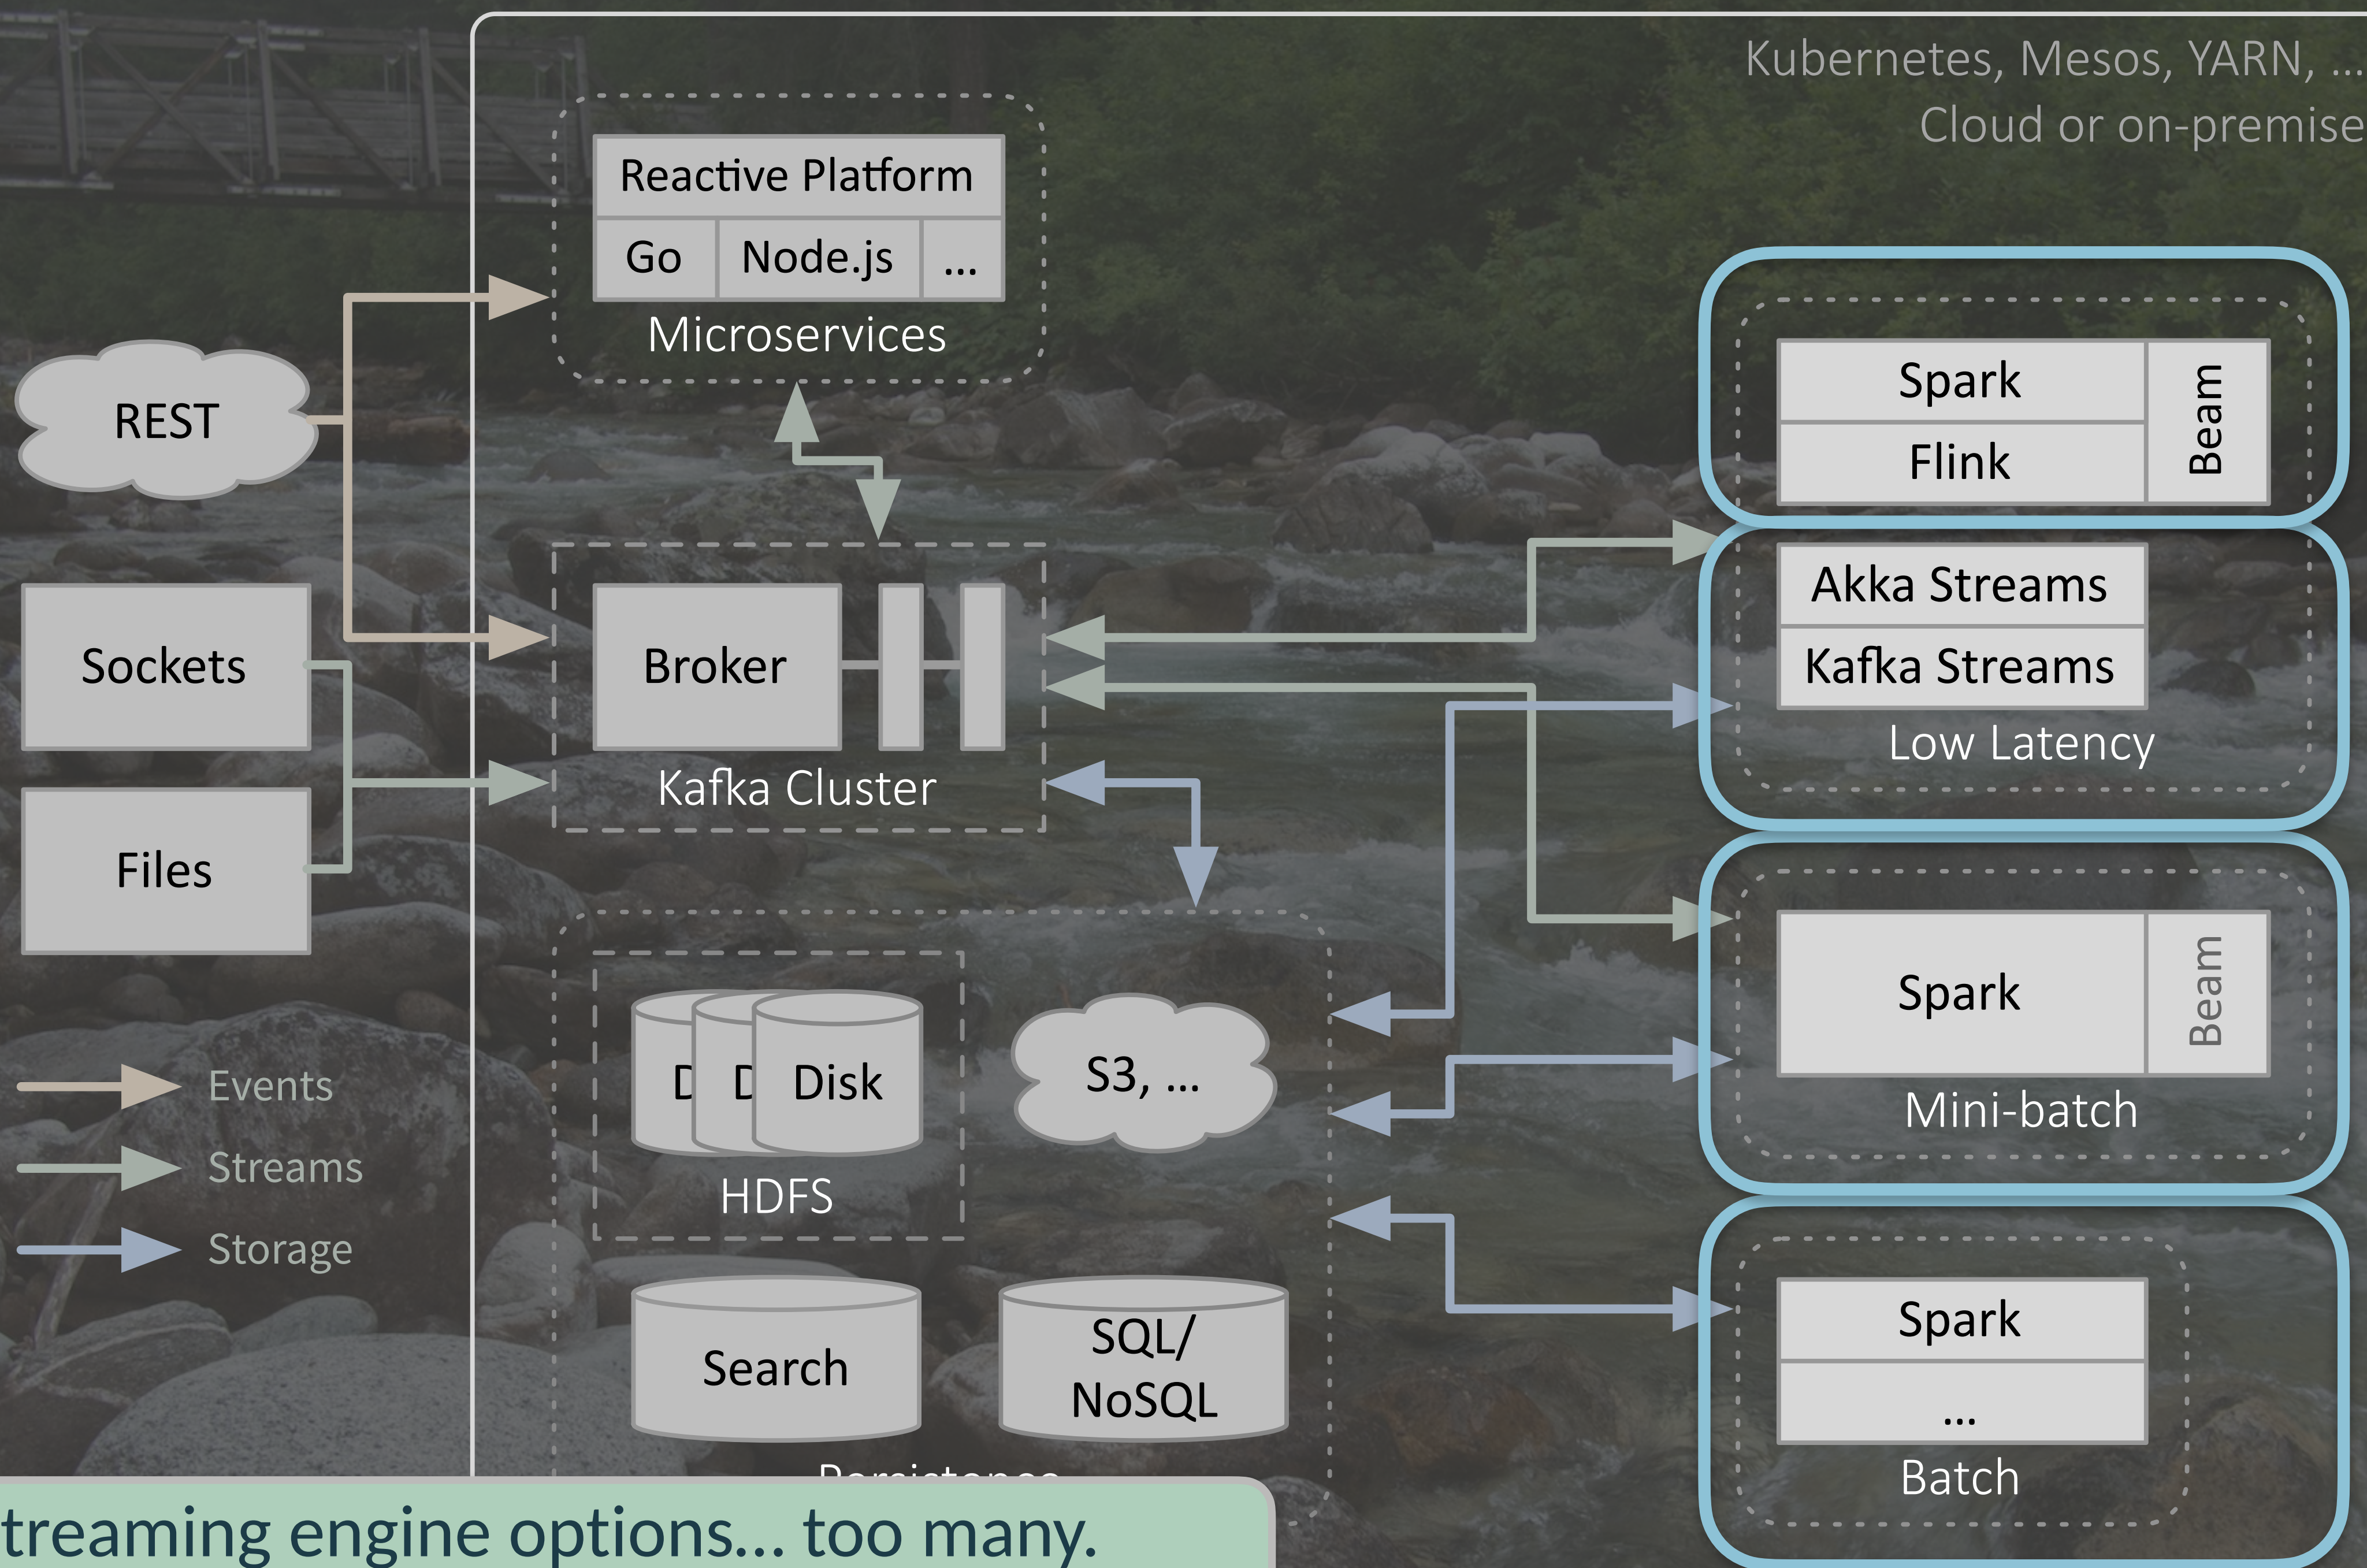
Messy and fragile;
what if "Service 1"
goes down?

After:



Simpler and more
robust! Loss of Service
1 means no data loss.

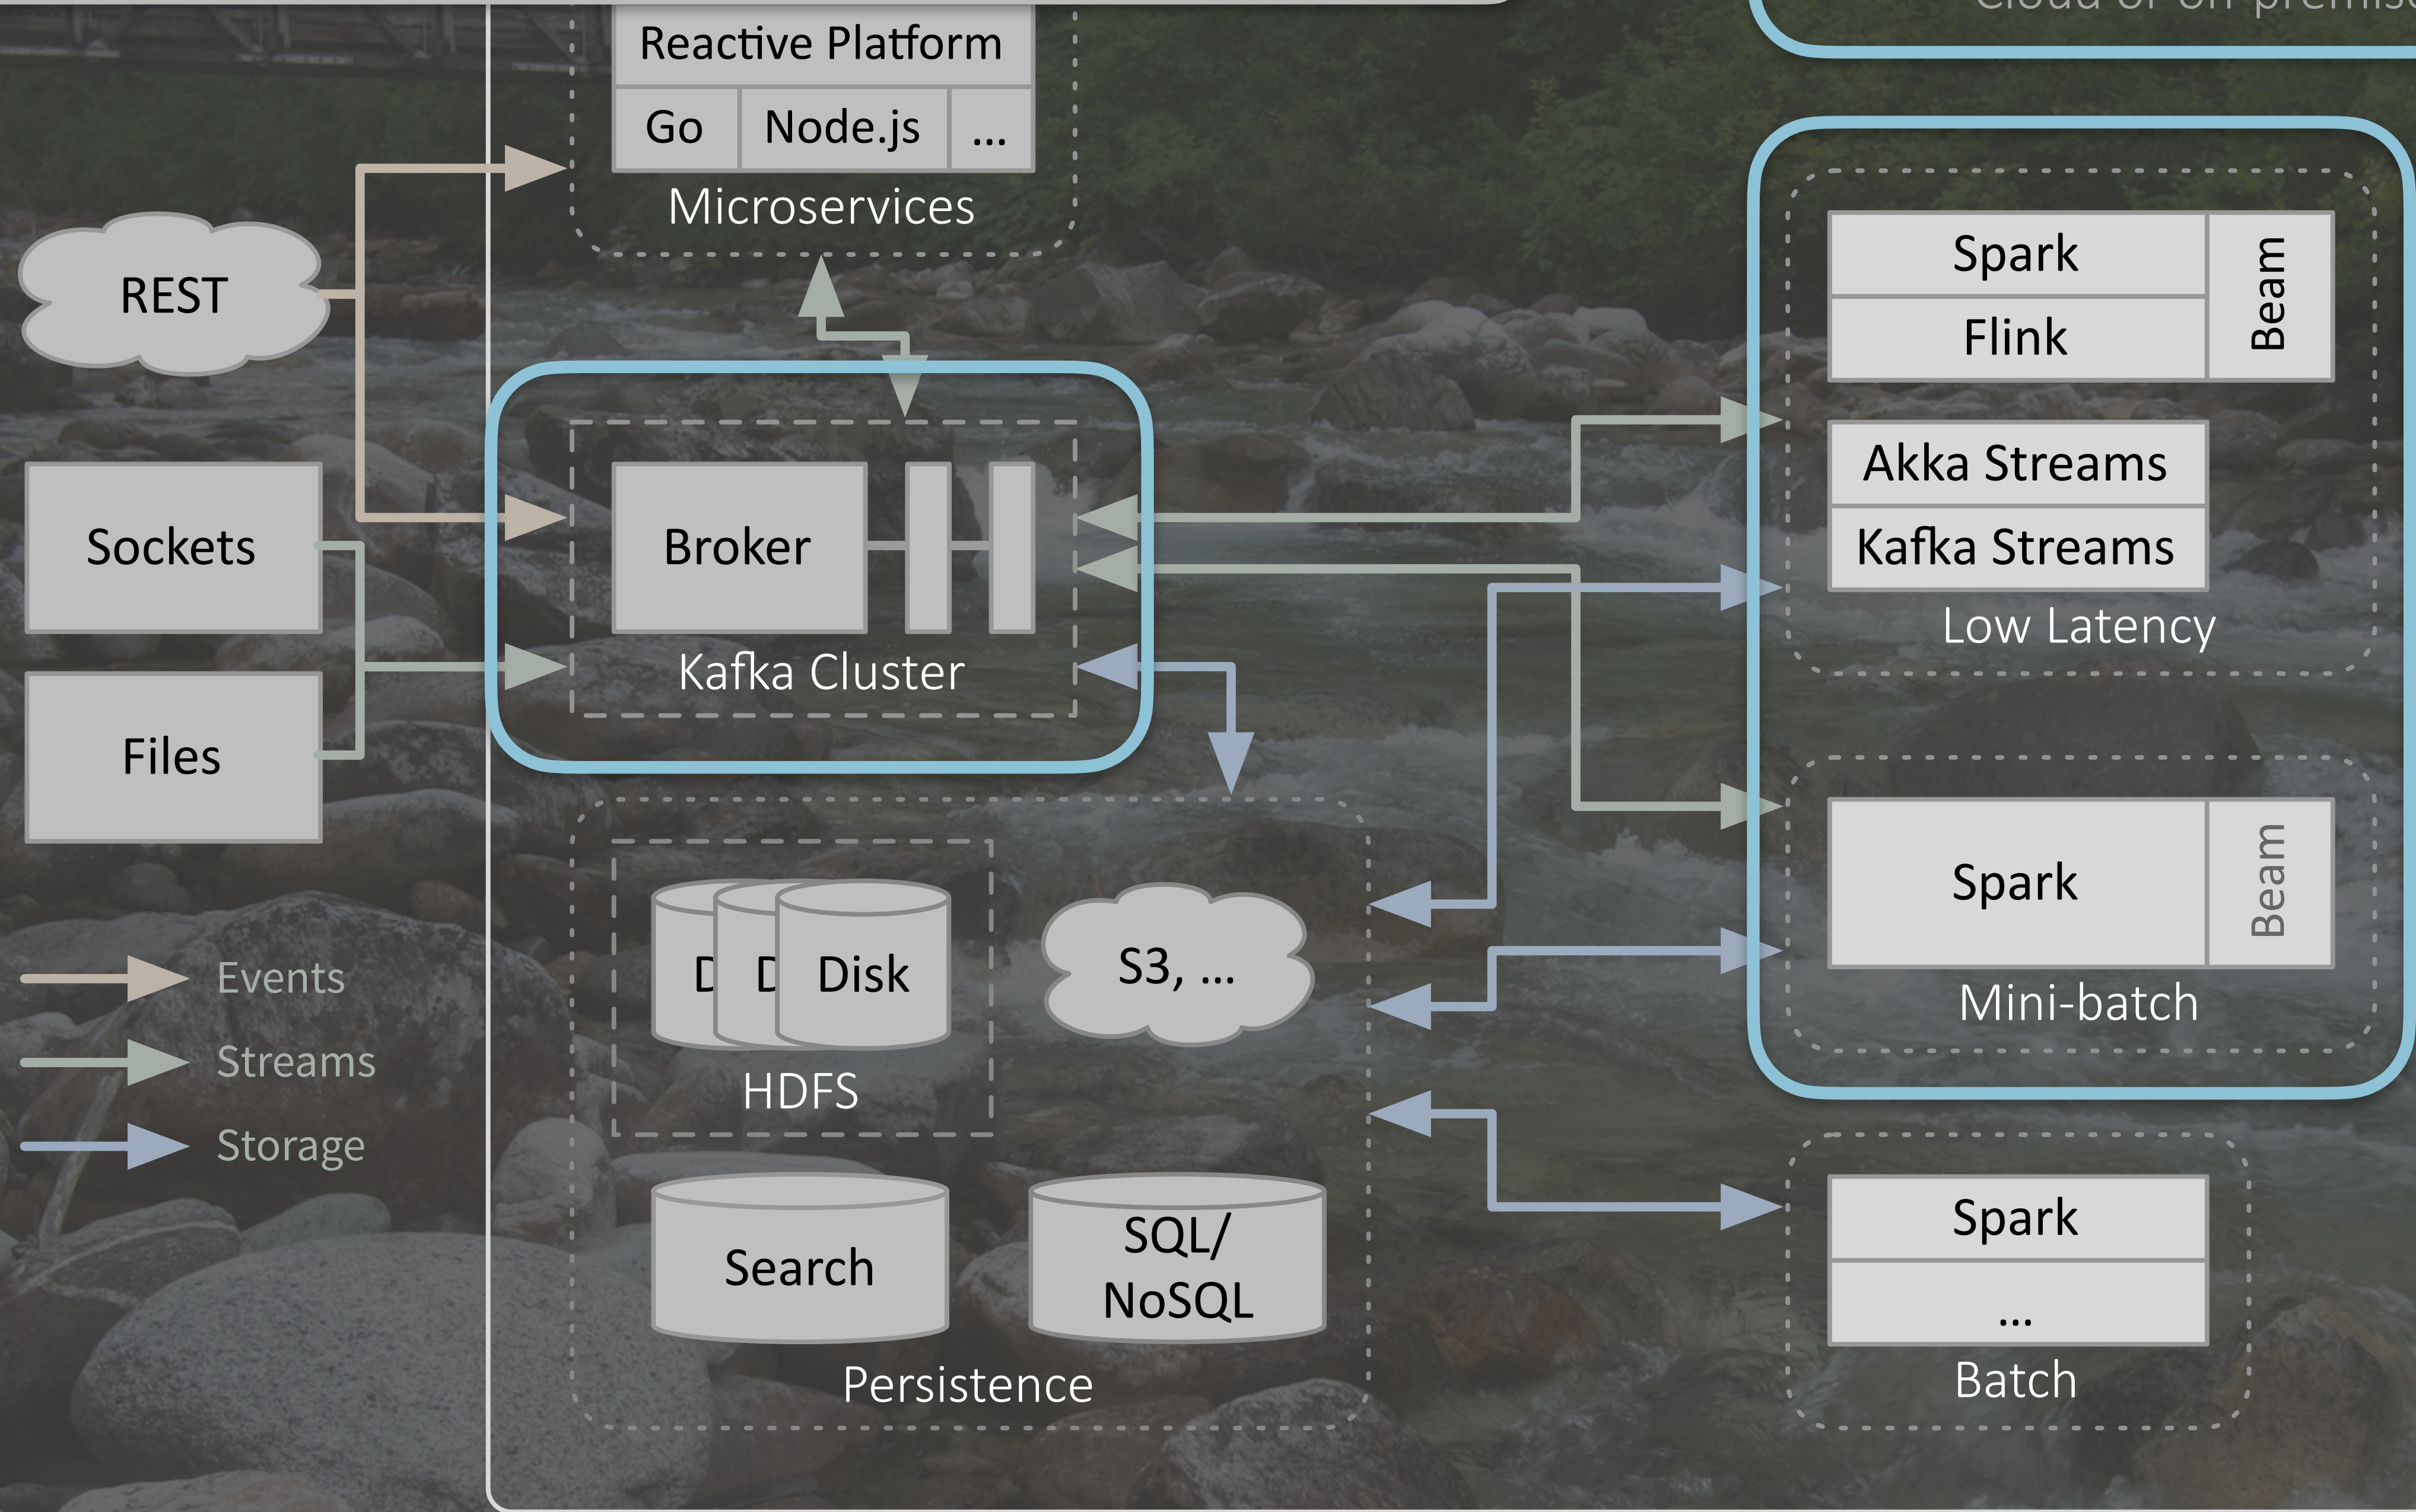


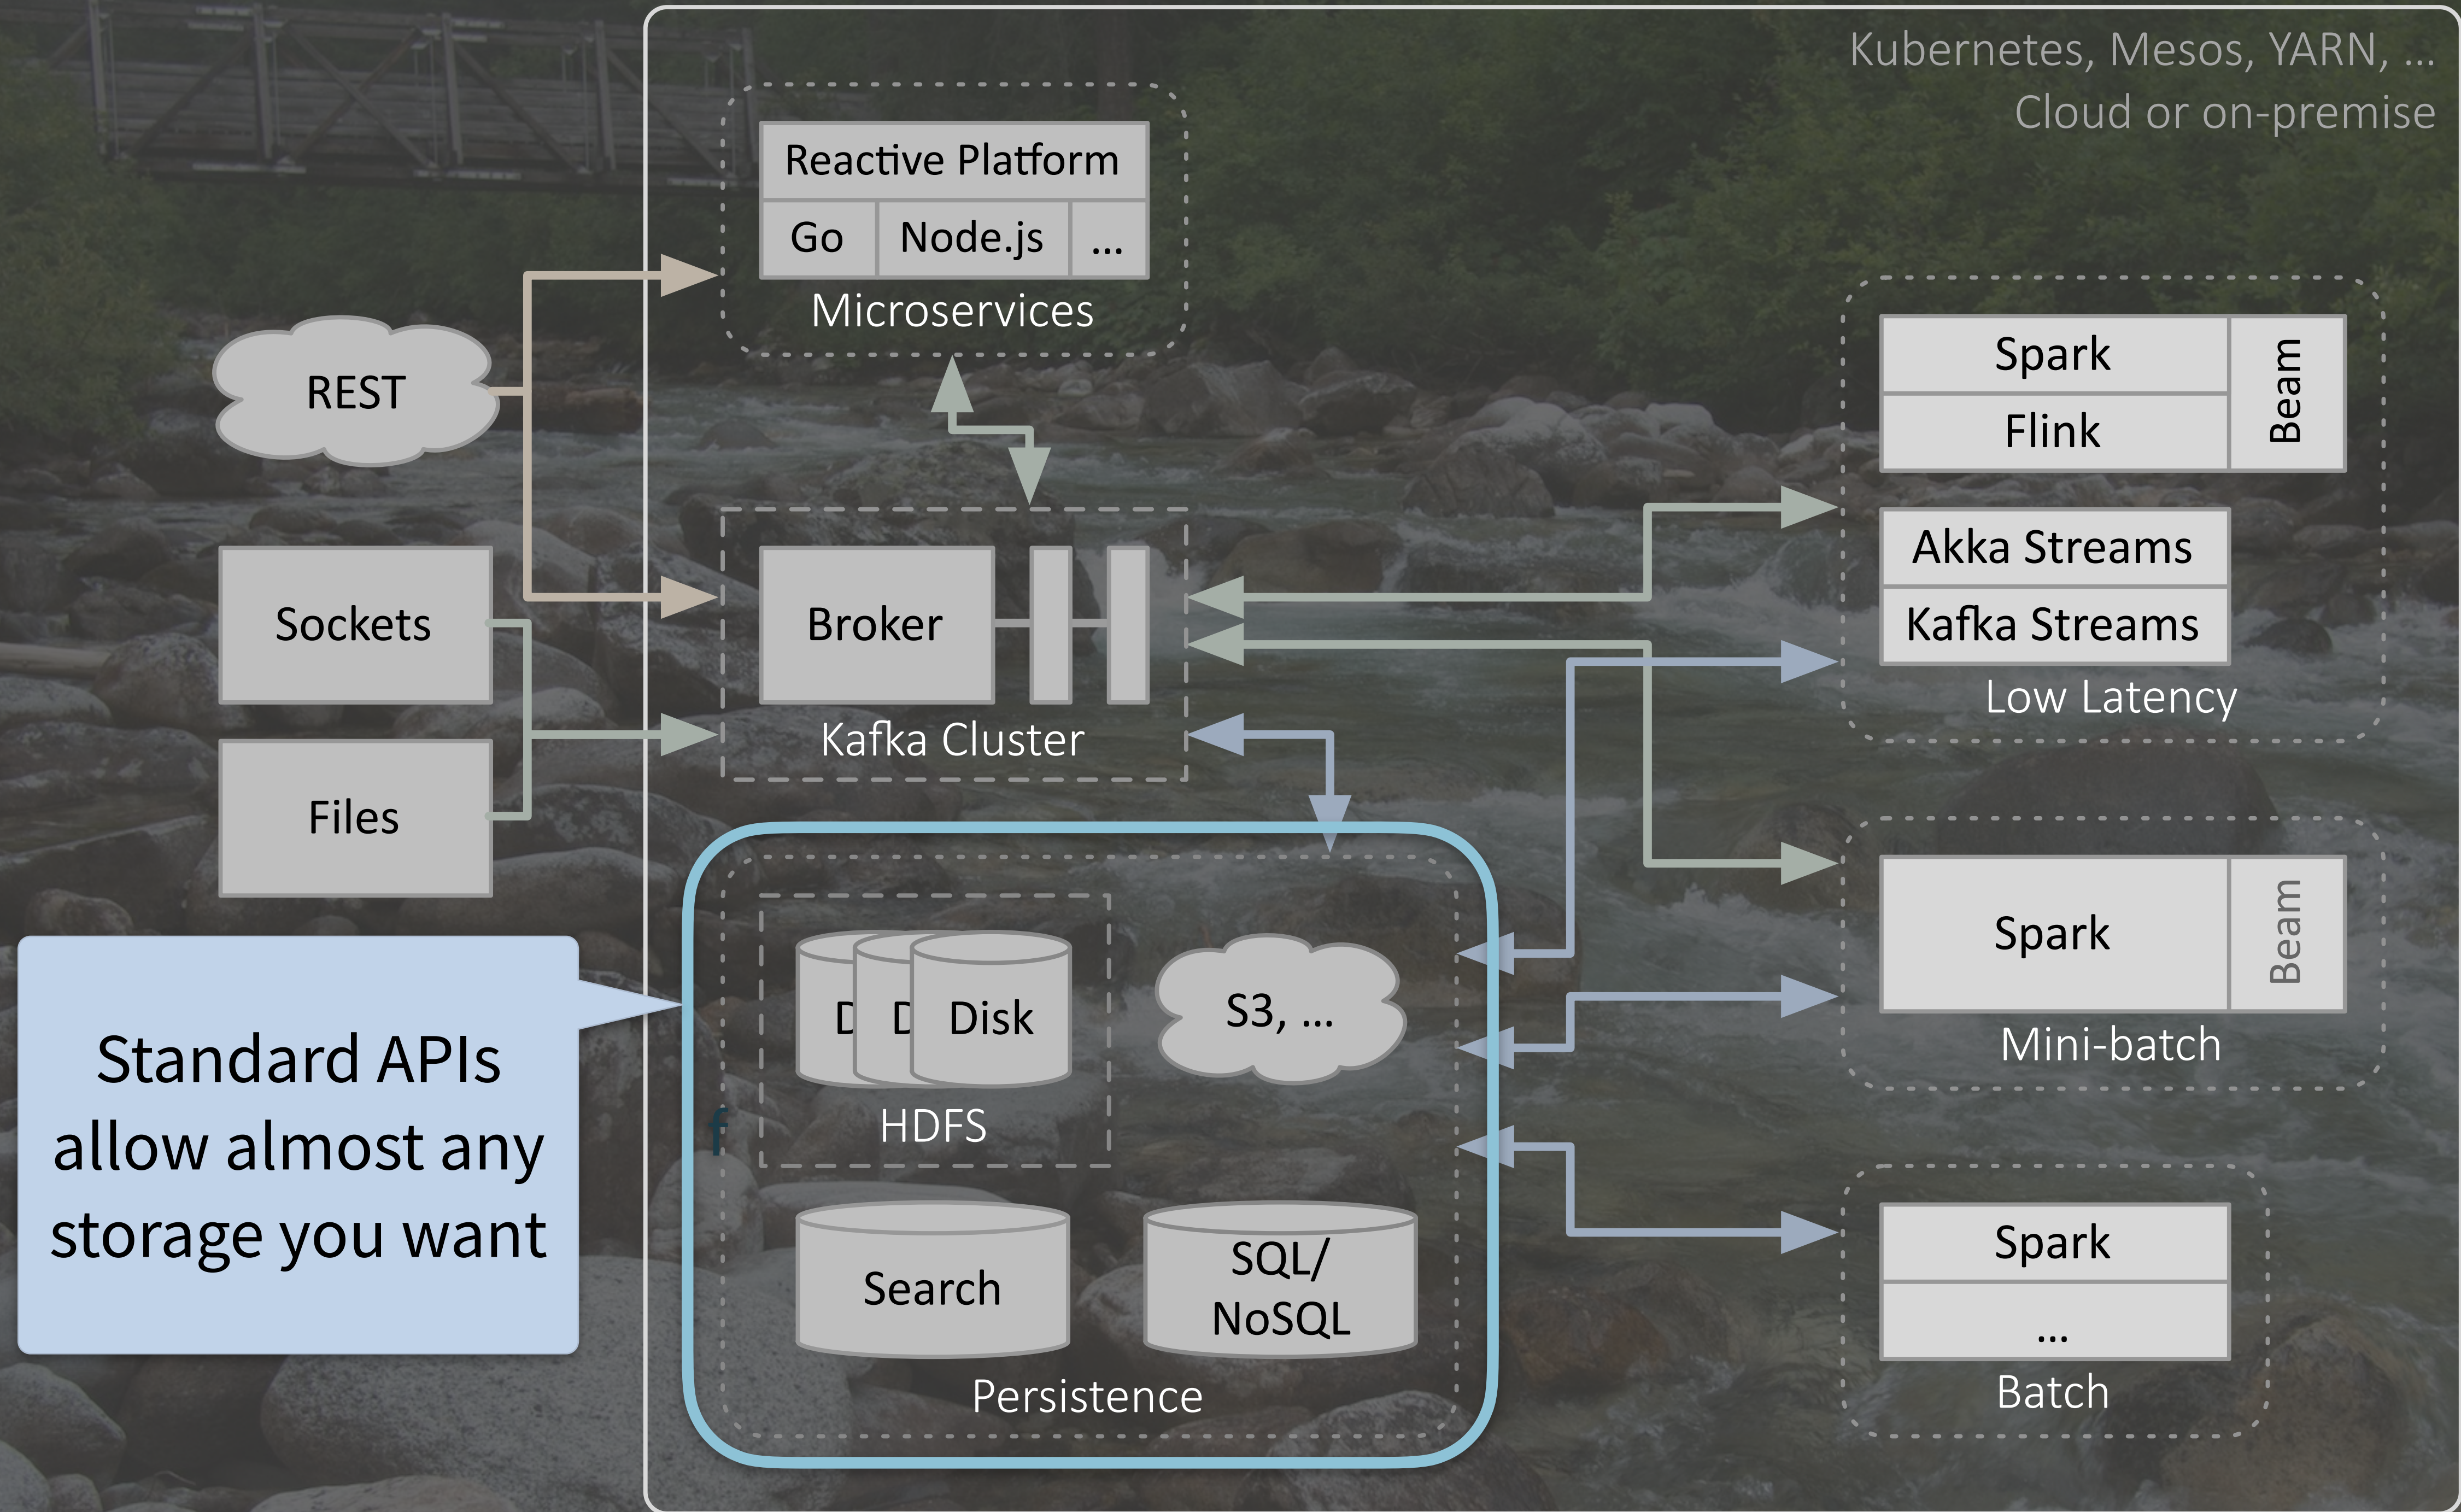


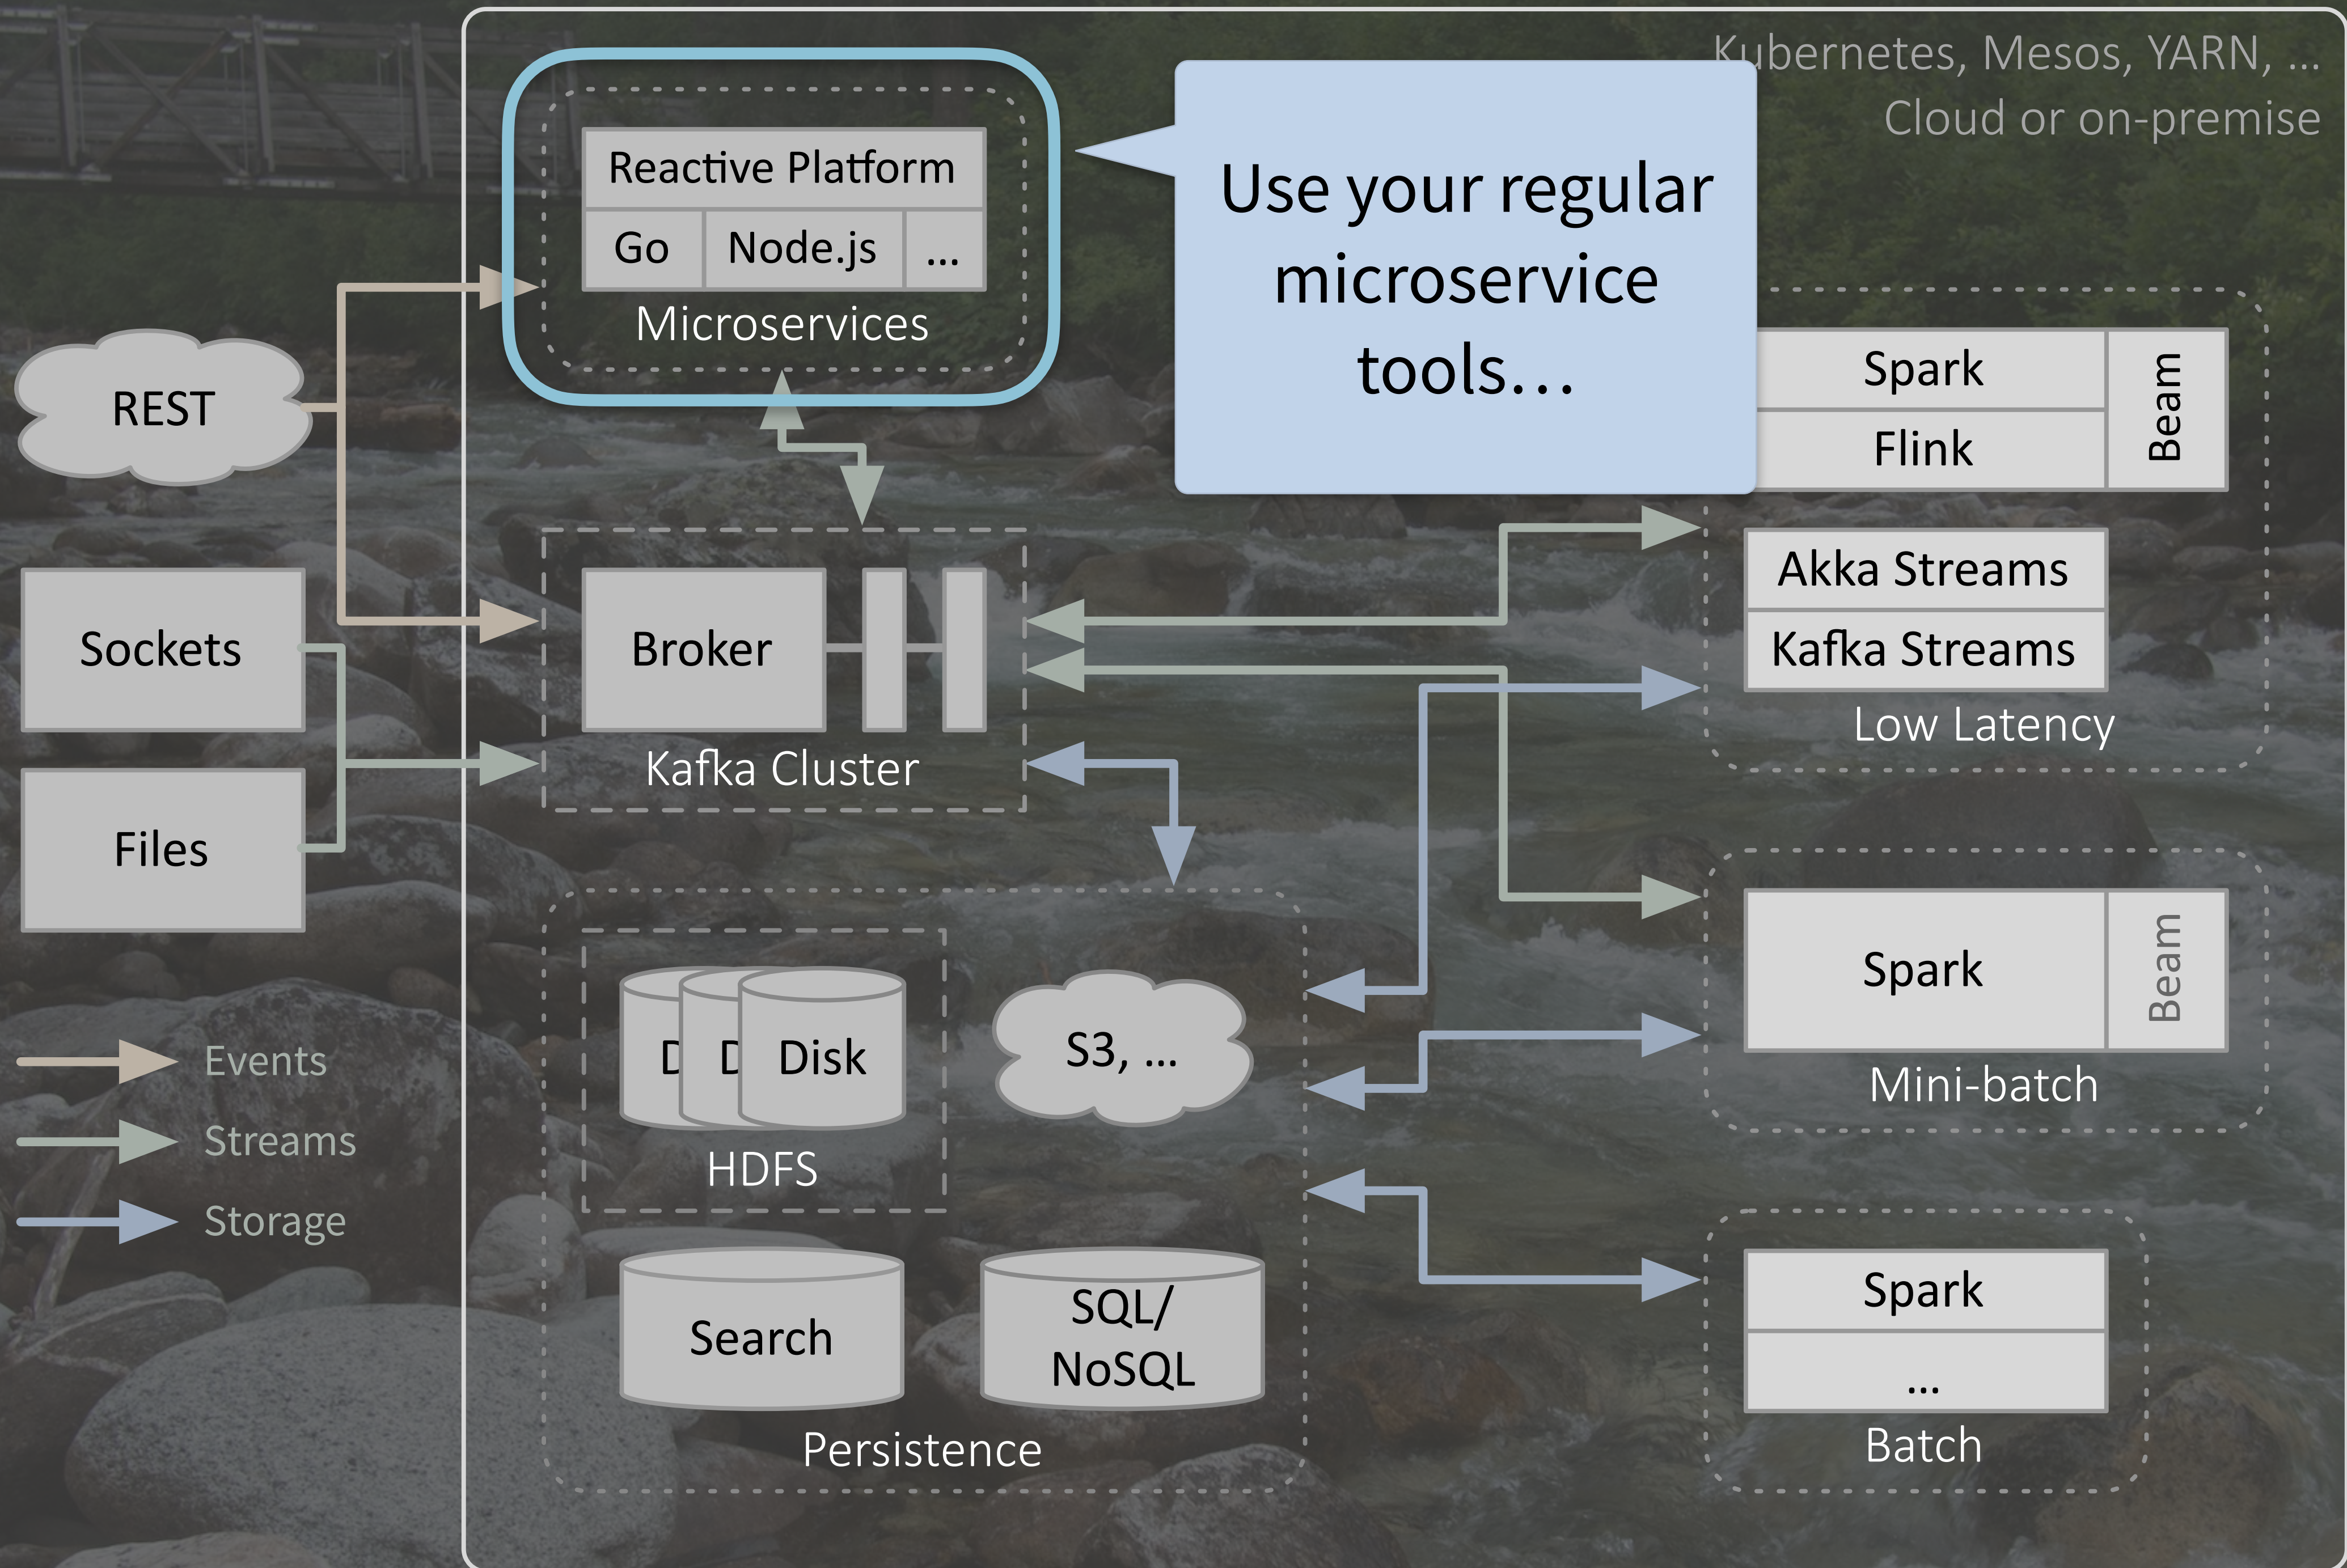
Lots of streaming engine options... too many.

The streaming analog of a deconstructed database!

Kubernetes, Mesos, YARN, ...
Cloud or on-premise






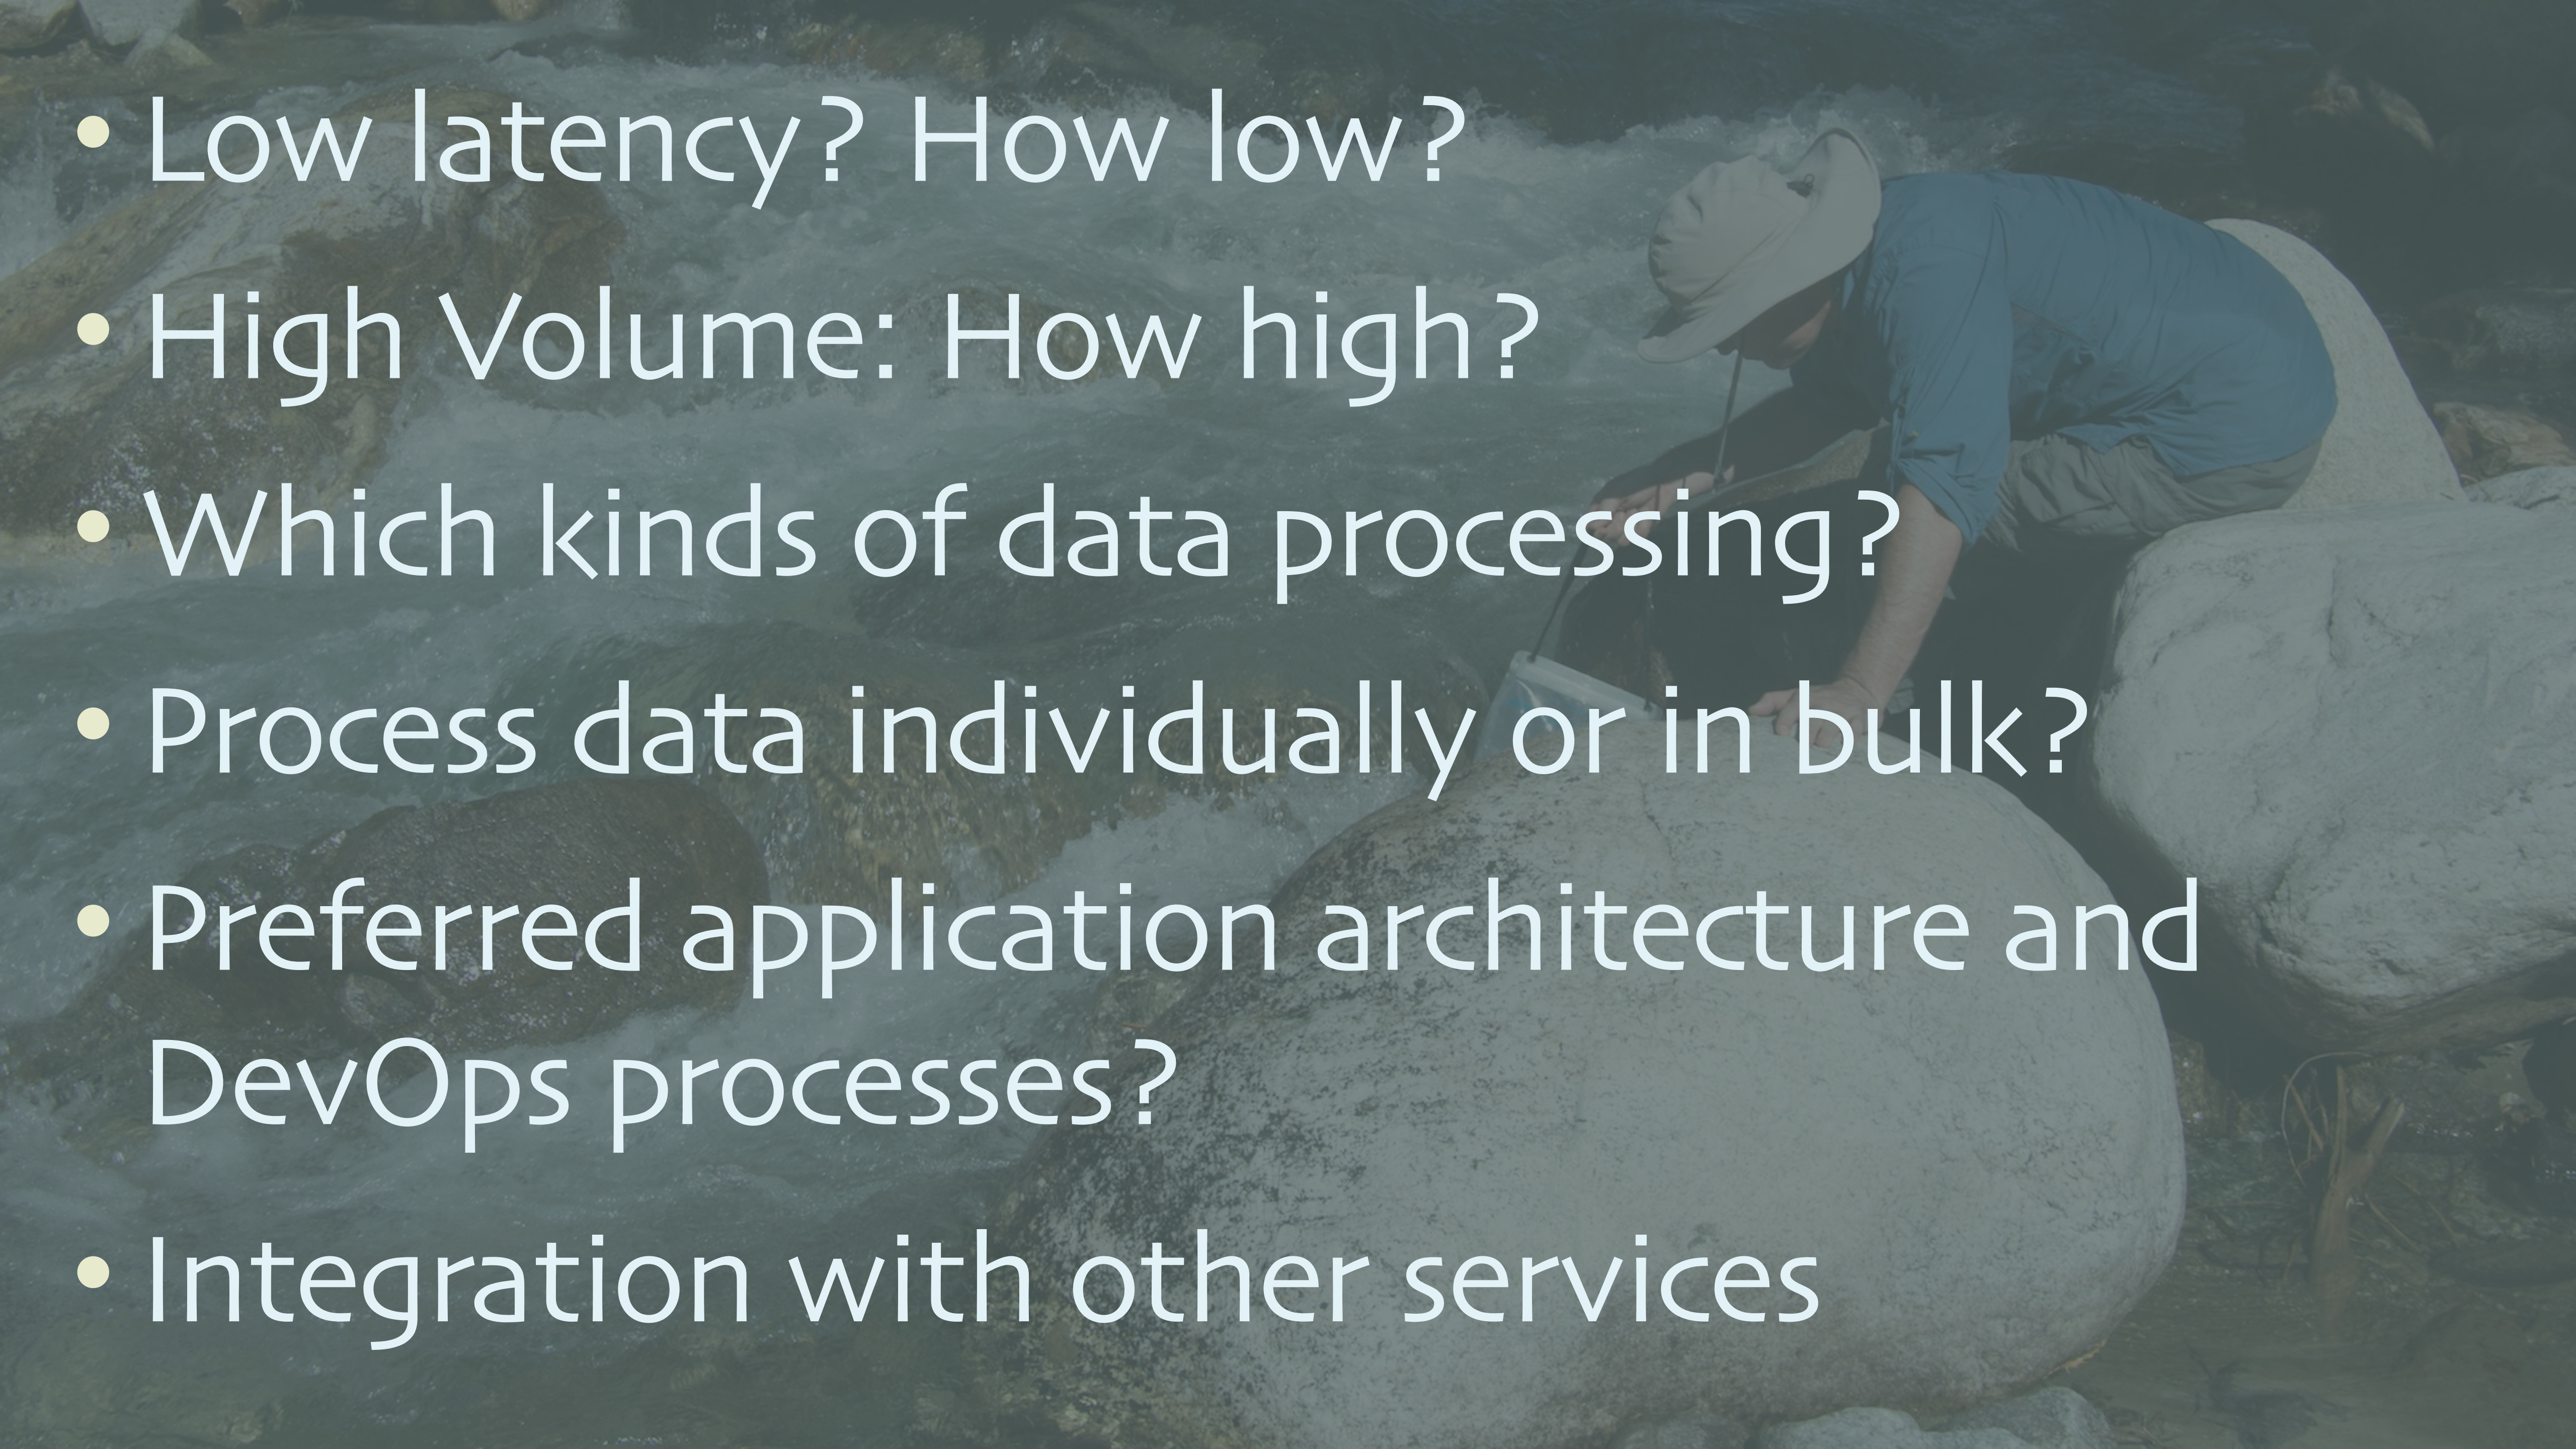




Streaming Engines

Features to Consider...



- 
- A person wearing a blue long-sleeved shirt, a light-colored hat, and dark pants is crouching on a large, smooth, light-colored rock. They are looking down at a small electronic device in their hands. The background shows a rocky stream with water flowing over the rocks. The image is overlaid with a semi-transparent dark blue filter.
- Low latency? How low?
 - High Volume: How high?
 - Which kinds of data processing?
 - Process data individually or in bulk?
 - Preferred application architecture and DevOps processes?
 - Integration with other services

- Low latency? How low?

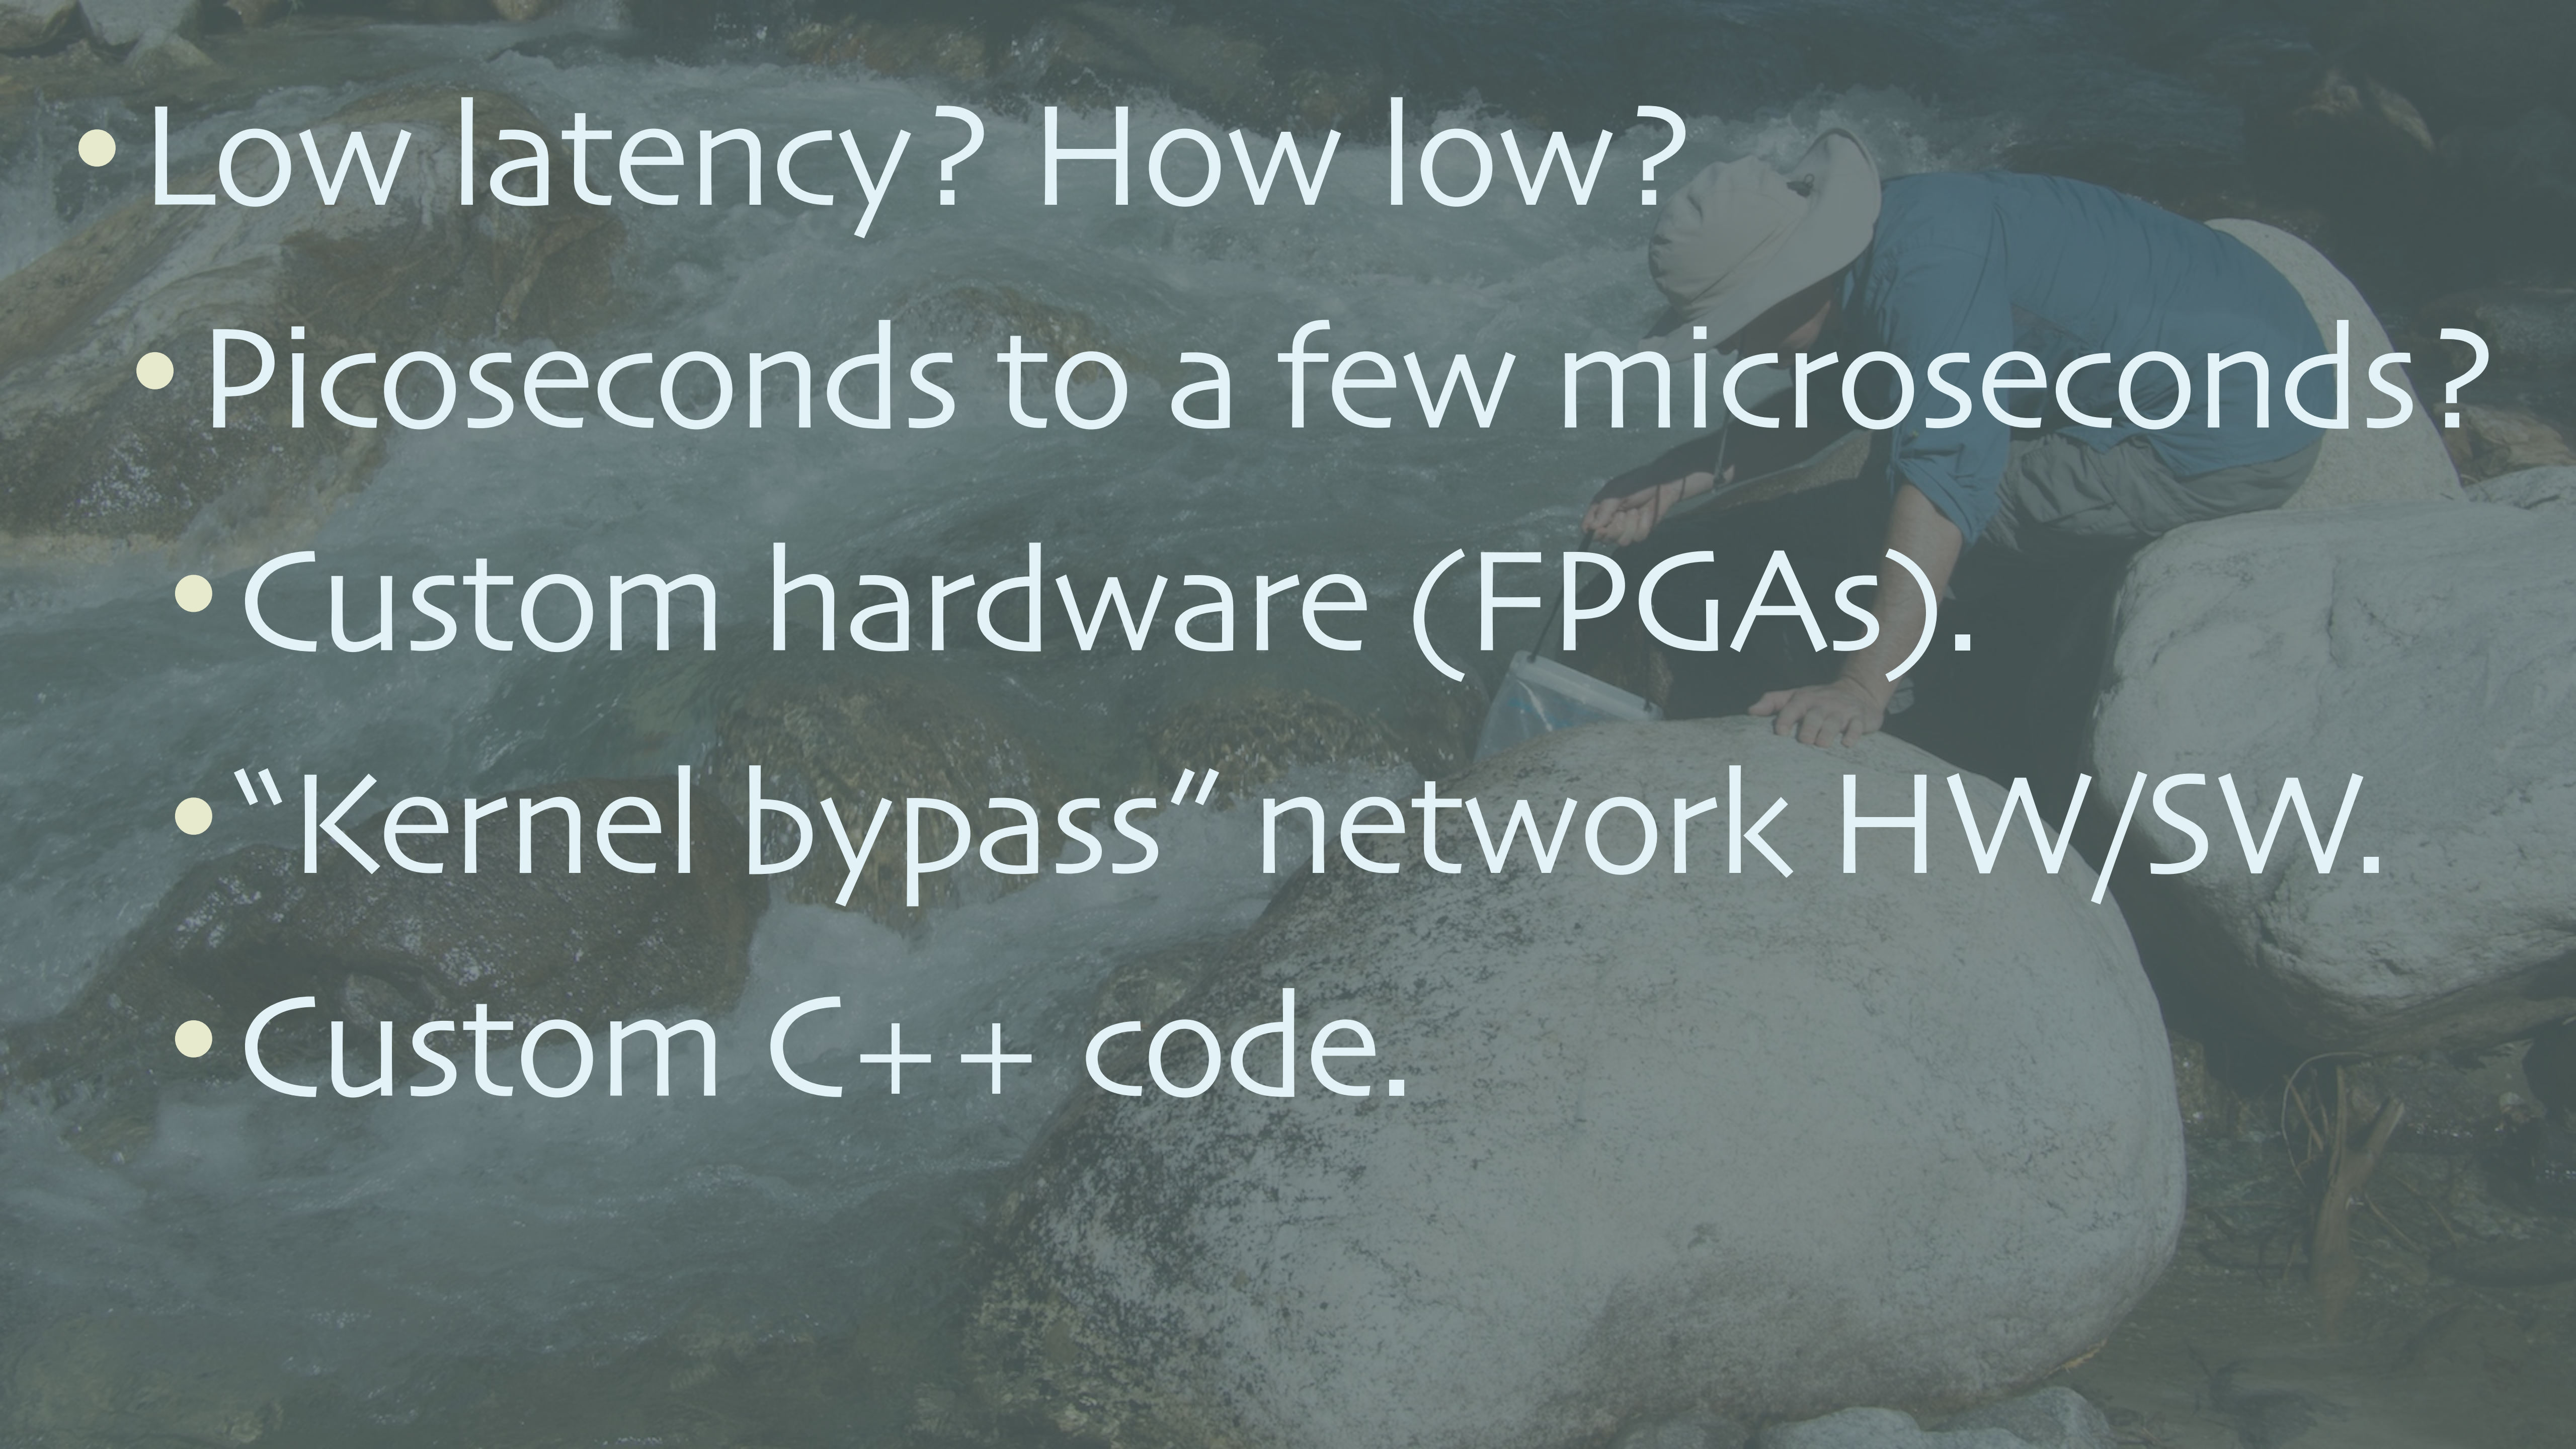


- Low latency? How low?
- Picoseconds to a few microseconds?




True "Real Time"

<http://www.spacex.com/news>

- 
- Low latency? How low?
 - Picoseconds to a few microseconds?
 - Custom hardware (FPGAs).
 - “Kernel bypass” network HW/SW.
 - Custom C++ code.

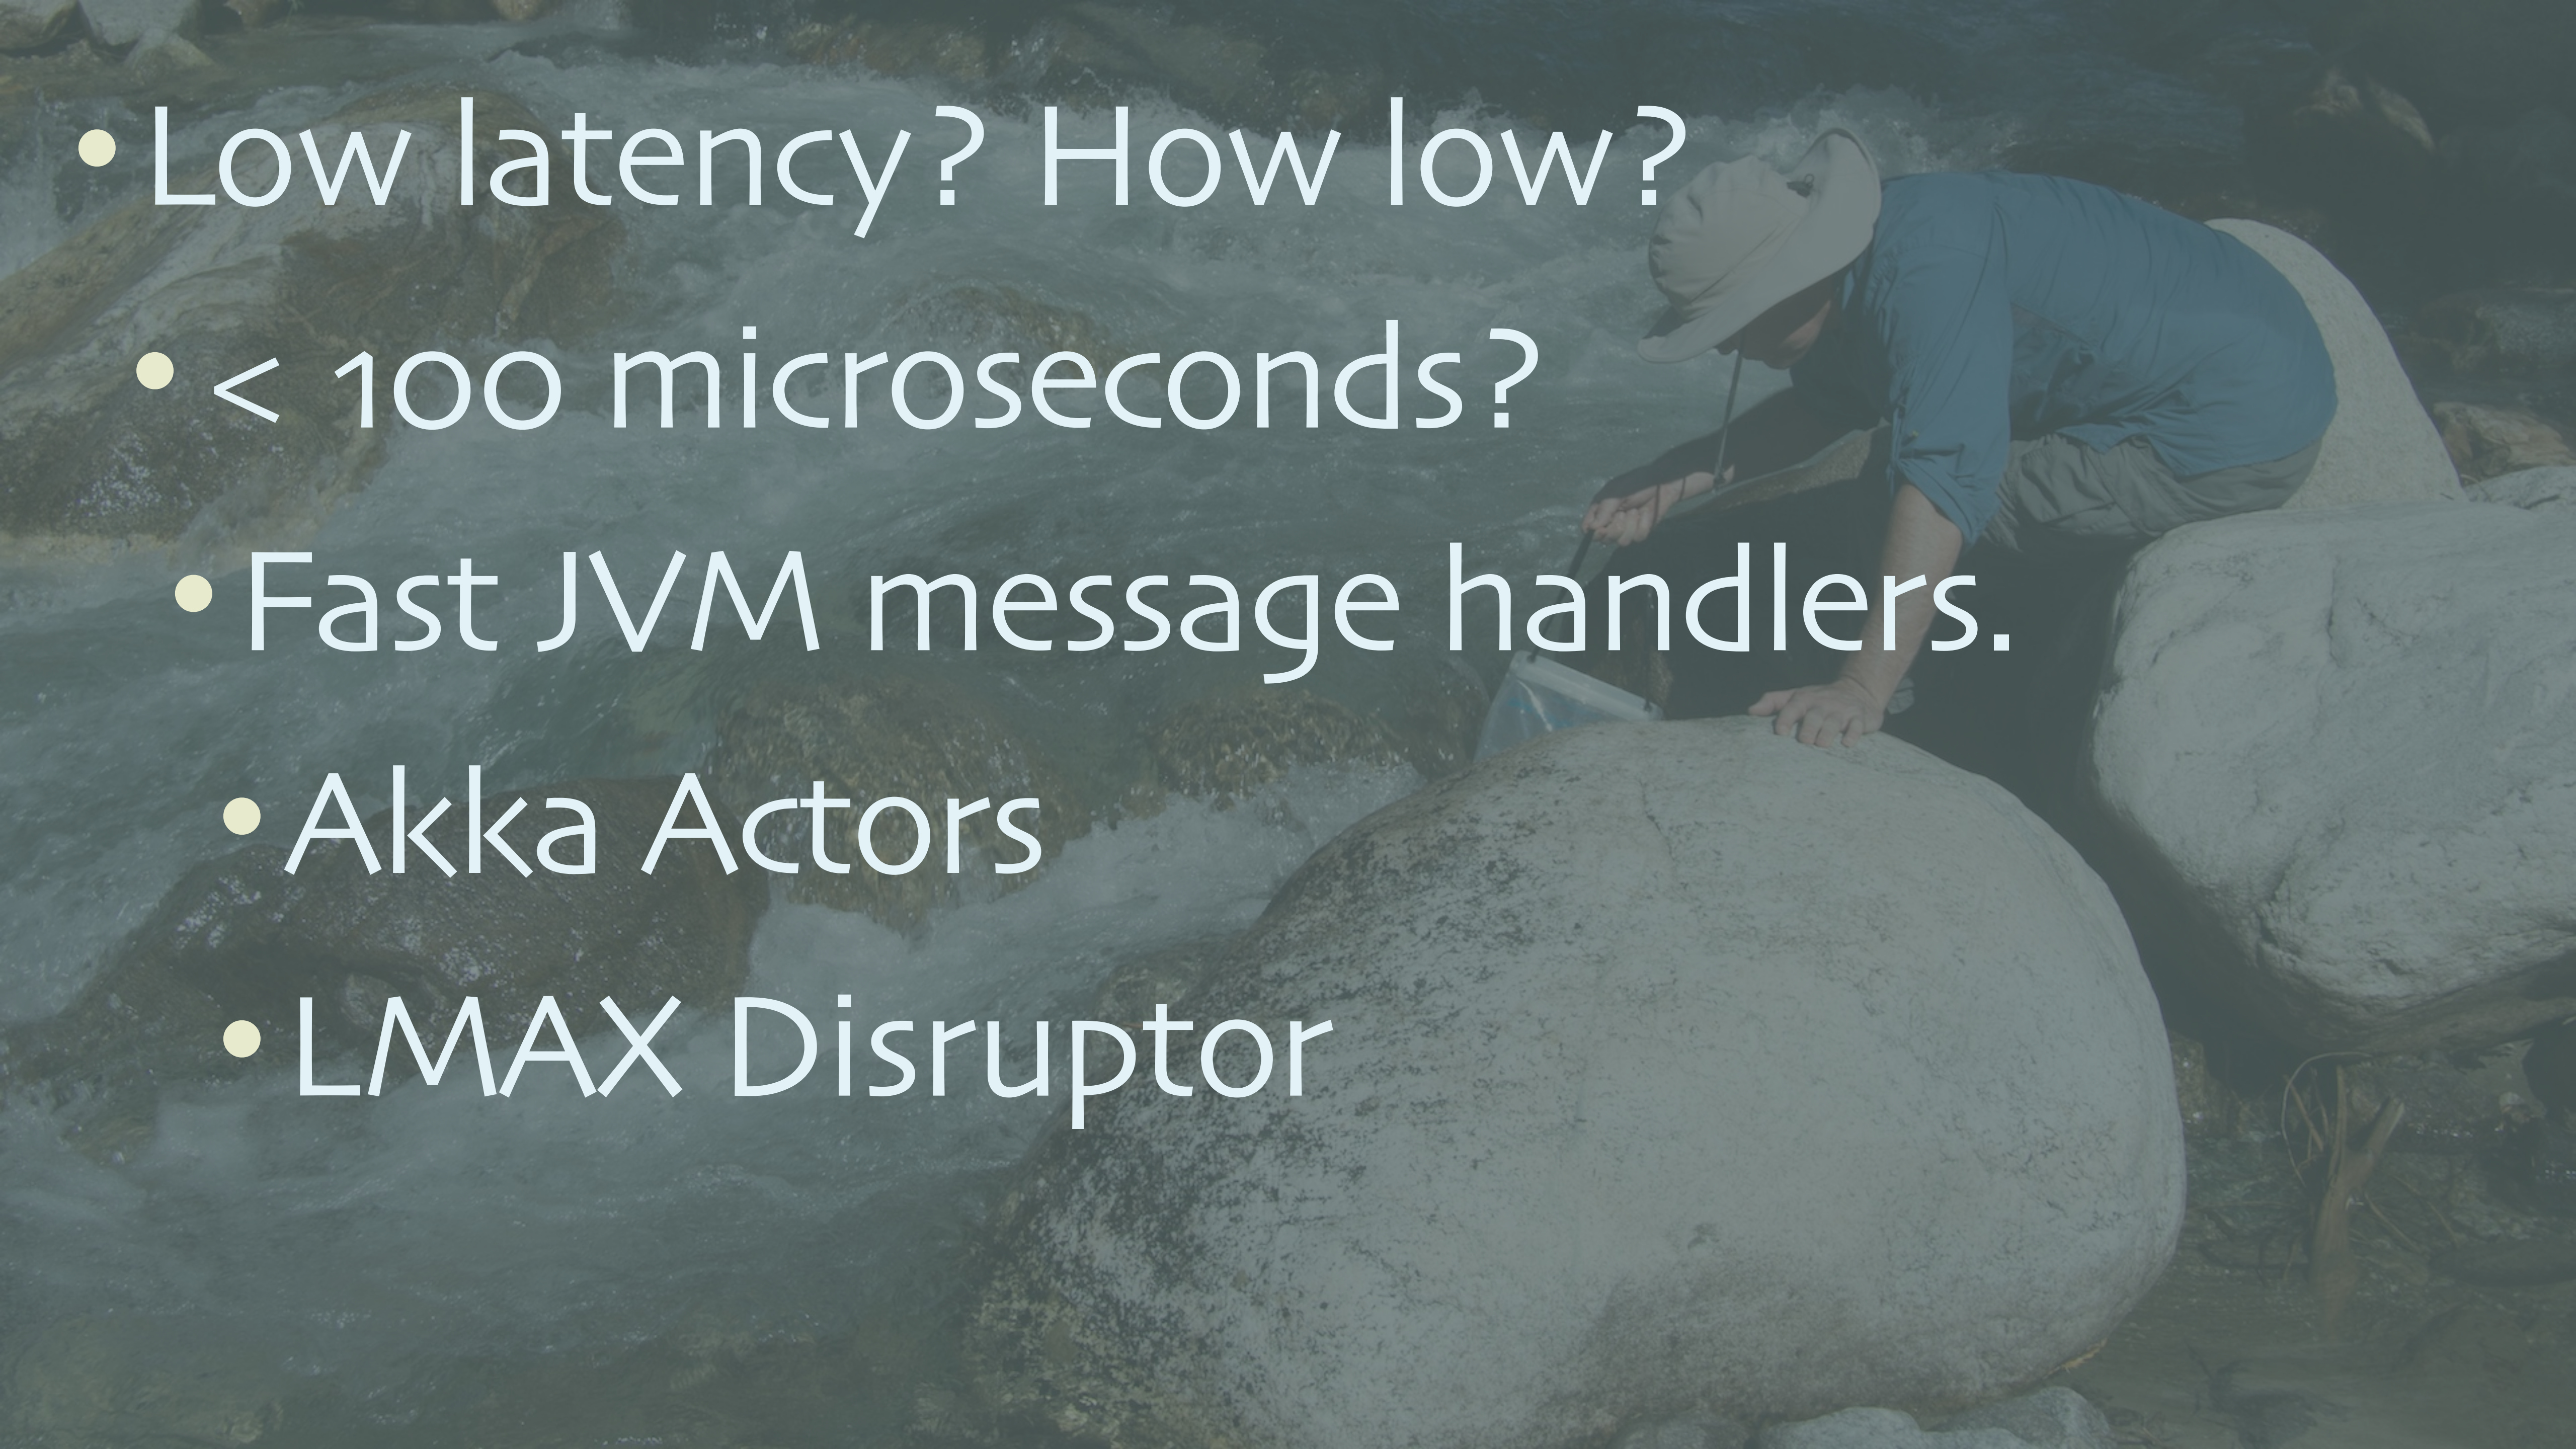
- Low latency? How low?
- < 100 microseconds?



<http://tradinghub.co/watch-list-for-mar-26th-2015/>

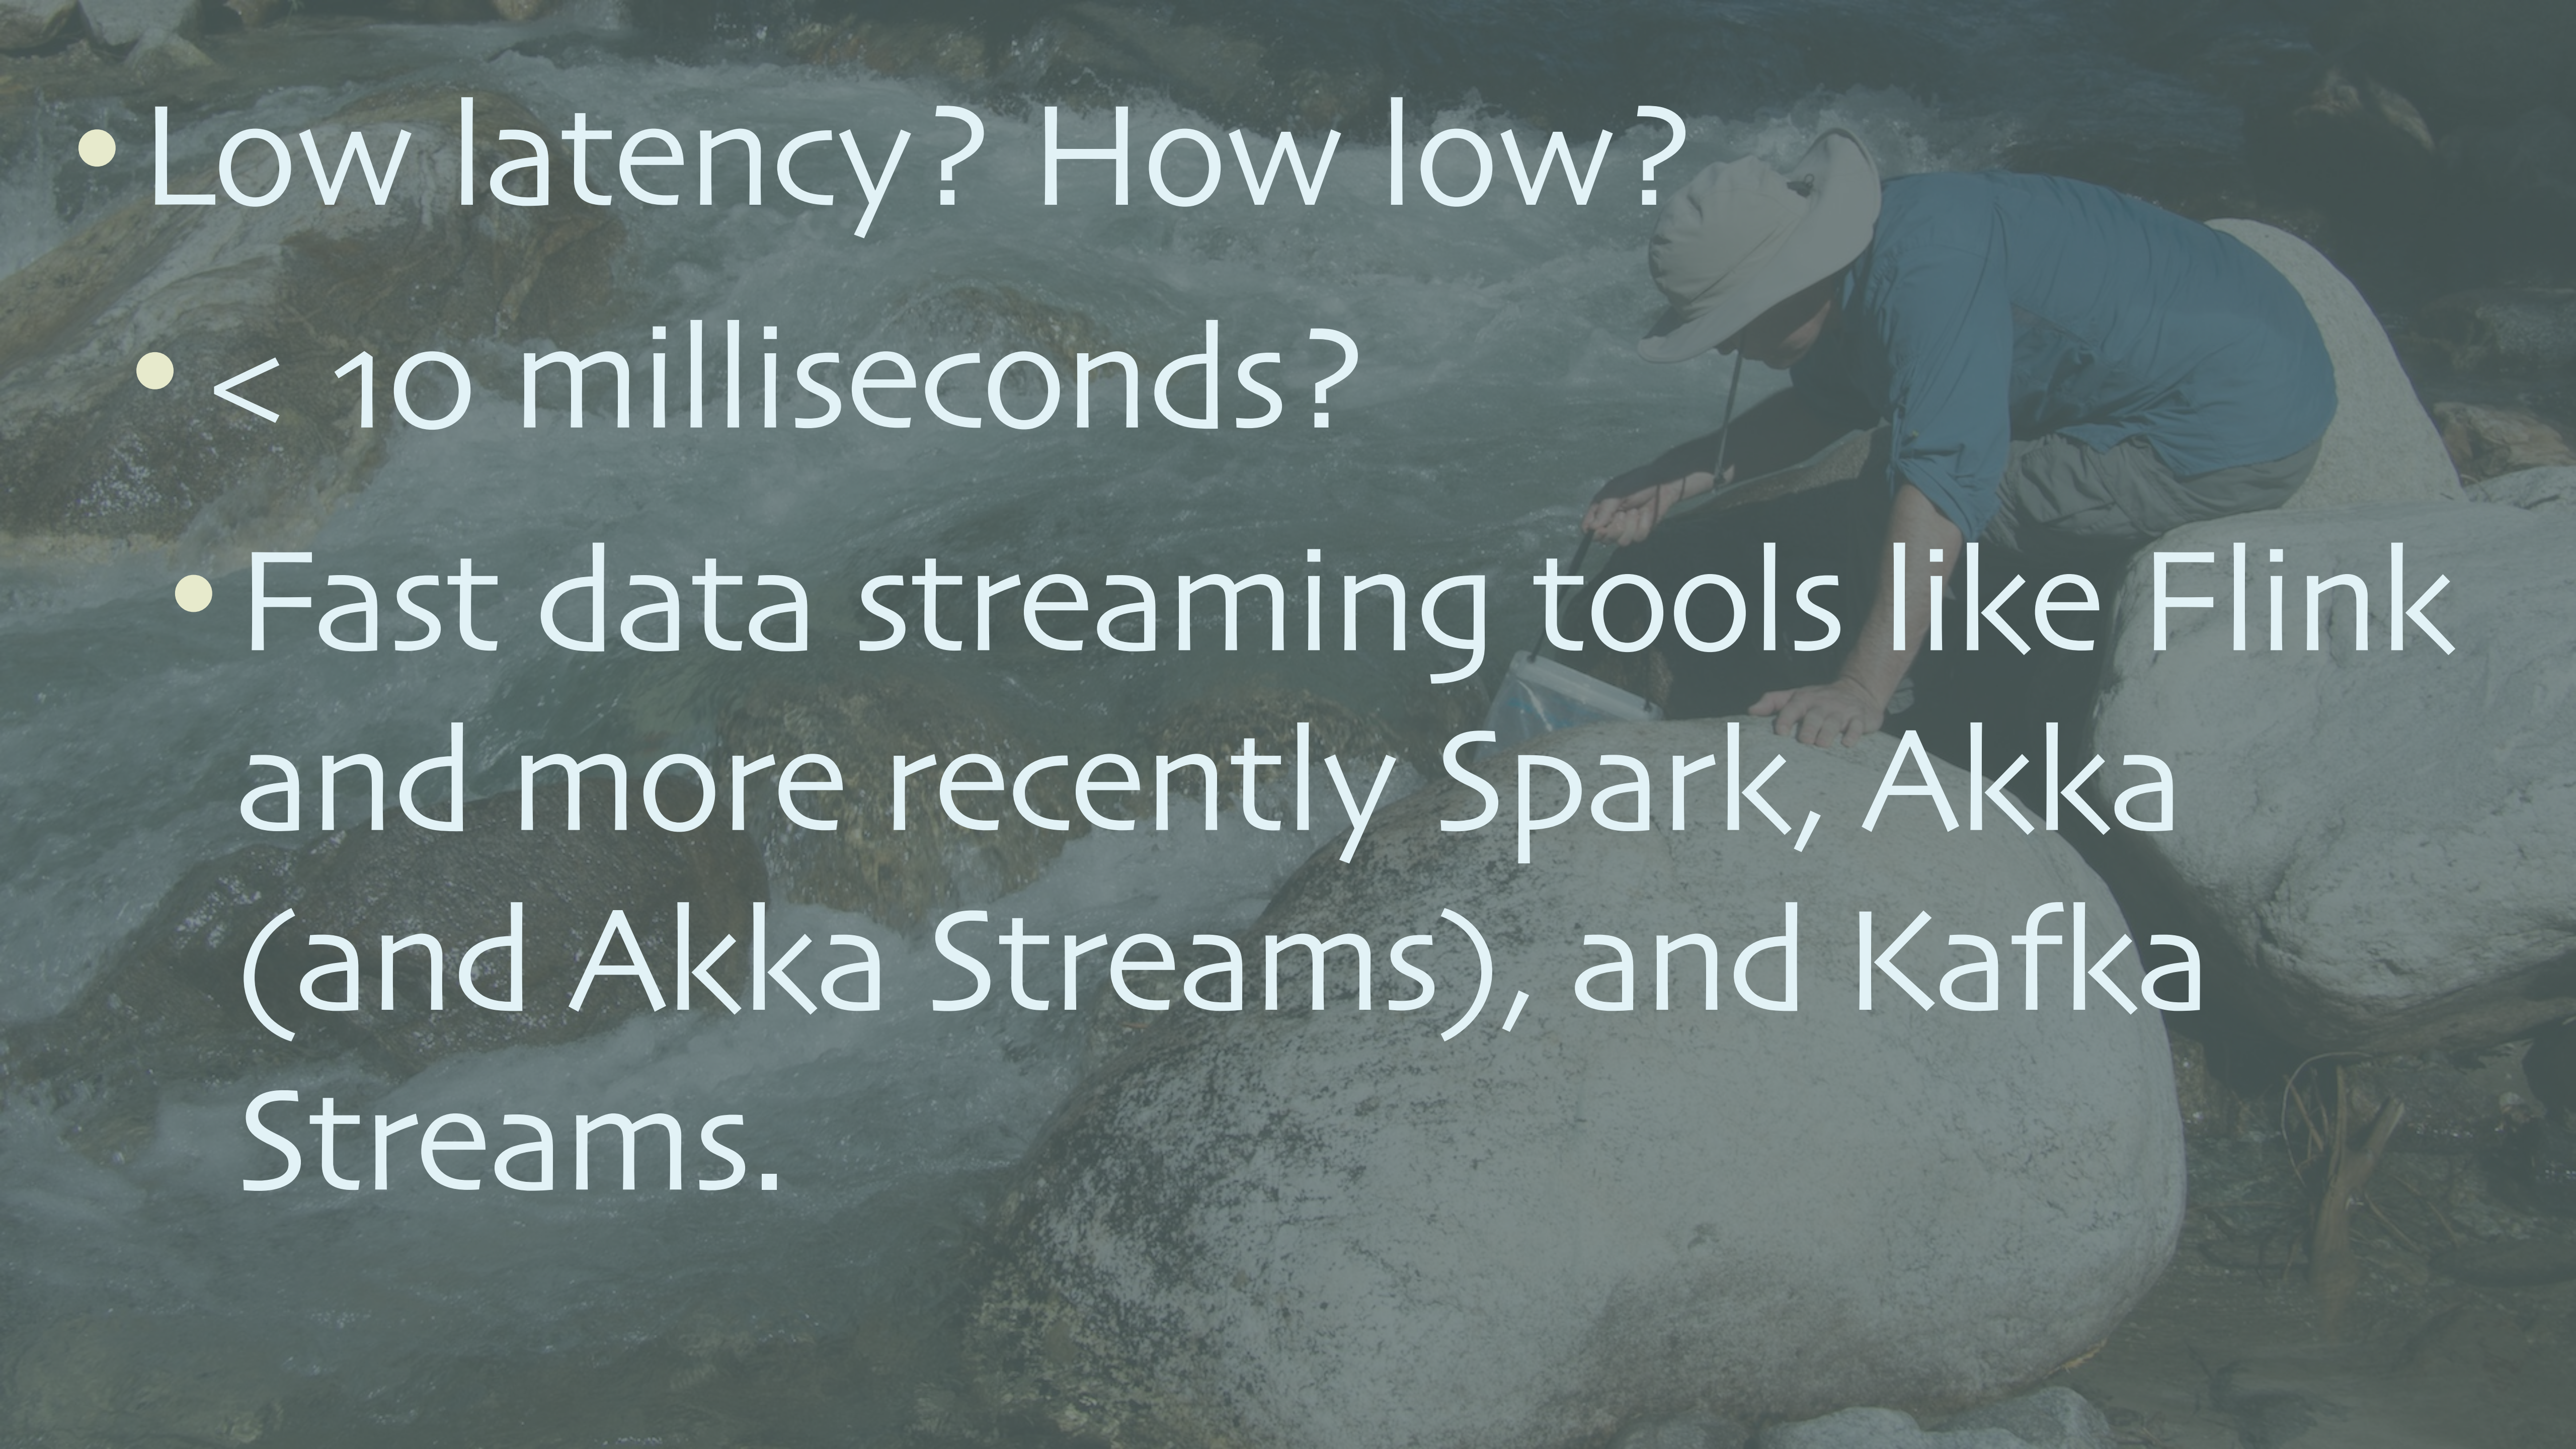


<http://www.usa.philips.com/>


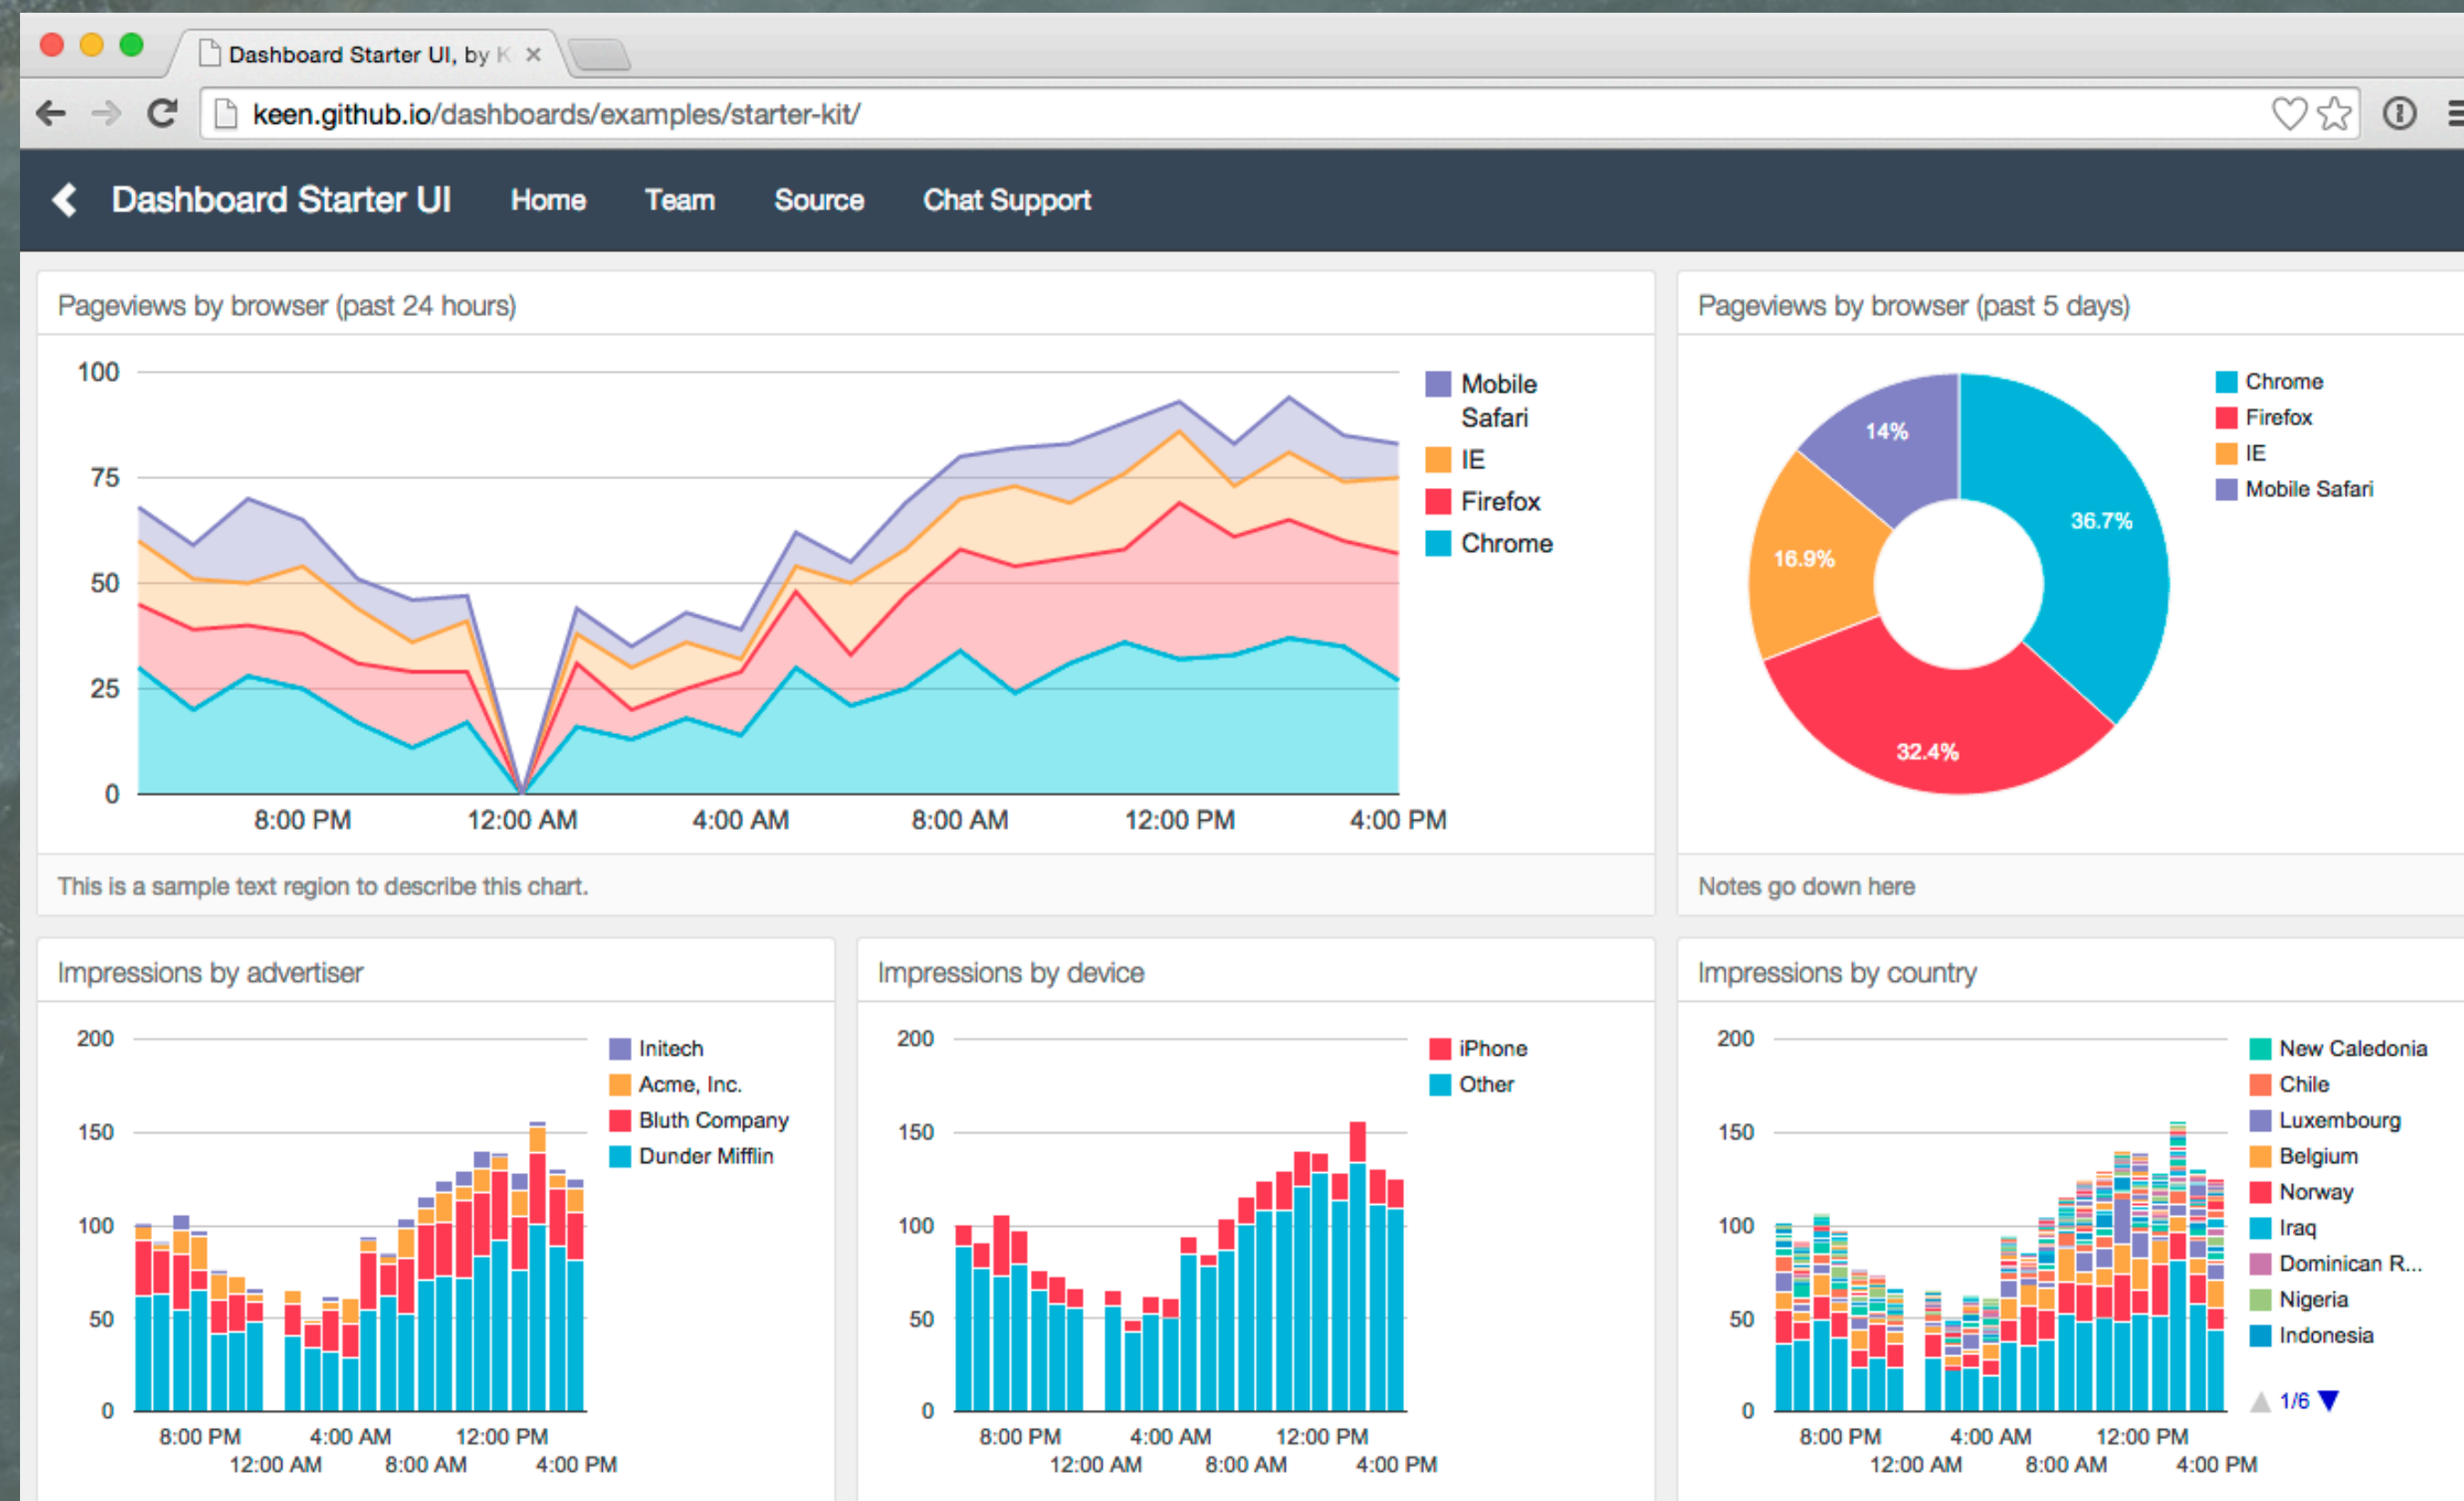
- 
- Low latency? How low?
 - < 100 microseconds?
 - Fast JVM message handlers.
 - Akka Actors
 - LMAX Disruptor

- Low latency? How low?
- < 10 milliseconds?



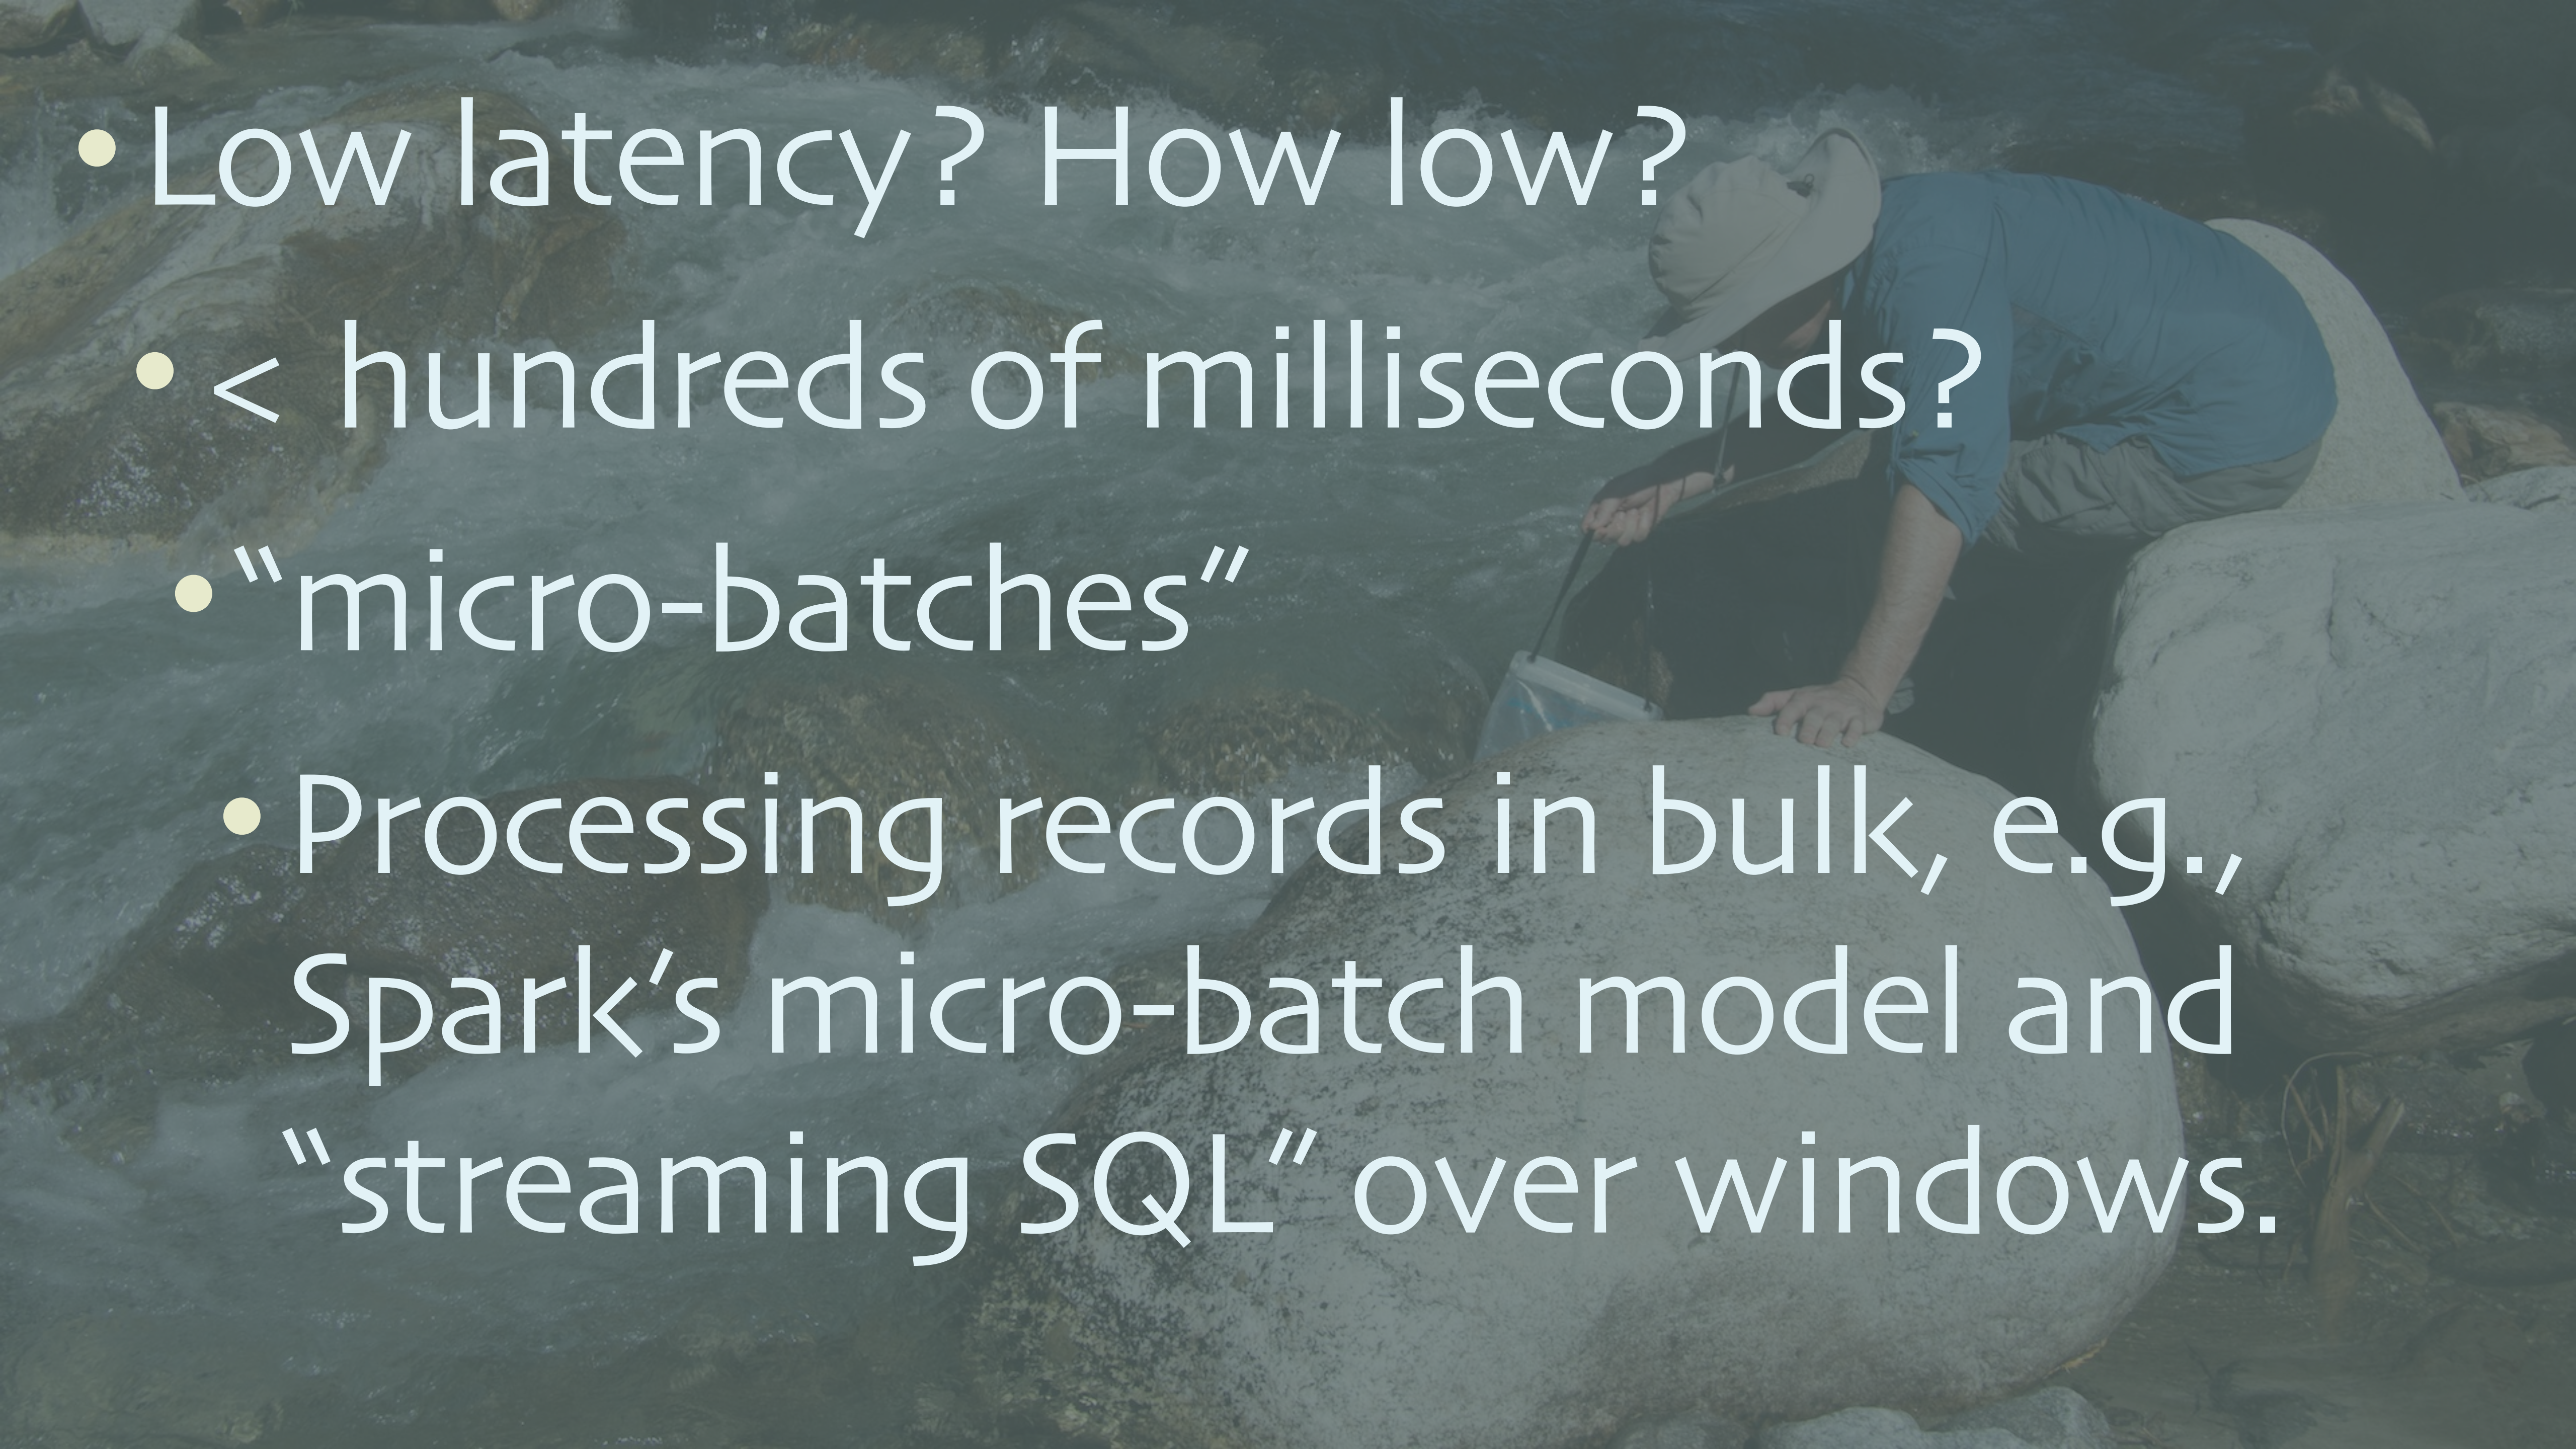
- 
- A person wearing a blue long-sleeved shirt, a light-colored hat, and dark pants is crouching on a large, smooth, light-colored rock in a stream. They are holding a small, clear plastic container in their hands. The water is flowing over the rocks, creating white foam. The background is a blurred view of the stream and more rocks.
- Low latency? How low?
 - < 10 milliseconds?
 - Fast data streaming tools like Flink and more recently Spark, Akka (and Akka Streams), and Kafka Streams.

- Low latency? How low?
- < hundreds of milliseconds?

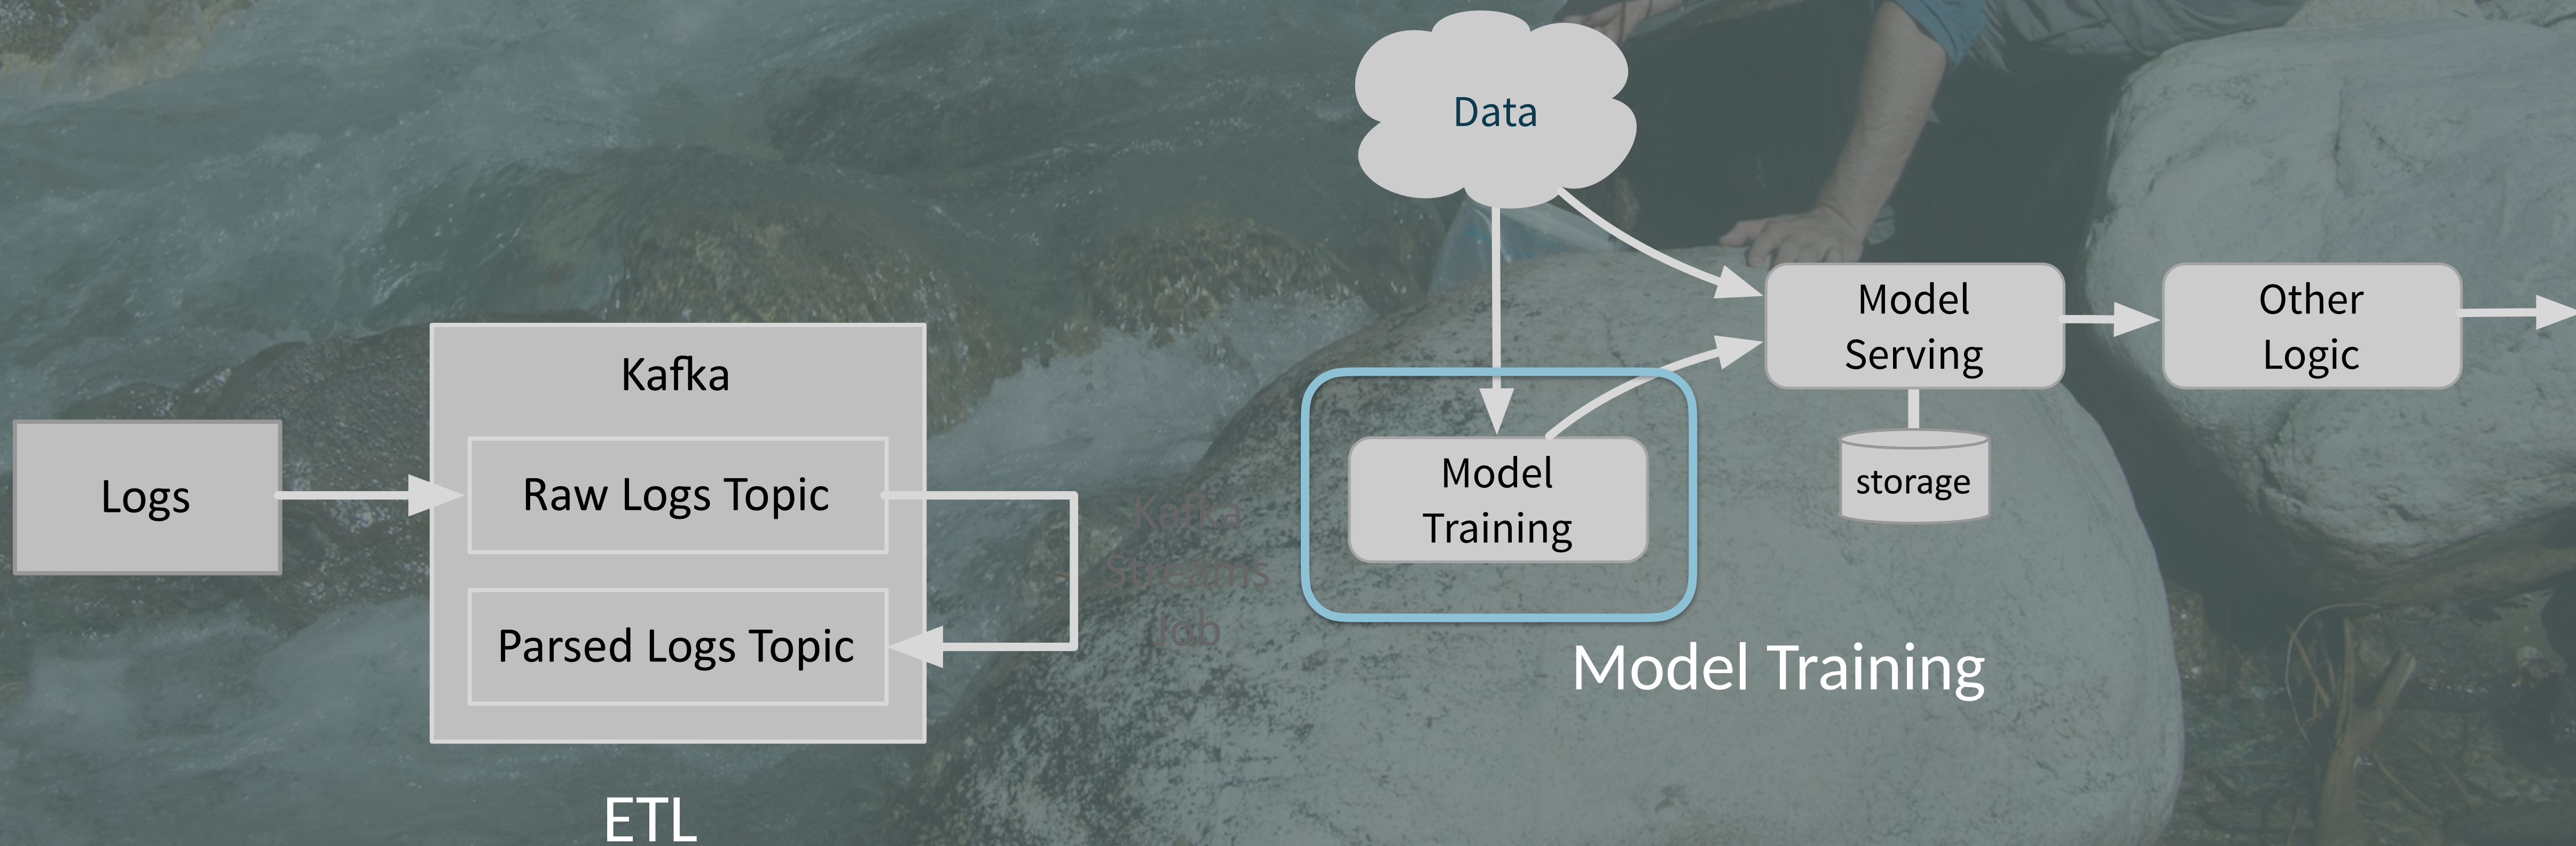


<https://www.coursera.org/learn/machine-learning>

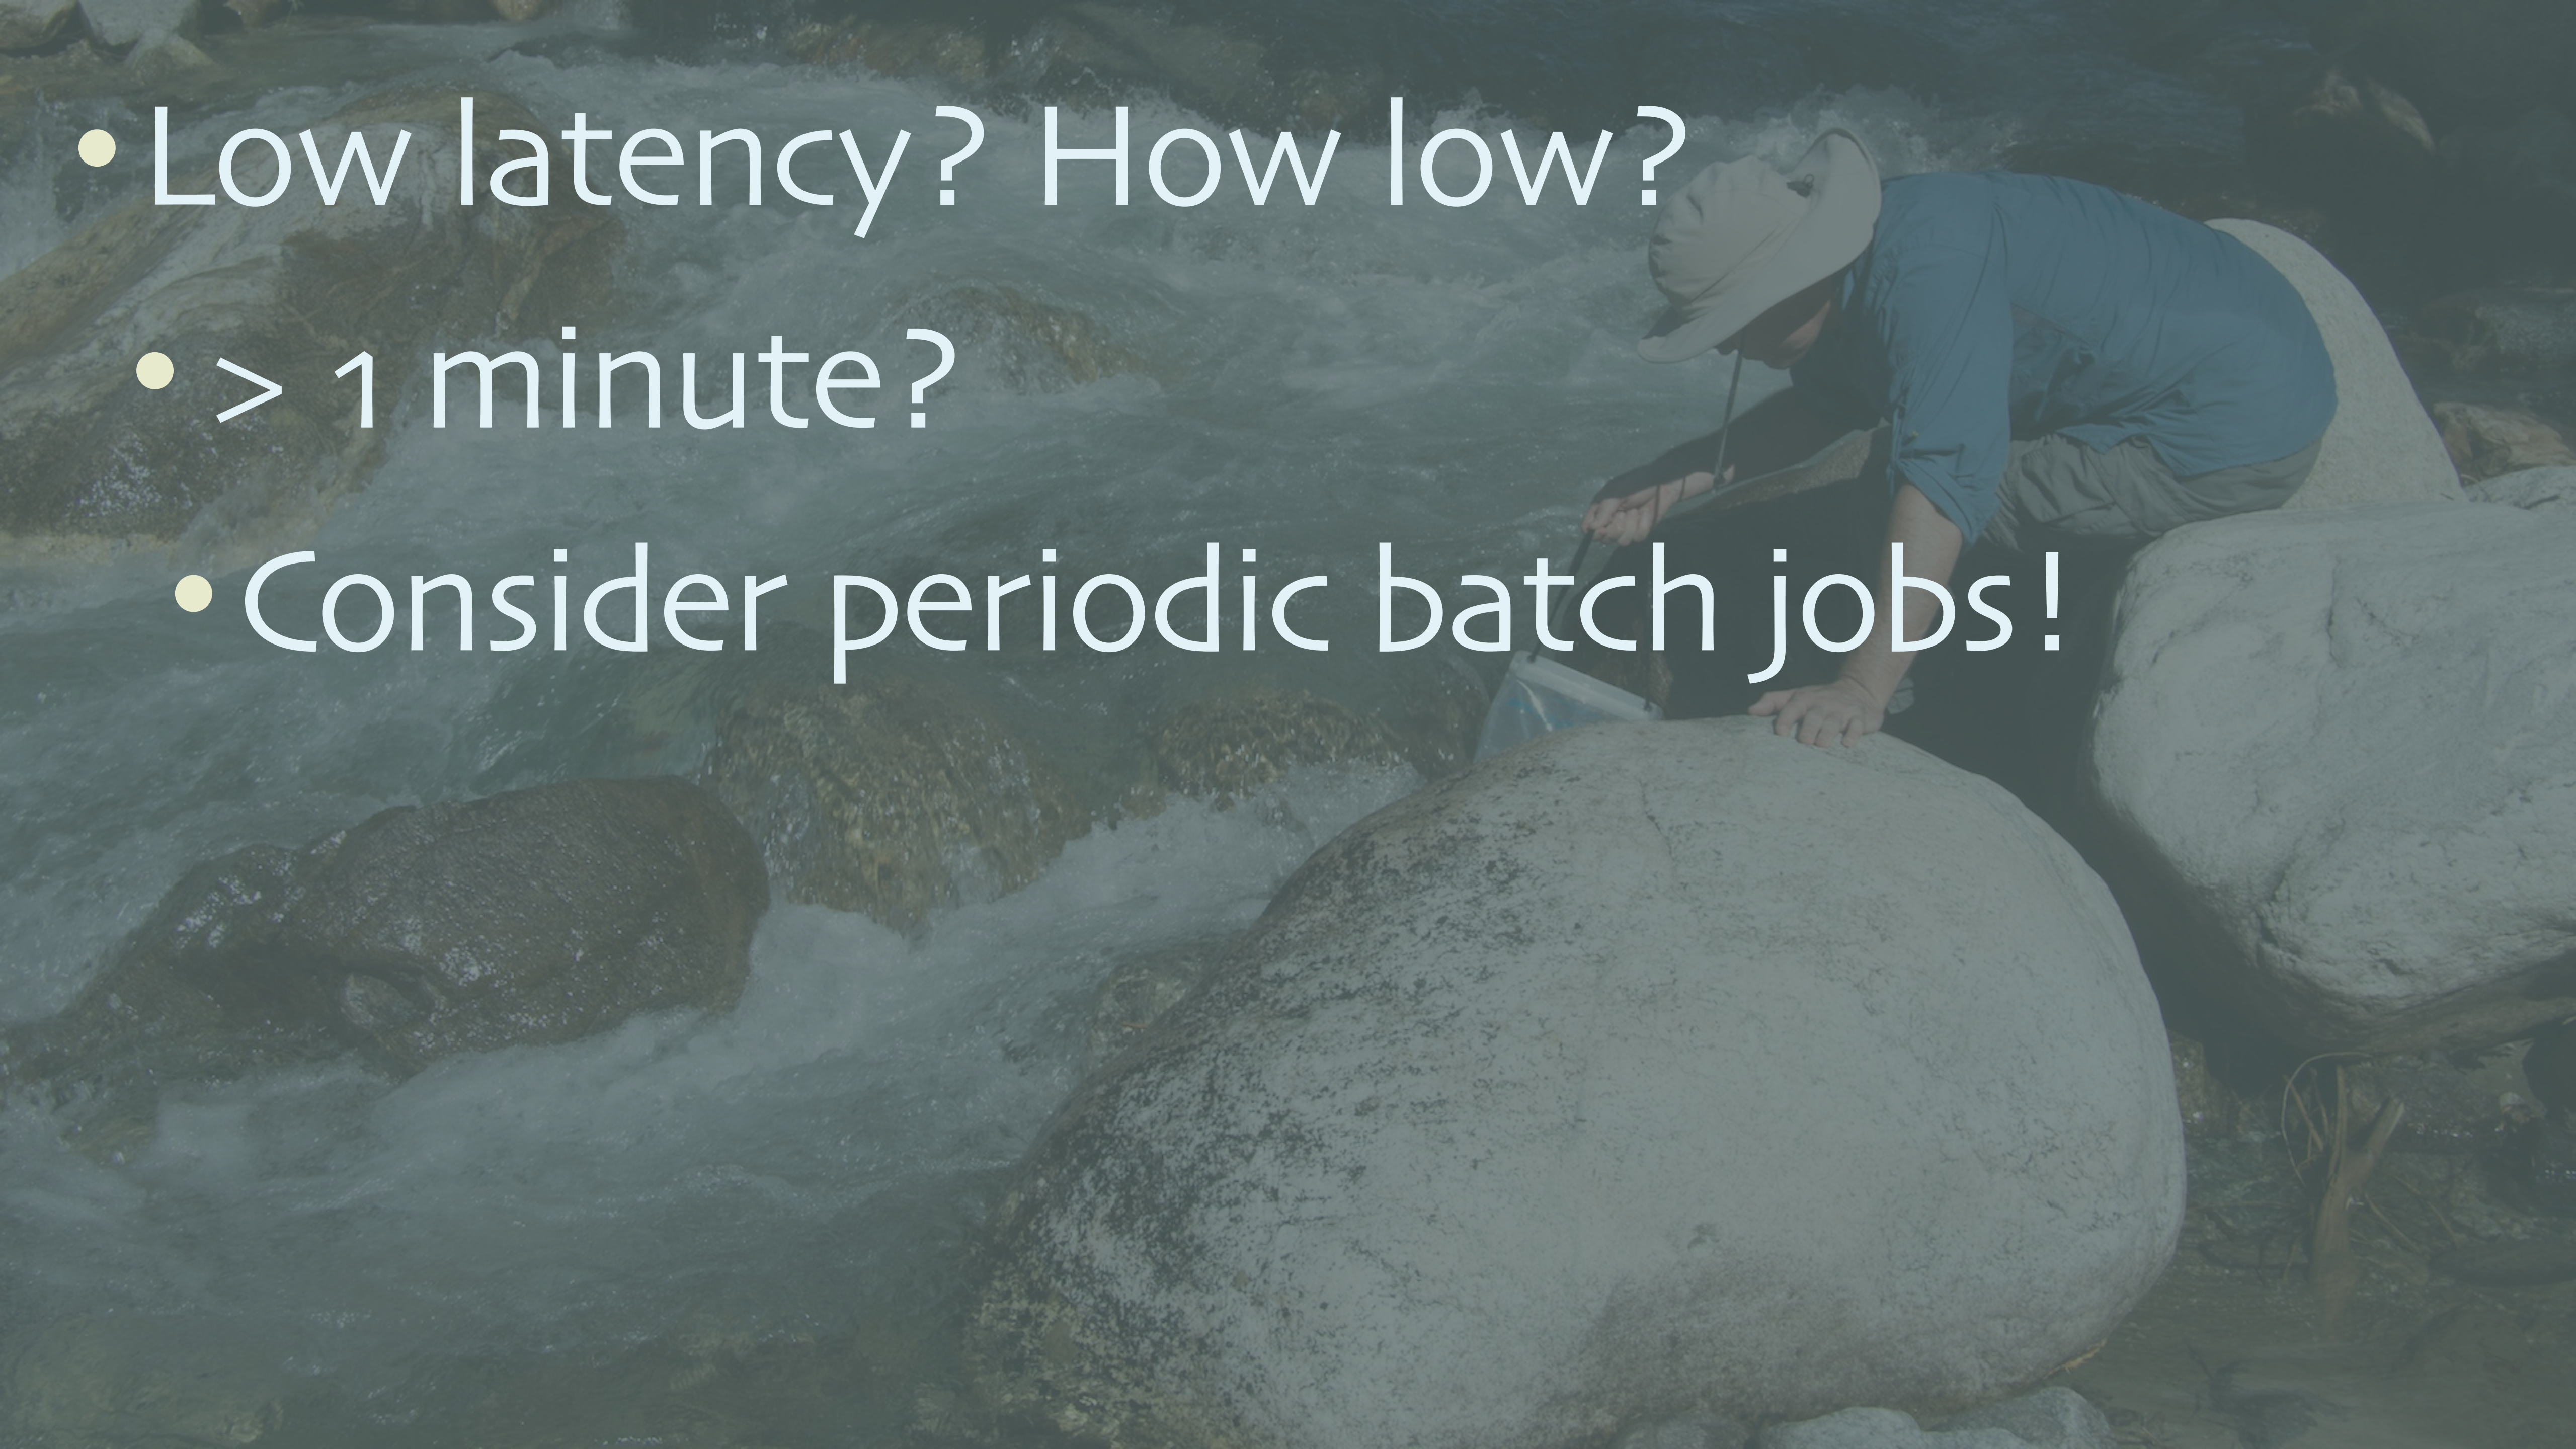
<https://github.com/keen/dashboards>

- 
- A person wearing a blue long-sleeved shirt, a light-colored hat, and dark pants is crouching on a large, light-colored rock in a river. They are holding a small white bucket with a black handle. The background shows the flowing water of the river and other rocks.
- Low latency? How low?
 - < hundreds of milliseconds?
 - “micro-batches”
 - Processing records in bulk, e.g., Spark’s micro-batch model and “streaming SQL” over windows.

- Low latency? How low?
- < 1 second to minutes?



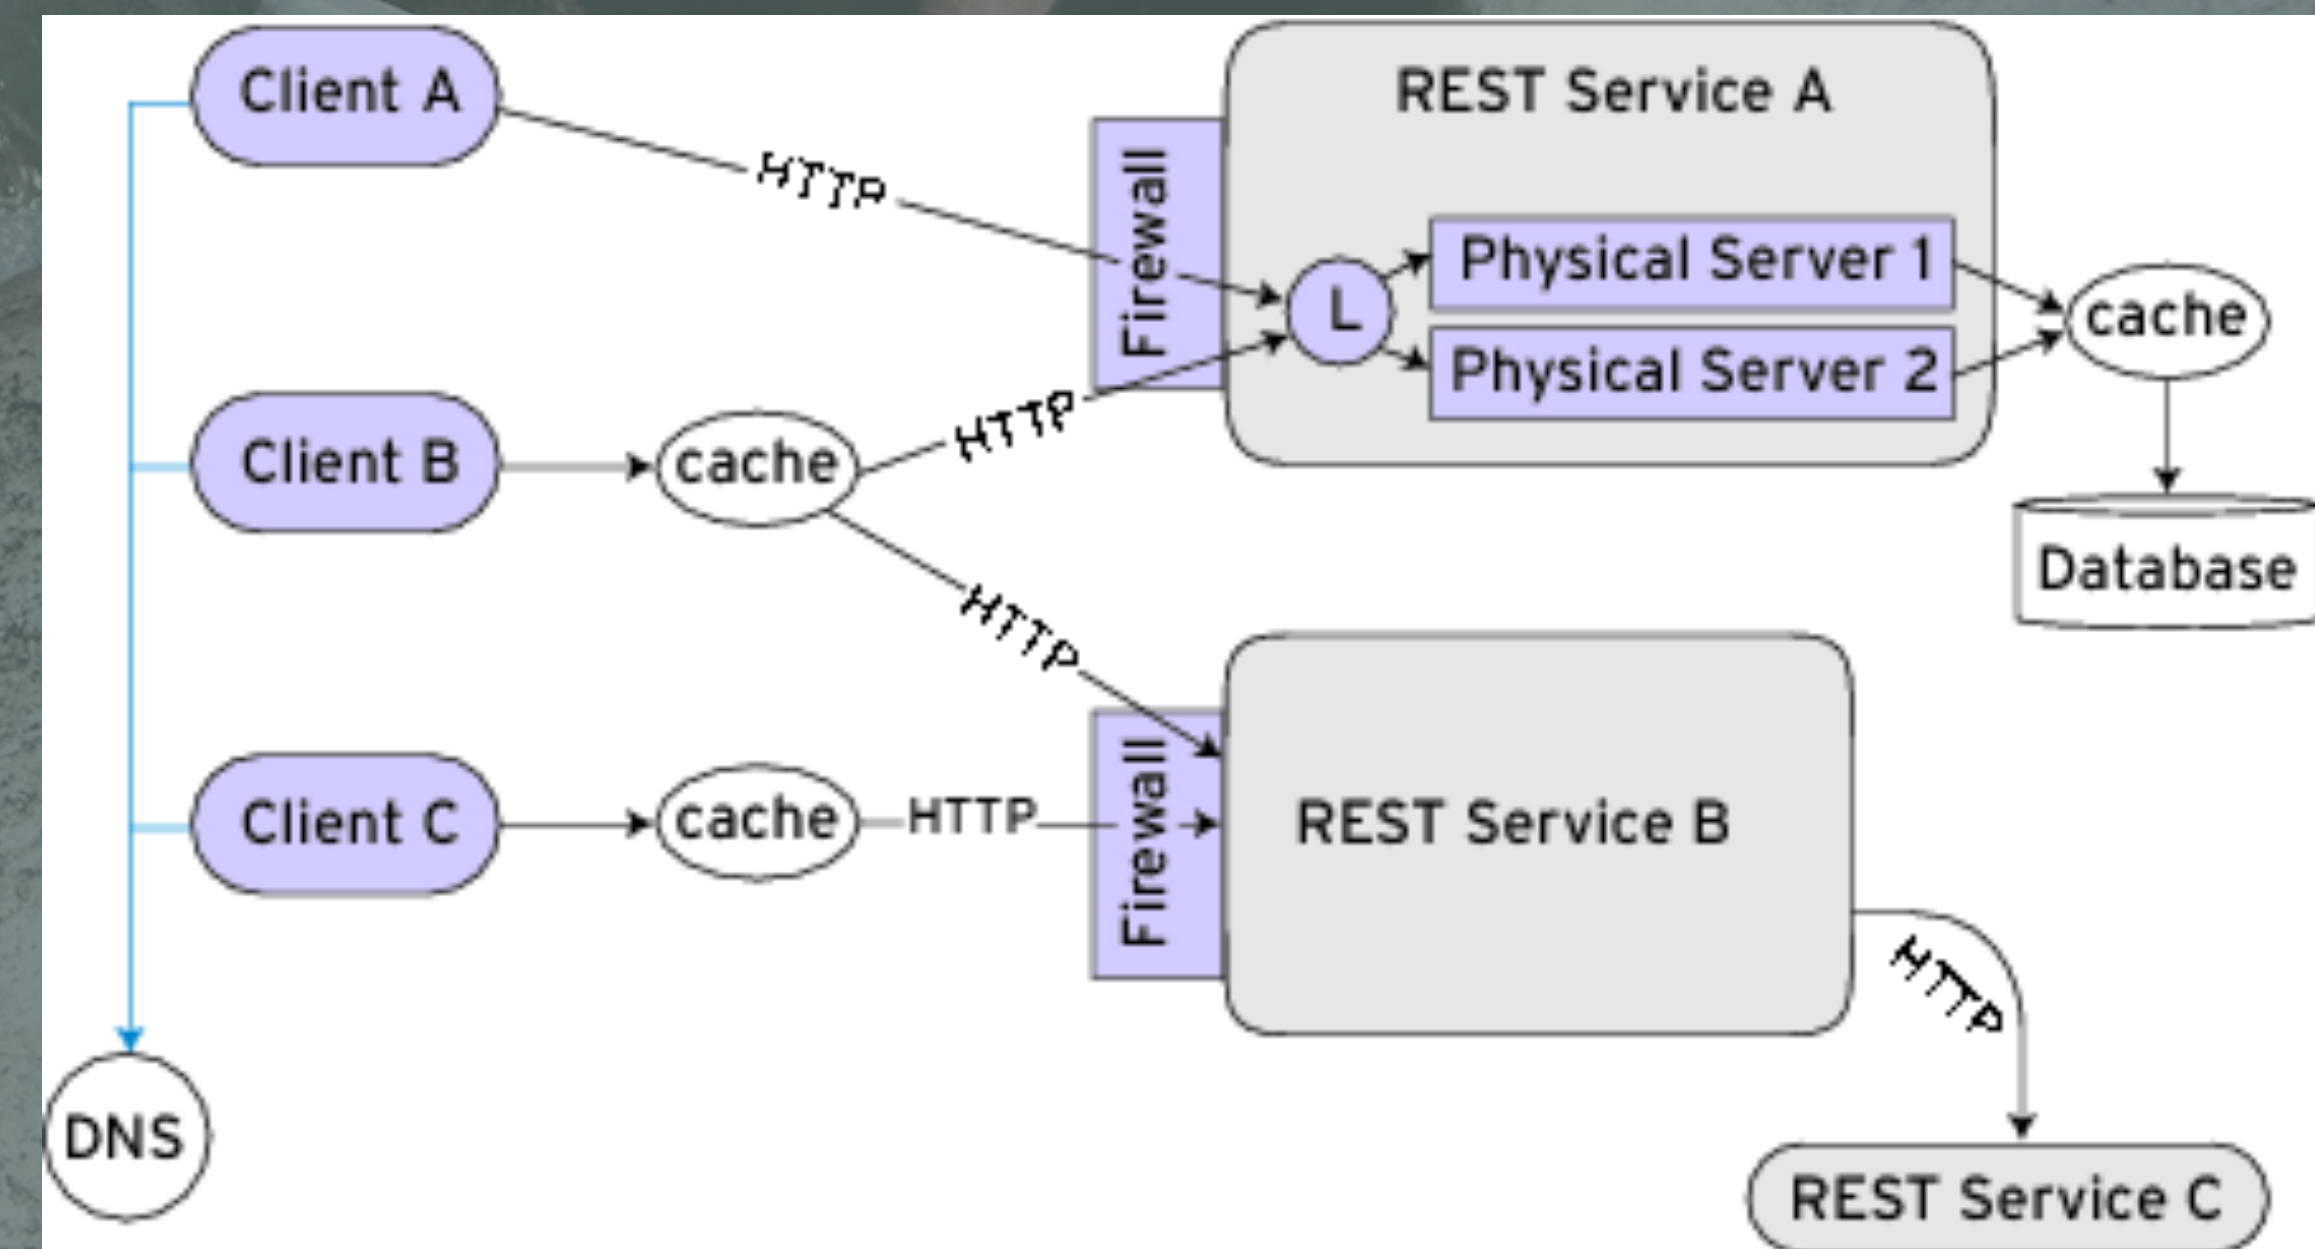
- Low latency? How low?
- > 1 minute?
- Consider periodic batch jobs!



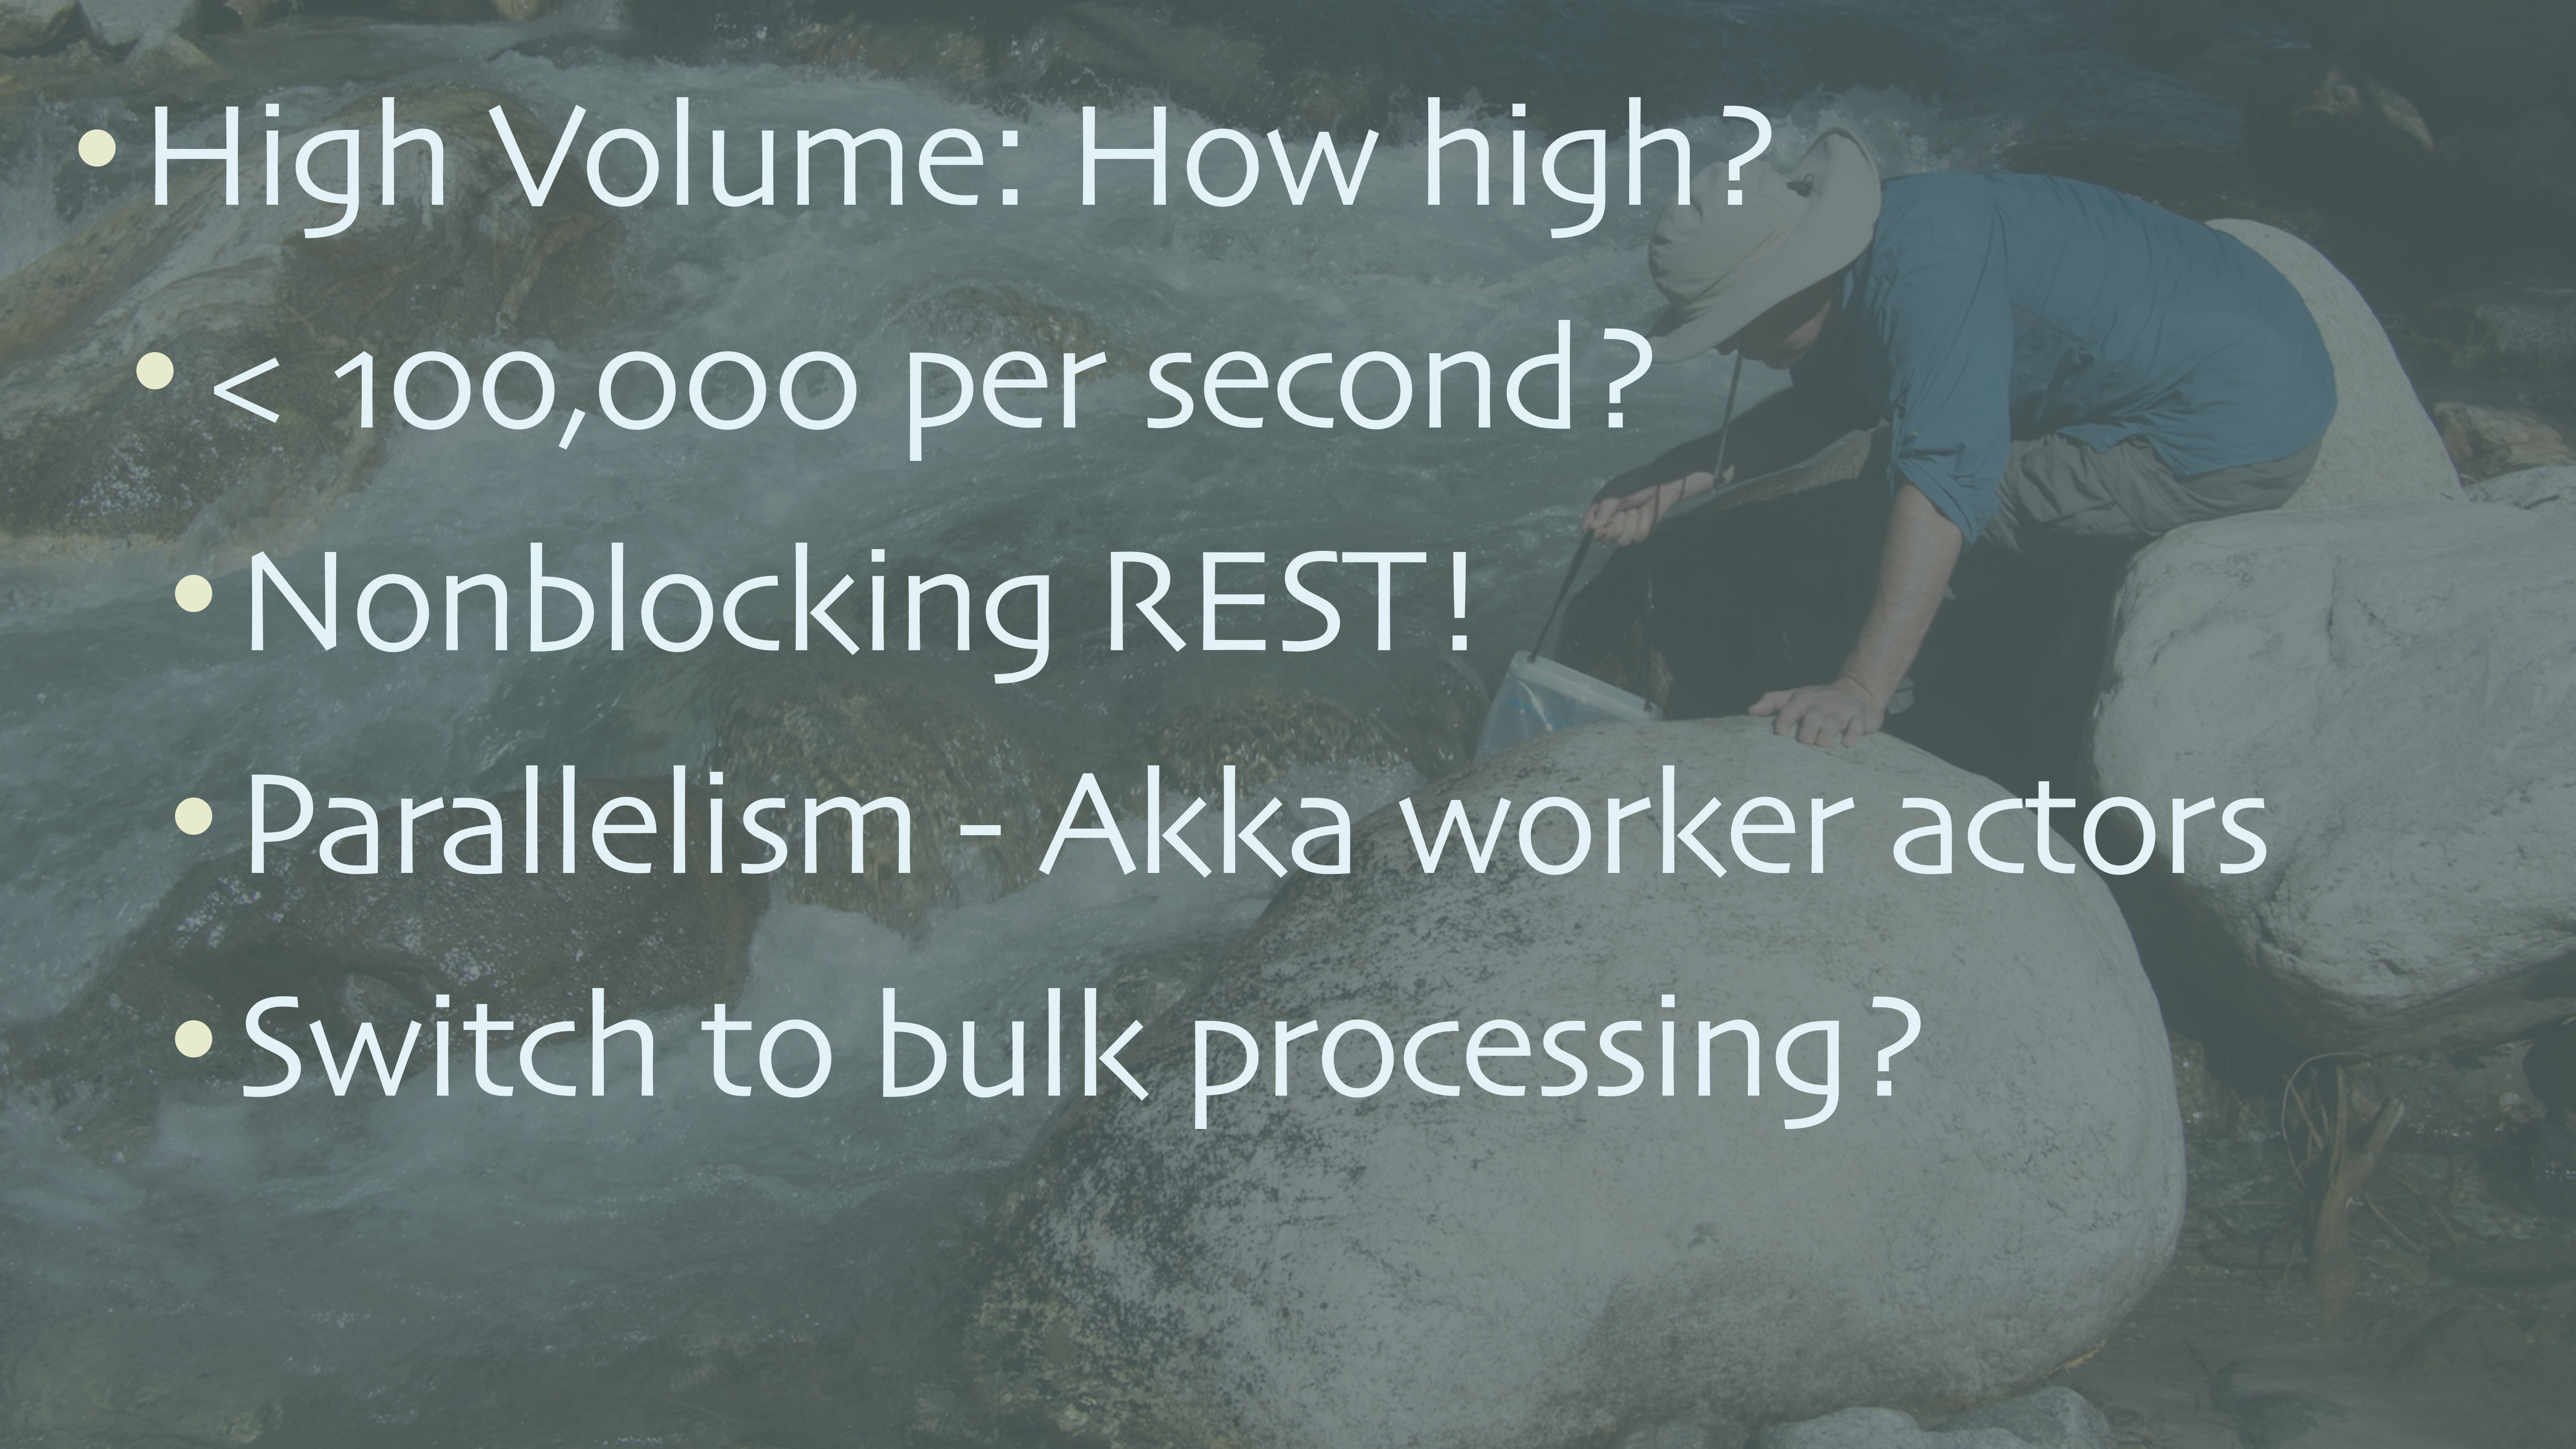
- High Volume: How high?



- High Volume: How high?
- < 10,000 events/second?
- REST
- One at a time...



<http://www.drdobbs.com/web-development/soa-web-services-and-restful-systems/199902676>

- 
- A person wearing a blue long-sleeved shirt, a light-colored hat, and dark pants is crouching on a large, smooth, light-colored rock in a stream. They are holding a white bucket with a black handle. The water is flowing over rocks, creating white foam. The background is a blurred view of the stream and more rocks.
- High Volume: How high?
 - < 100,000 per second?
 - Nonblocking REST!
 - Parallelism - Akka worker actors
 - Switch to bulk processing?

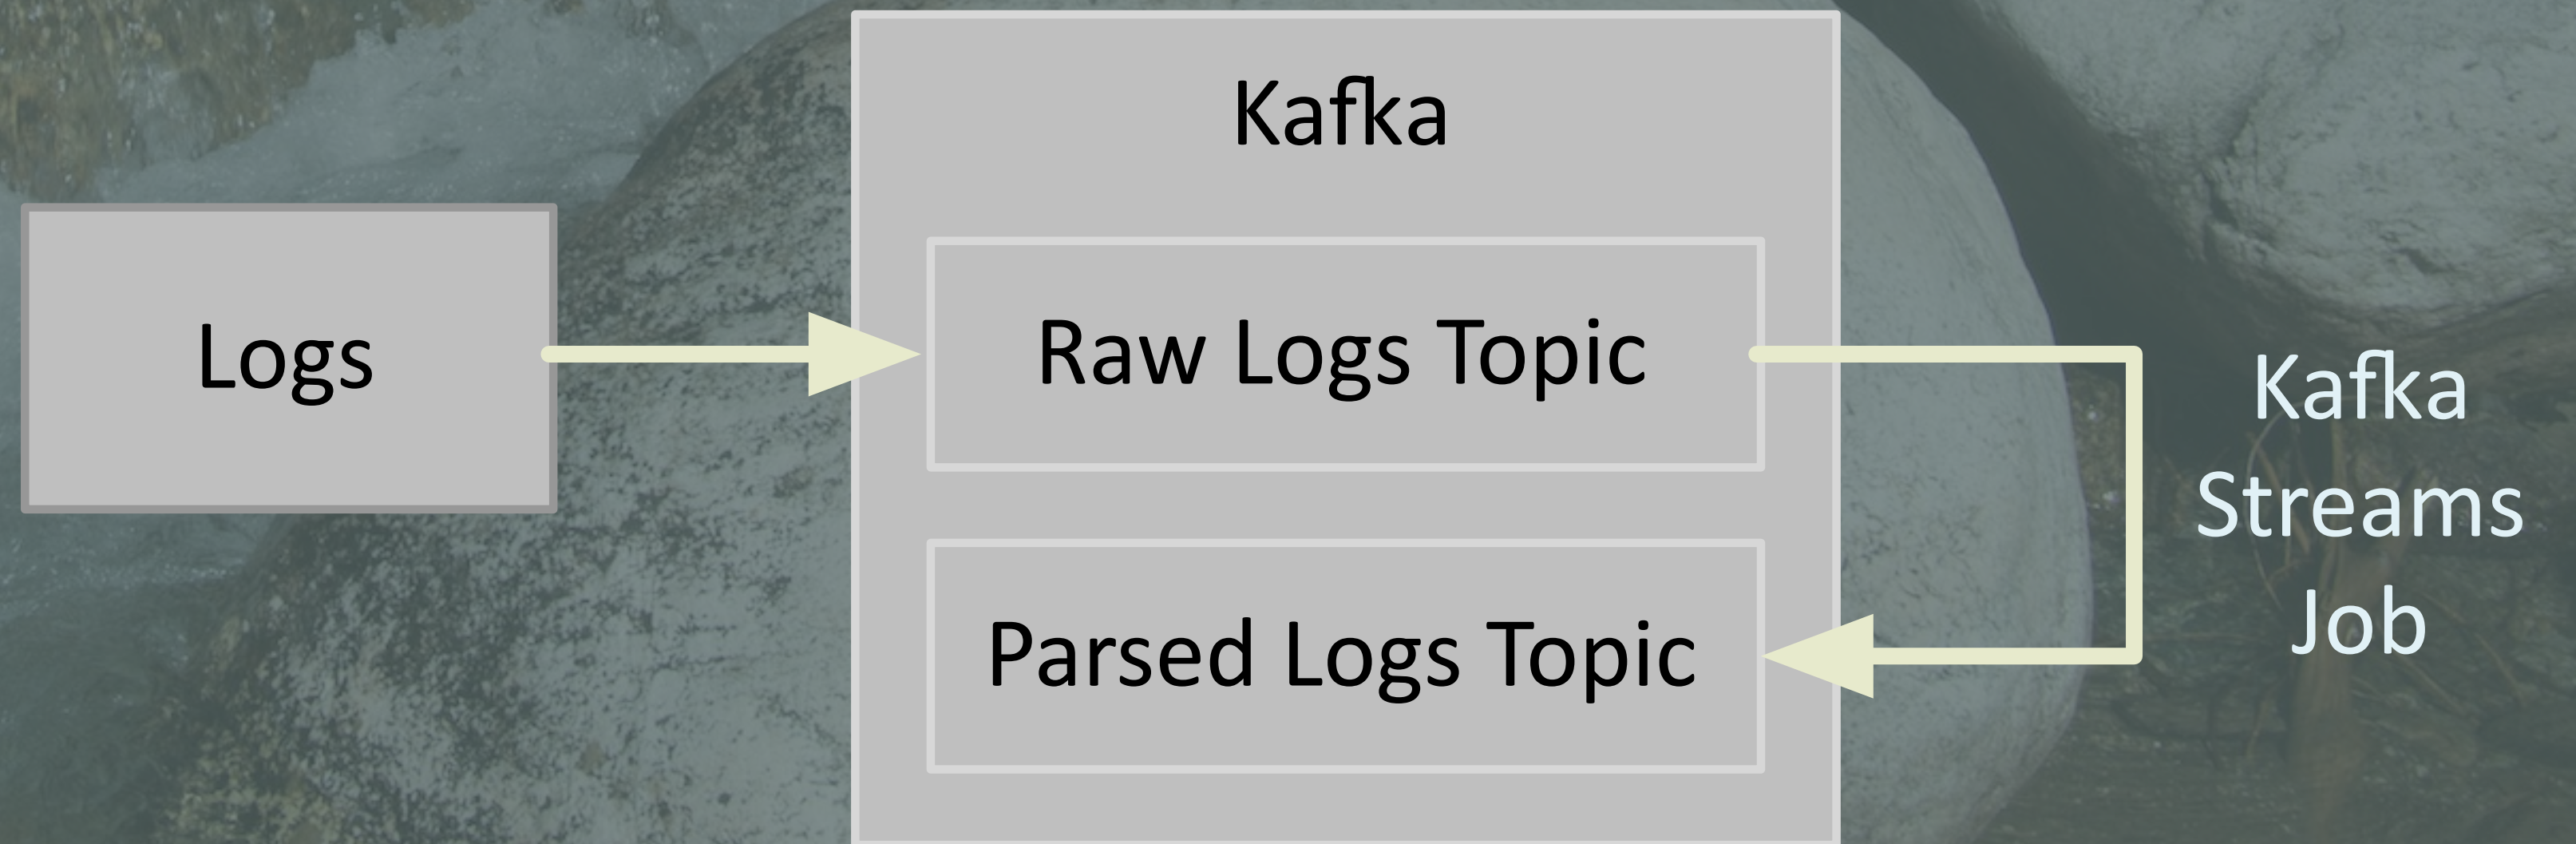
- High Volume: How high?
- 1,000,000s per second?
- Flink or Spark Streaming
- Process in bulk



- Which kinds of data processing?



- Which kinds of data processing?
- Extract, transform, and load (ETL)?



- Which kinds of data processing?
- “Dataflow” pipelines

```
val sc = new SparkContext("local[*]", "Inverted Idx")
sc.textFile("data/crawl")
  .map { line => val Array(path, text) = line.split("\t", 2);
    (path, text)
  } flatMap {
    case (path, text) => text.split("""\W+""").map((_, path))
  } map {
    case (w, p) => ((w, p), 1)
  } reduceByKey {
    case (n1, n2) => n1 + n2
  }
```



- Which kinds of data processing?
- SQL?

```
SELECT COUNT(*)  
FROM my-iot-data  
GROUP BY zip-code
```

```
val input = spark.read.  
  format("parquet").  
  stream("my-iot-data")
```

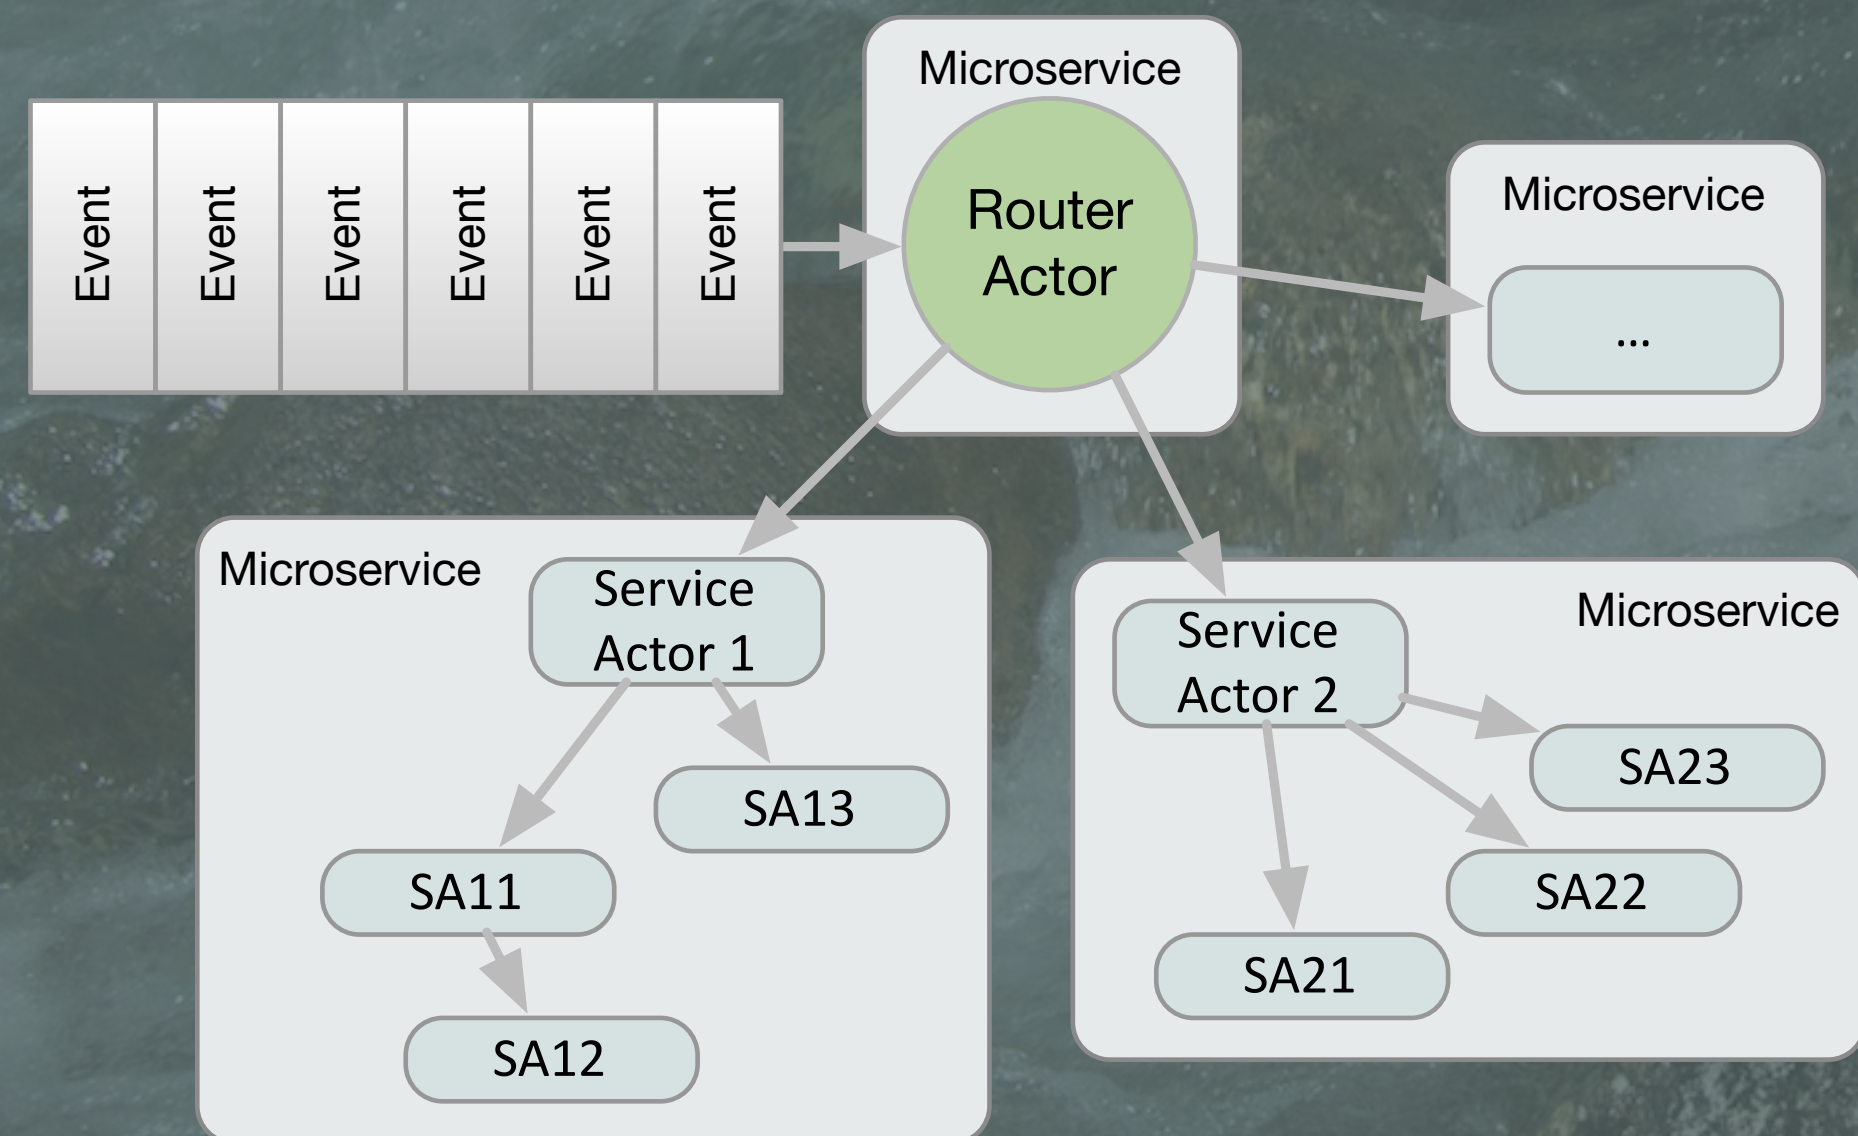
```
input.groupBy("zip-code").  
  count()
```


- Which kinds of data processing?
- Train and serve ML models?




- Process data individually or in bulk?

Event-driven μ -services



“Record-centric” μ -services

```
SELECT COUNT(*)  
FROM my-iot-data  
GROUP BY zip-code
```



Events

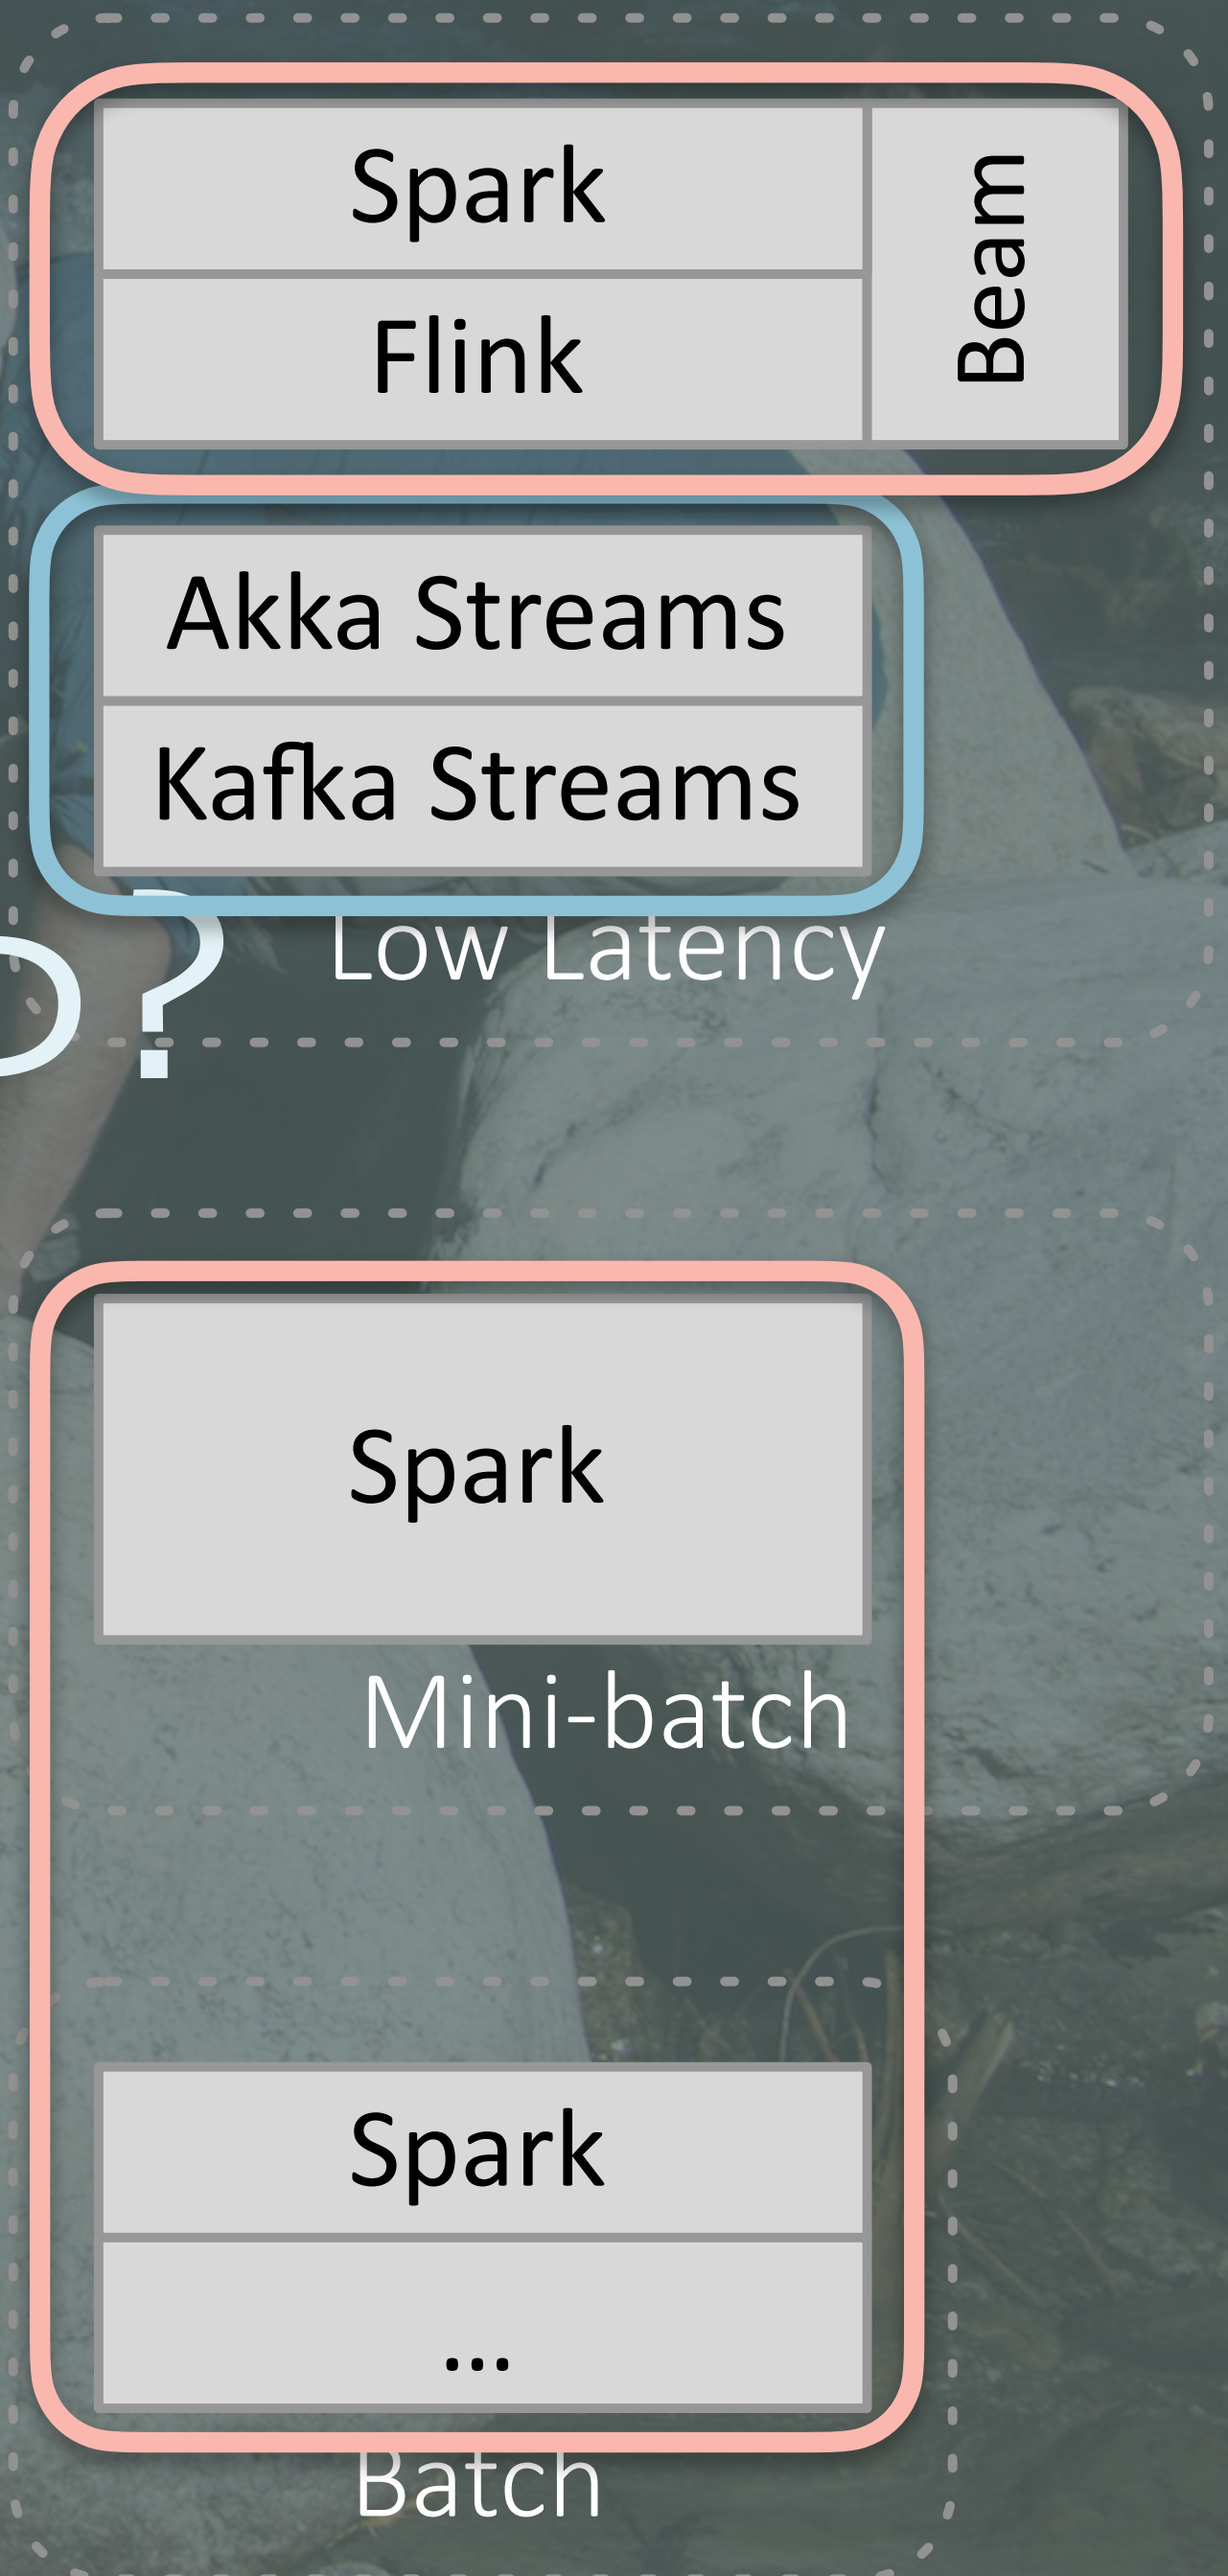
Records

- Preferred application architecture?

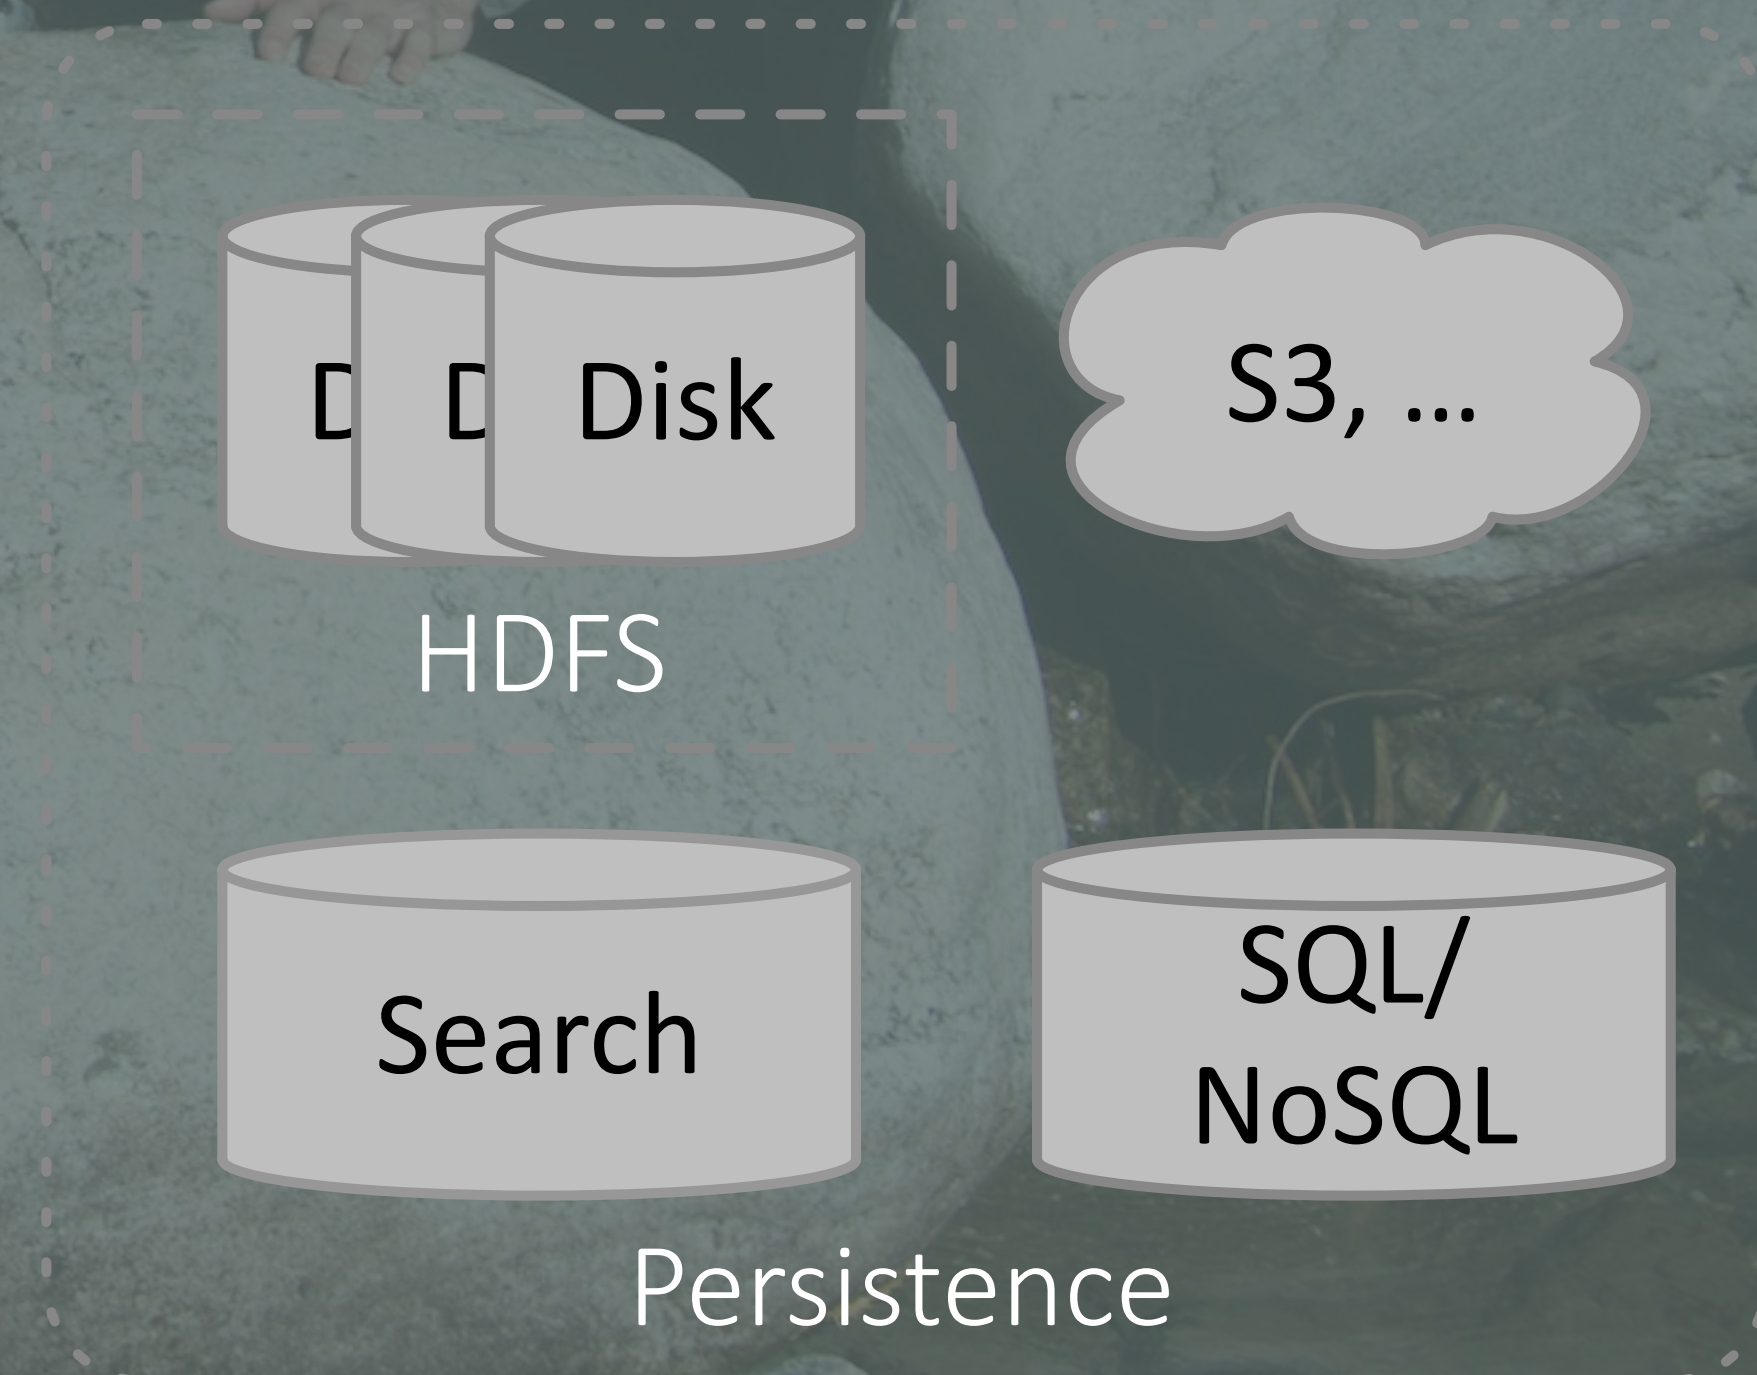
- Streaming library in an app?

- or, distributed services running your job?

Already have a microservices-based, DevOps CI/CD workflow? Stream processing with microservices may fit better into your environment!



- Integration with other tools.
- Akka, Flink, & Spark integrate with Databases, Kafka, file systems, REST, ...
- Kafka Streams only read & write Kafka topics.





Best of Breed Streaming Engines

The background of the slide is a photograph of a waterfall cascading over dark, layered rocks. The water is white and frothy as it falls, creating a sense of motion and power. The overall tone is dark and moody, with the white water providing a strong contrast.

Spark	Beam
Flink	

Akka Streams
Kafka Streams

Low Latency

Spark

Mini-batch

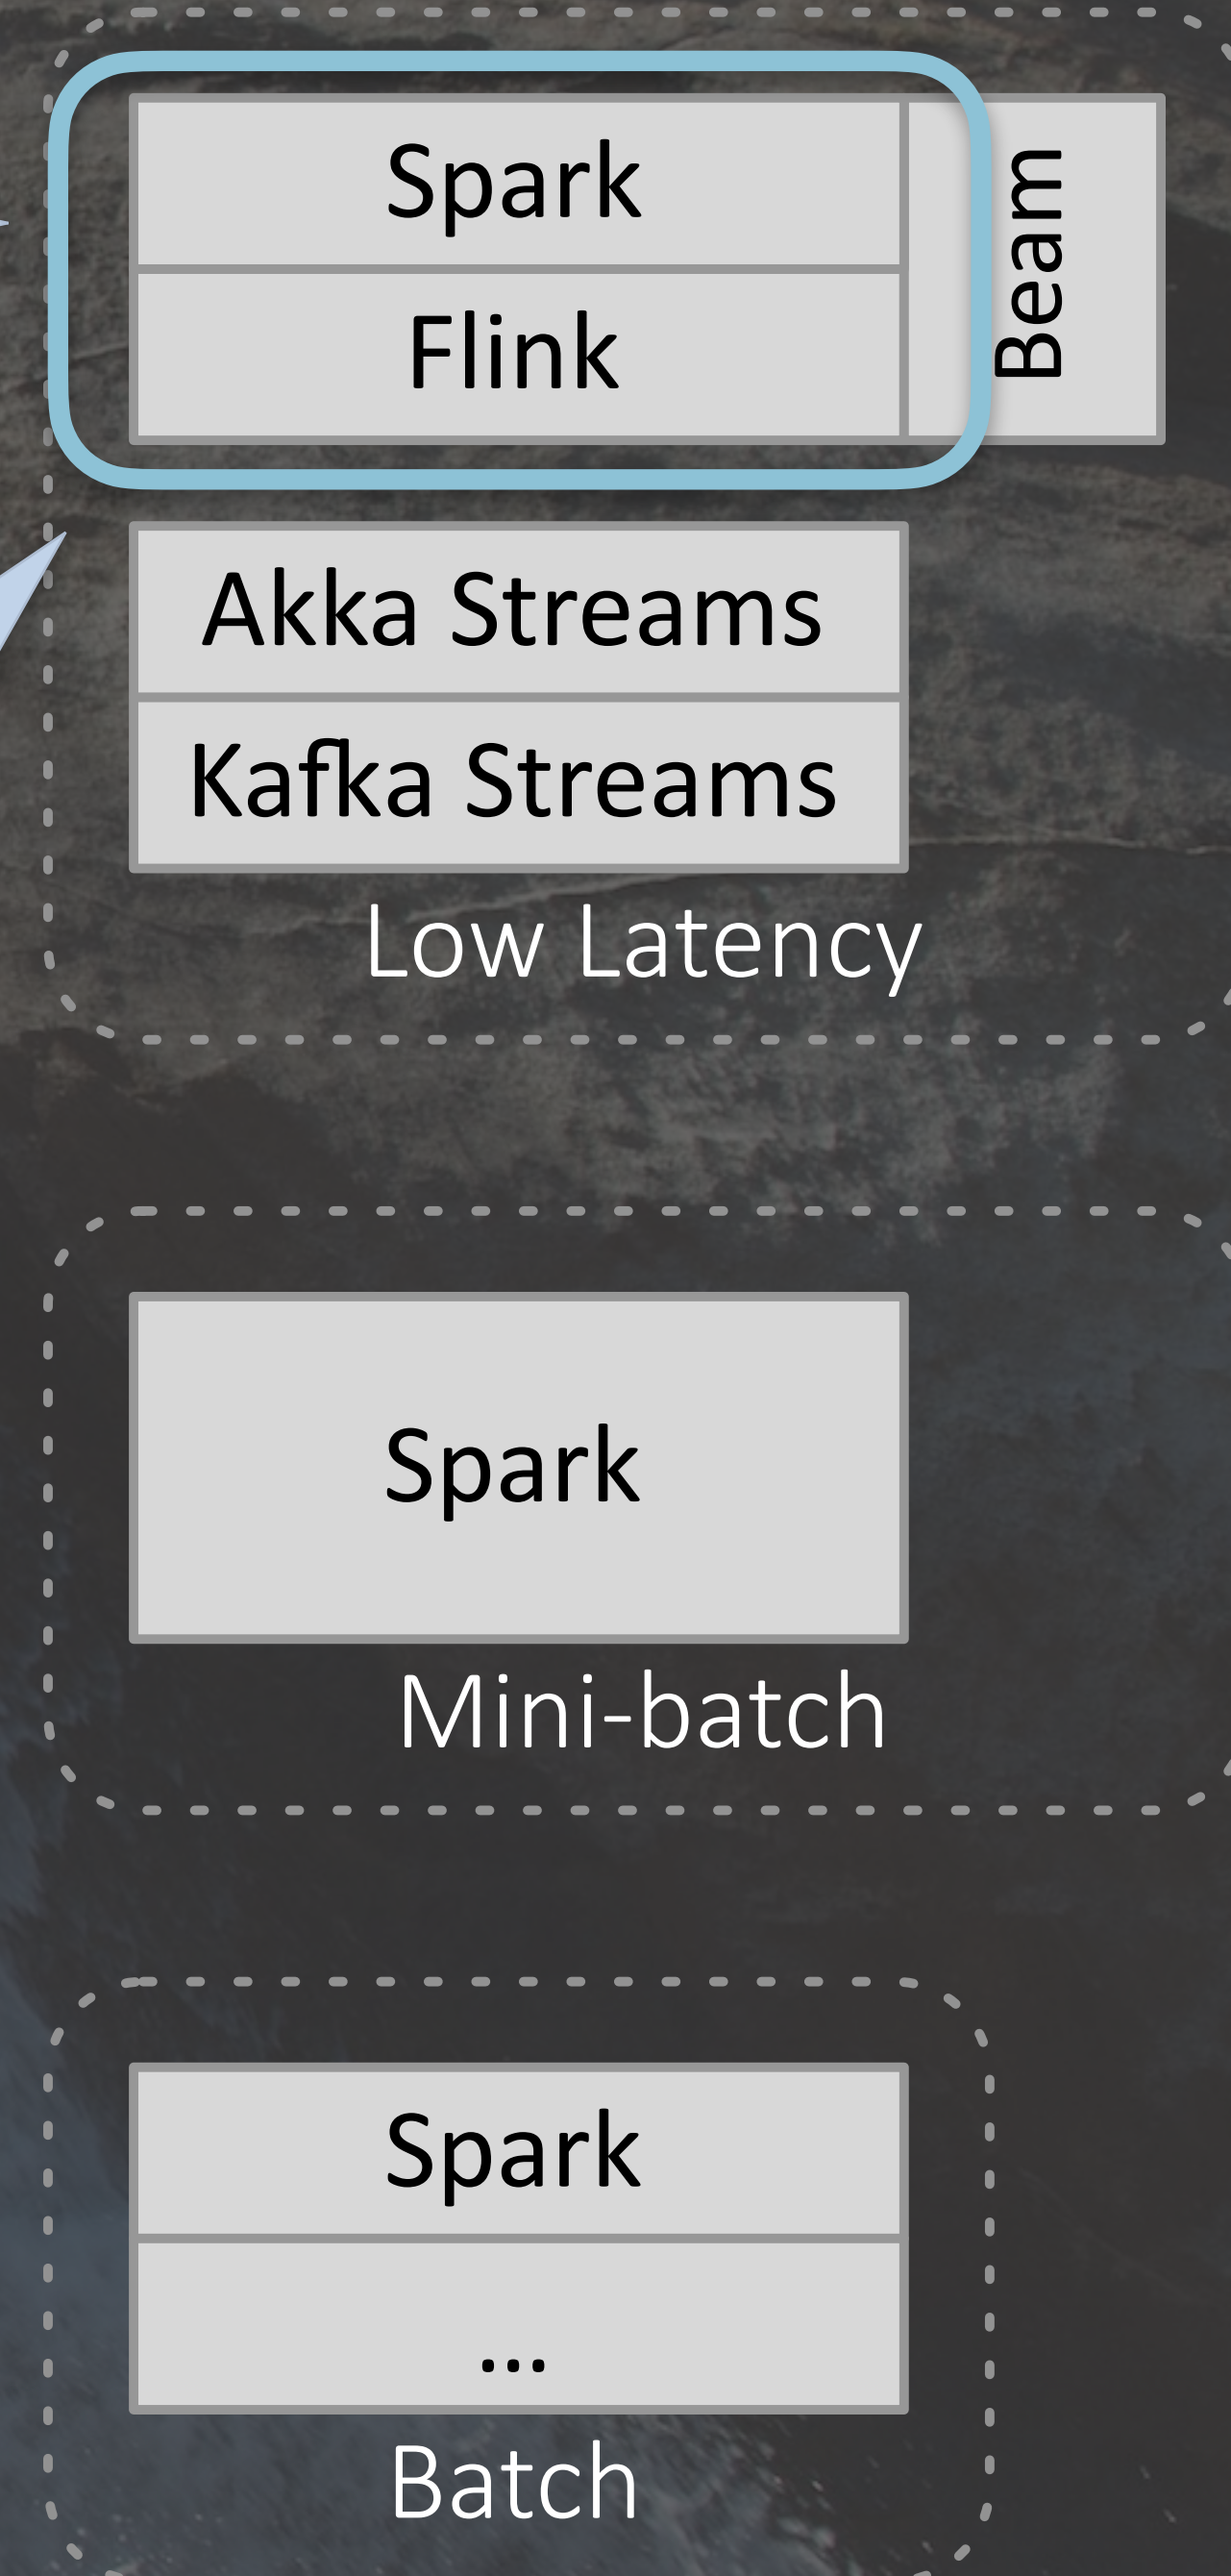
Spark
...

Batch

The streaming engines form two groups:

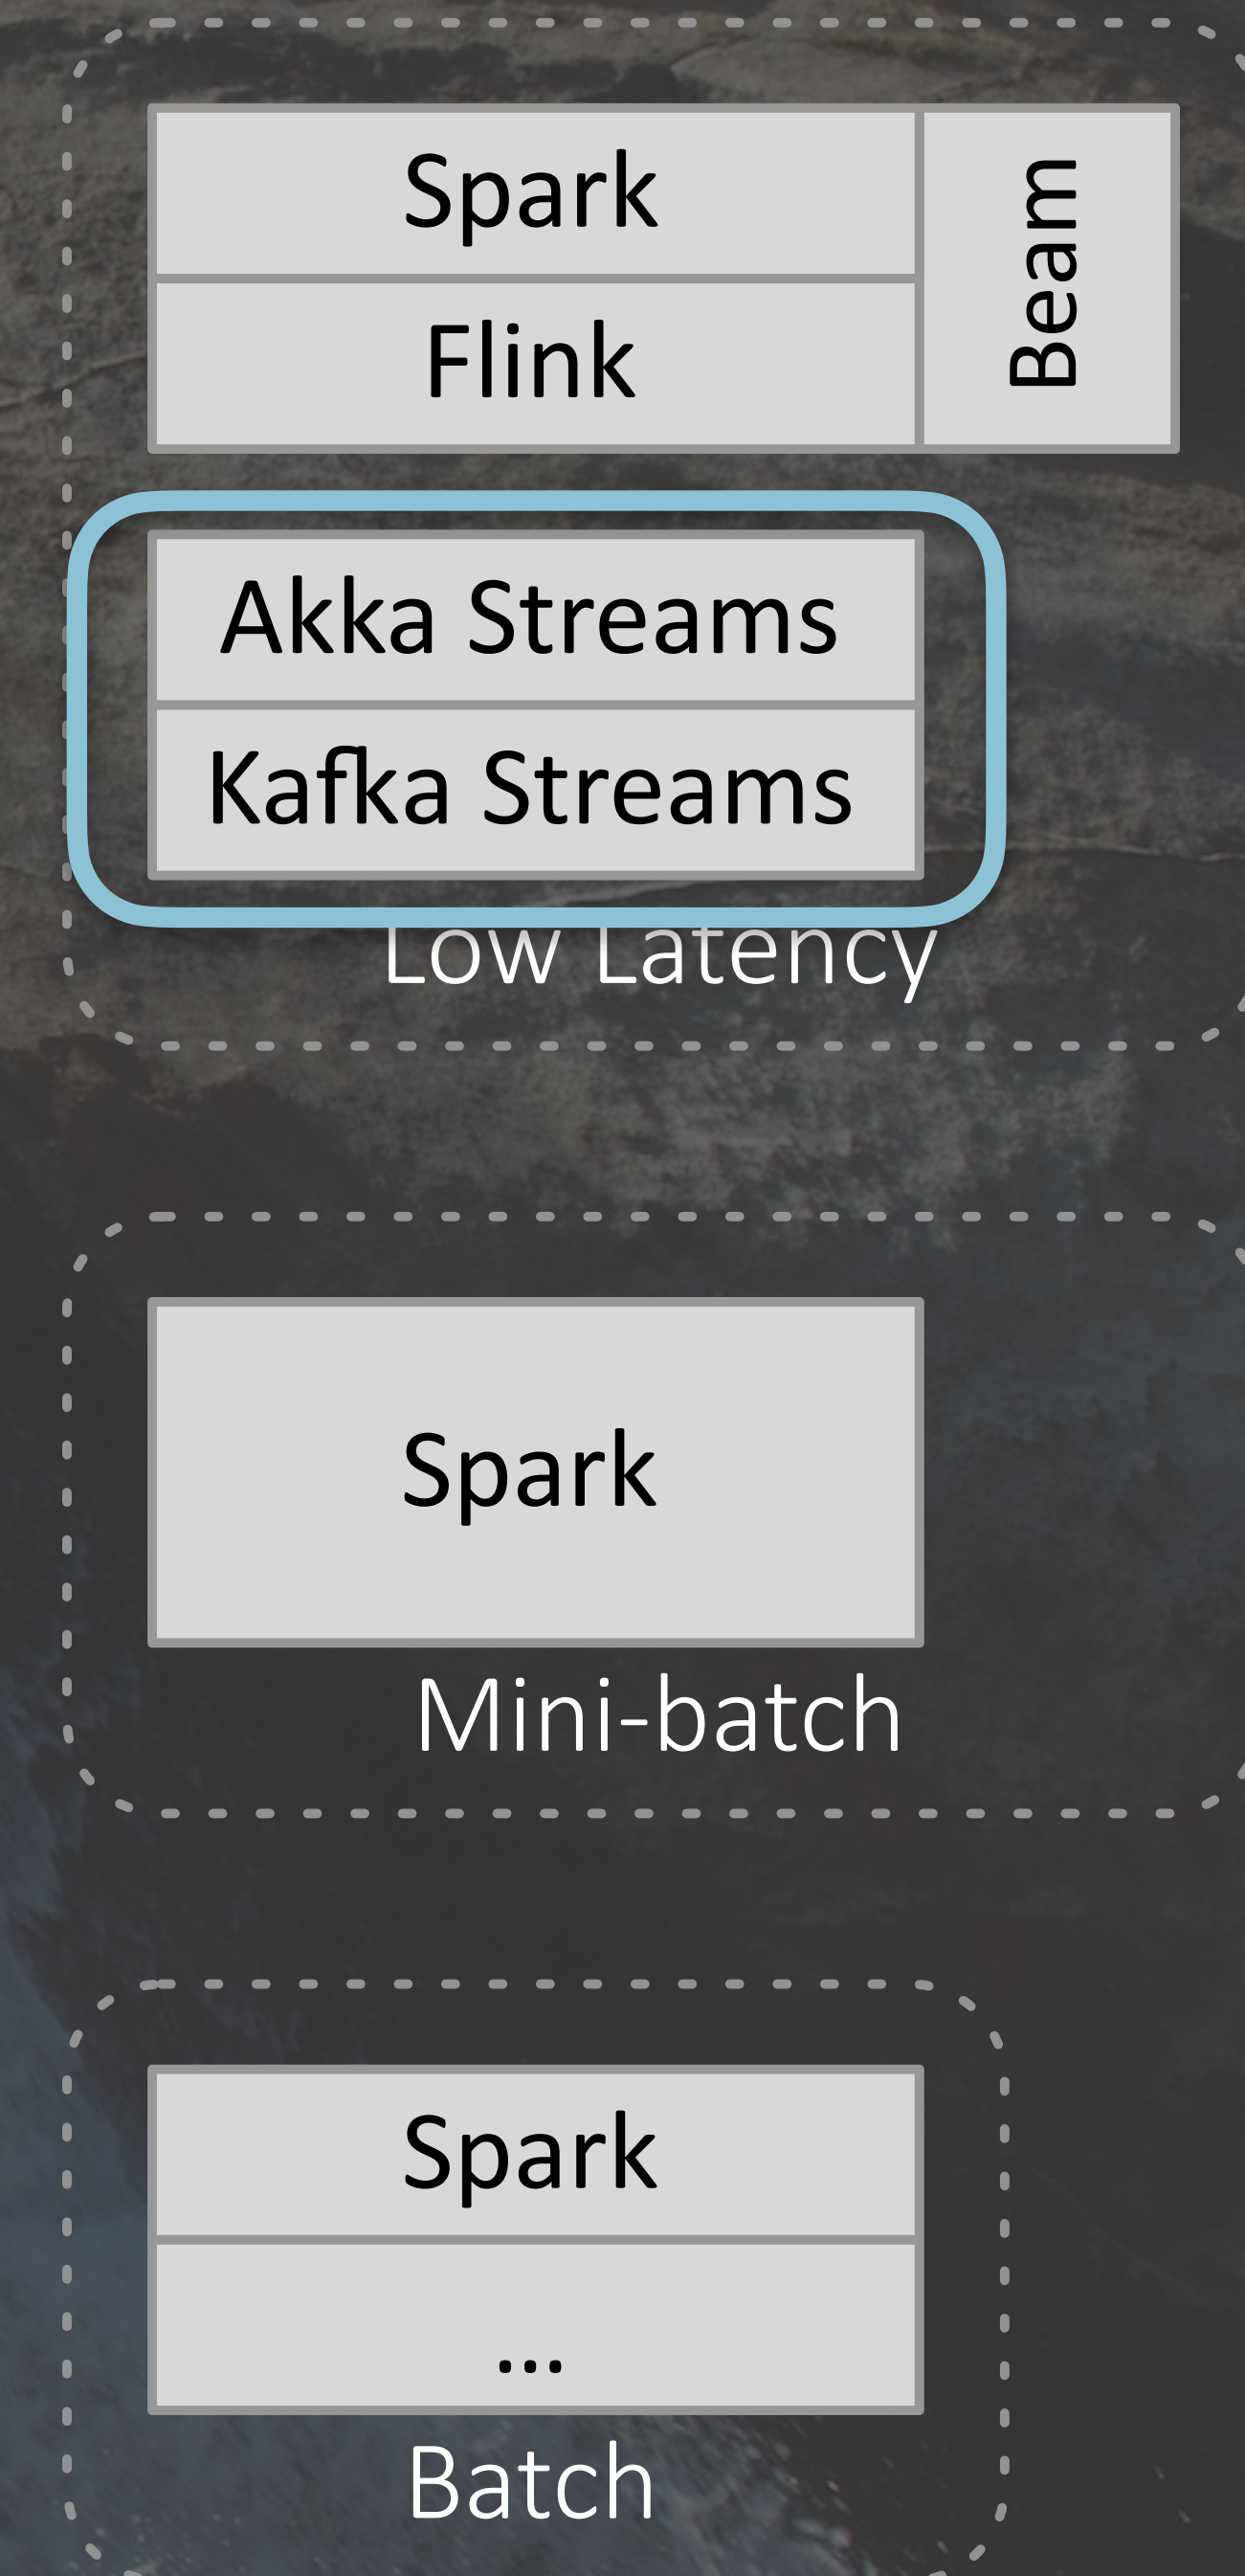
Run as distributed services

You submit jobs, they are partitioned into tasks

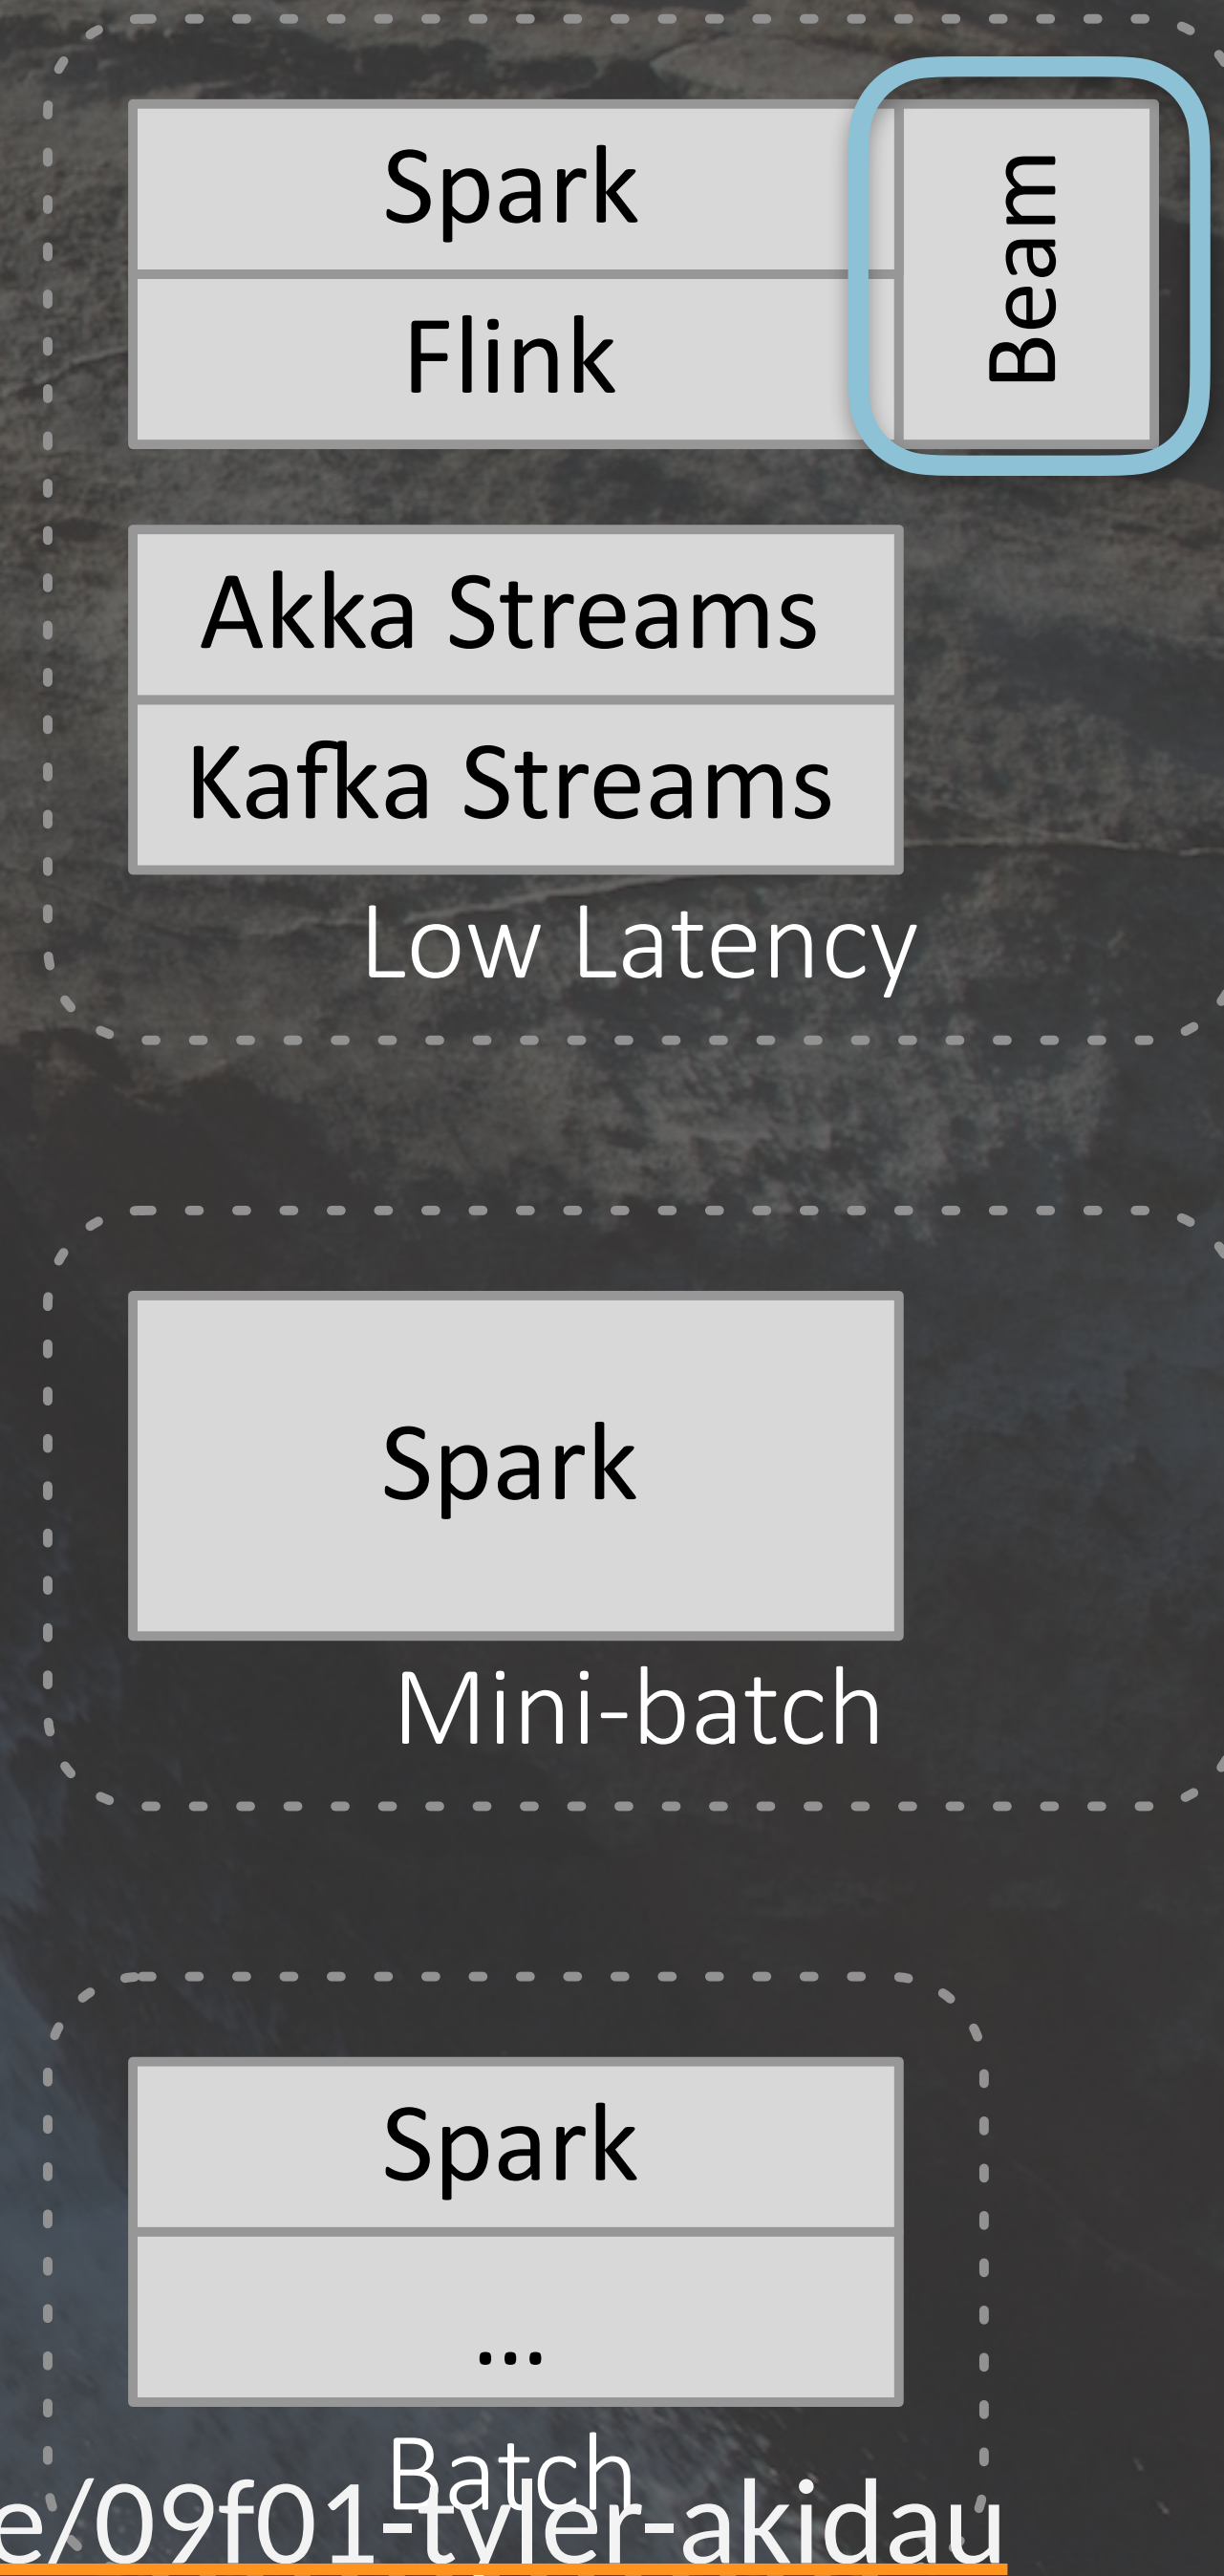


The streaming engines form two groups:

Libraries you embed in your microservices

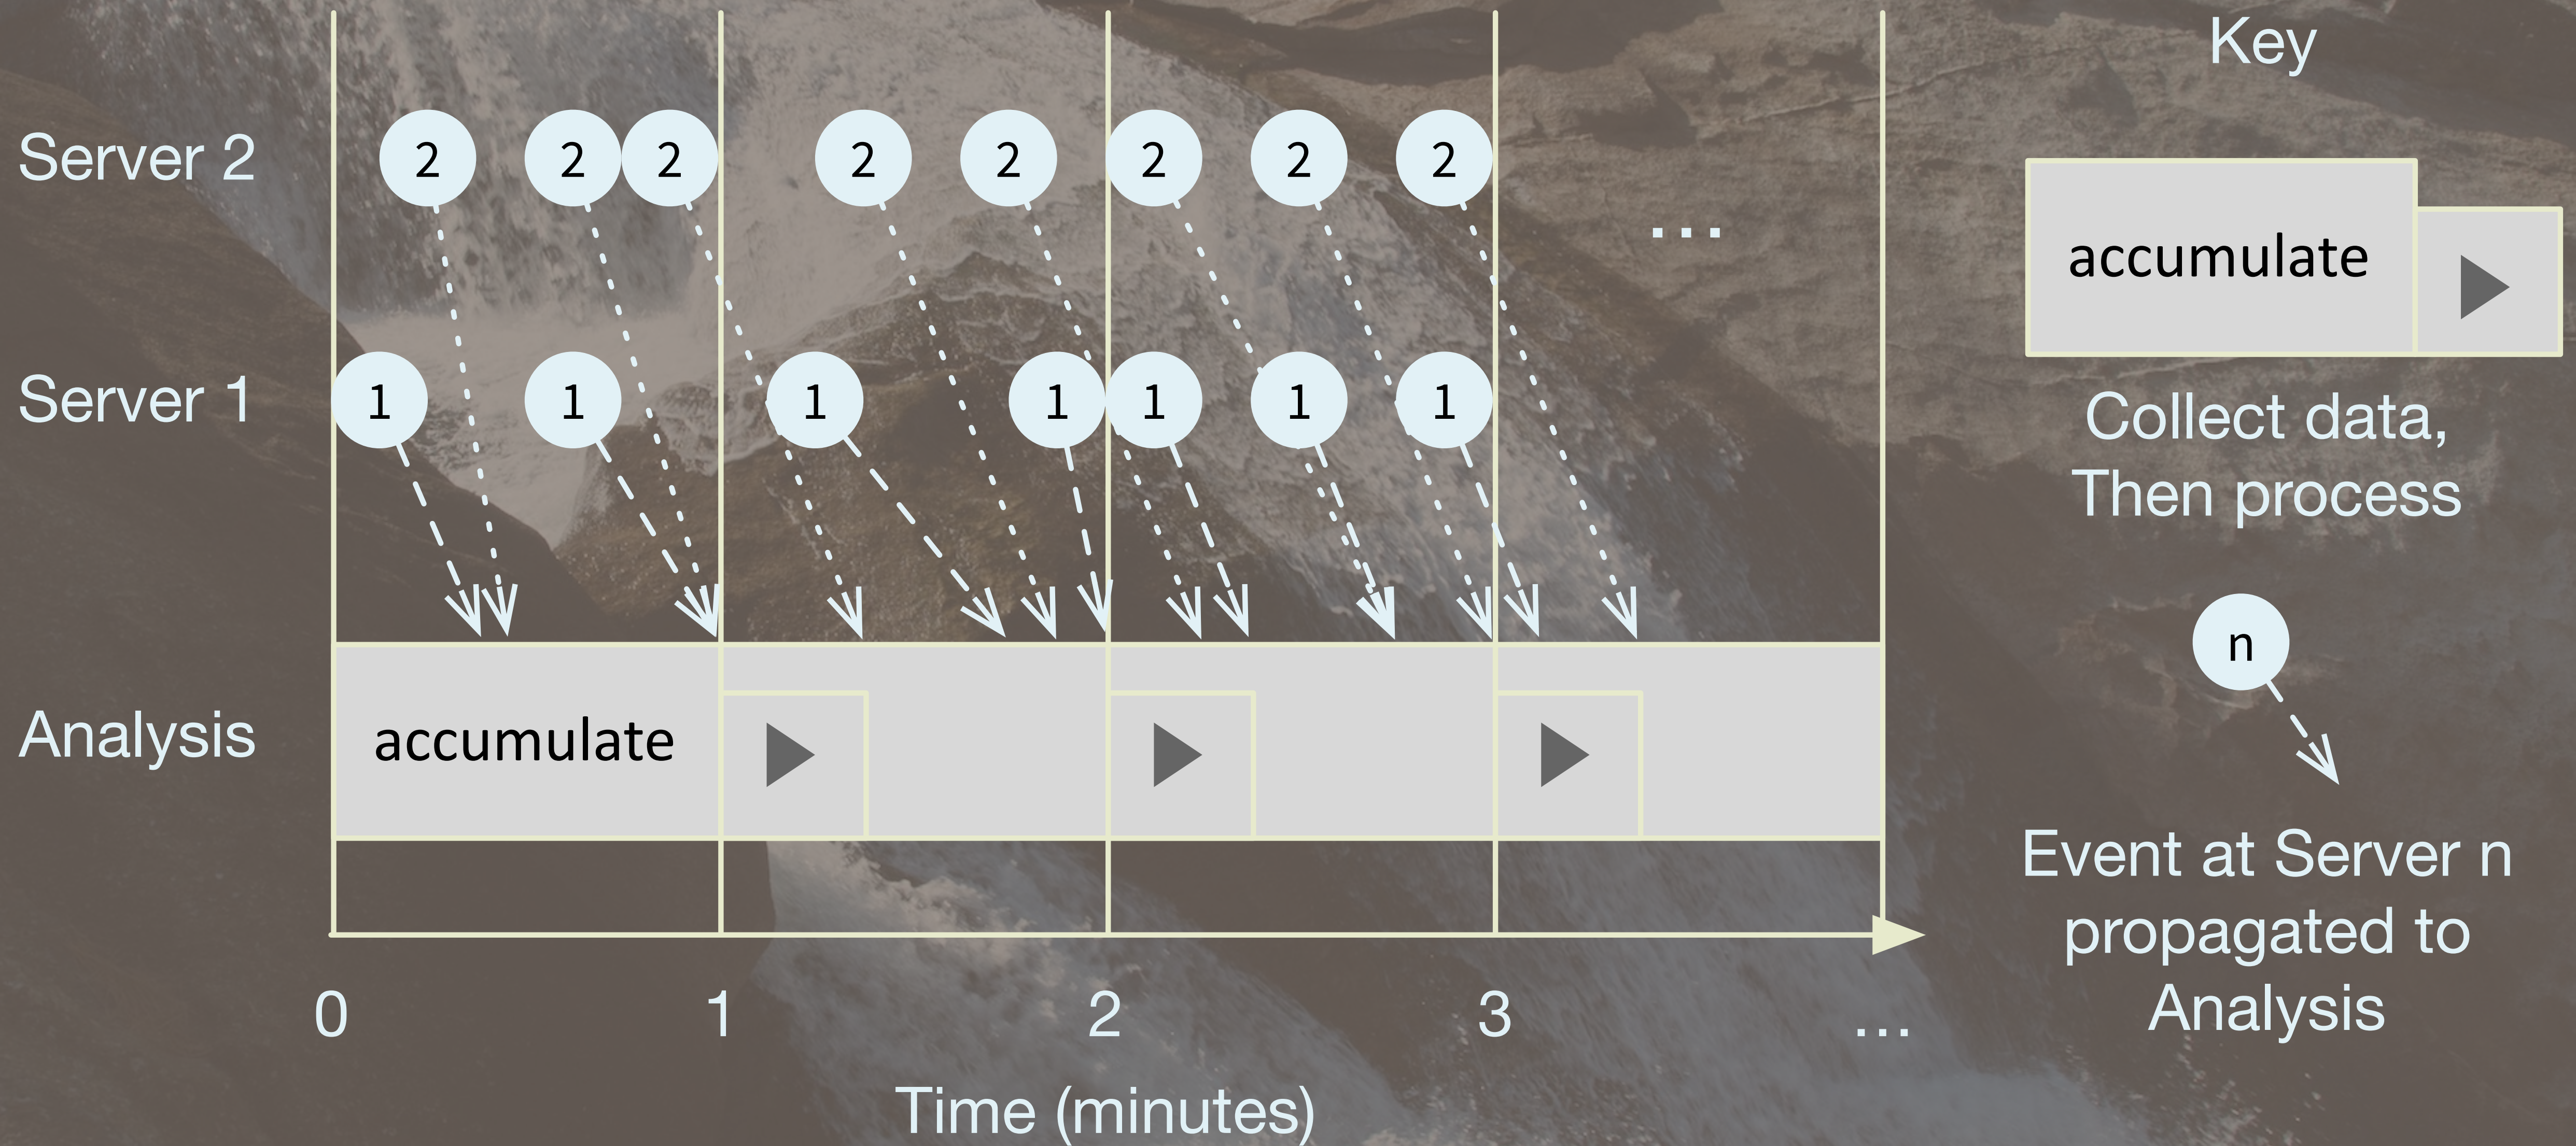


- Apache Beam
- (Google Dataflow)
- Requires a “runner”
- Most sophisticated streaming semantics

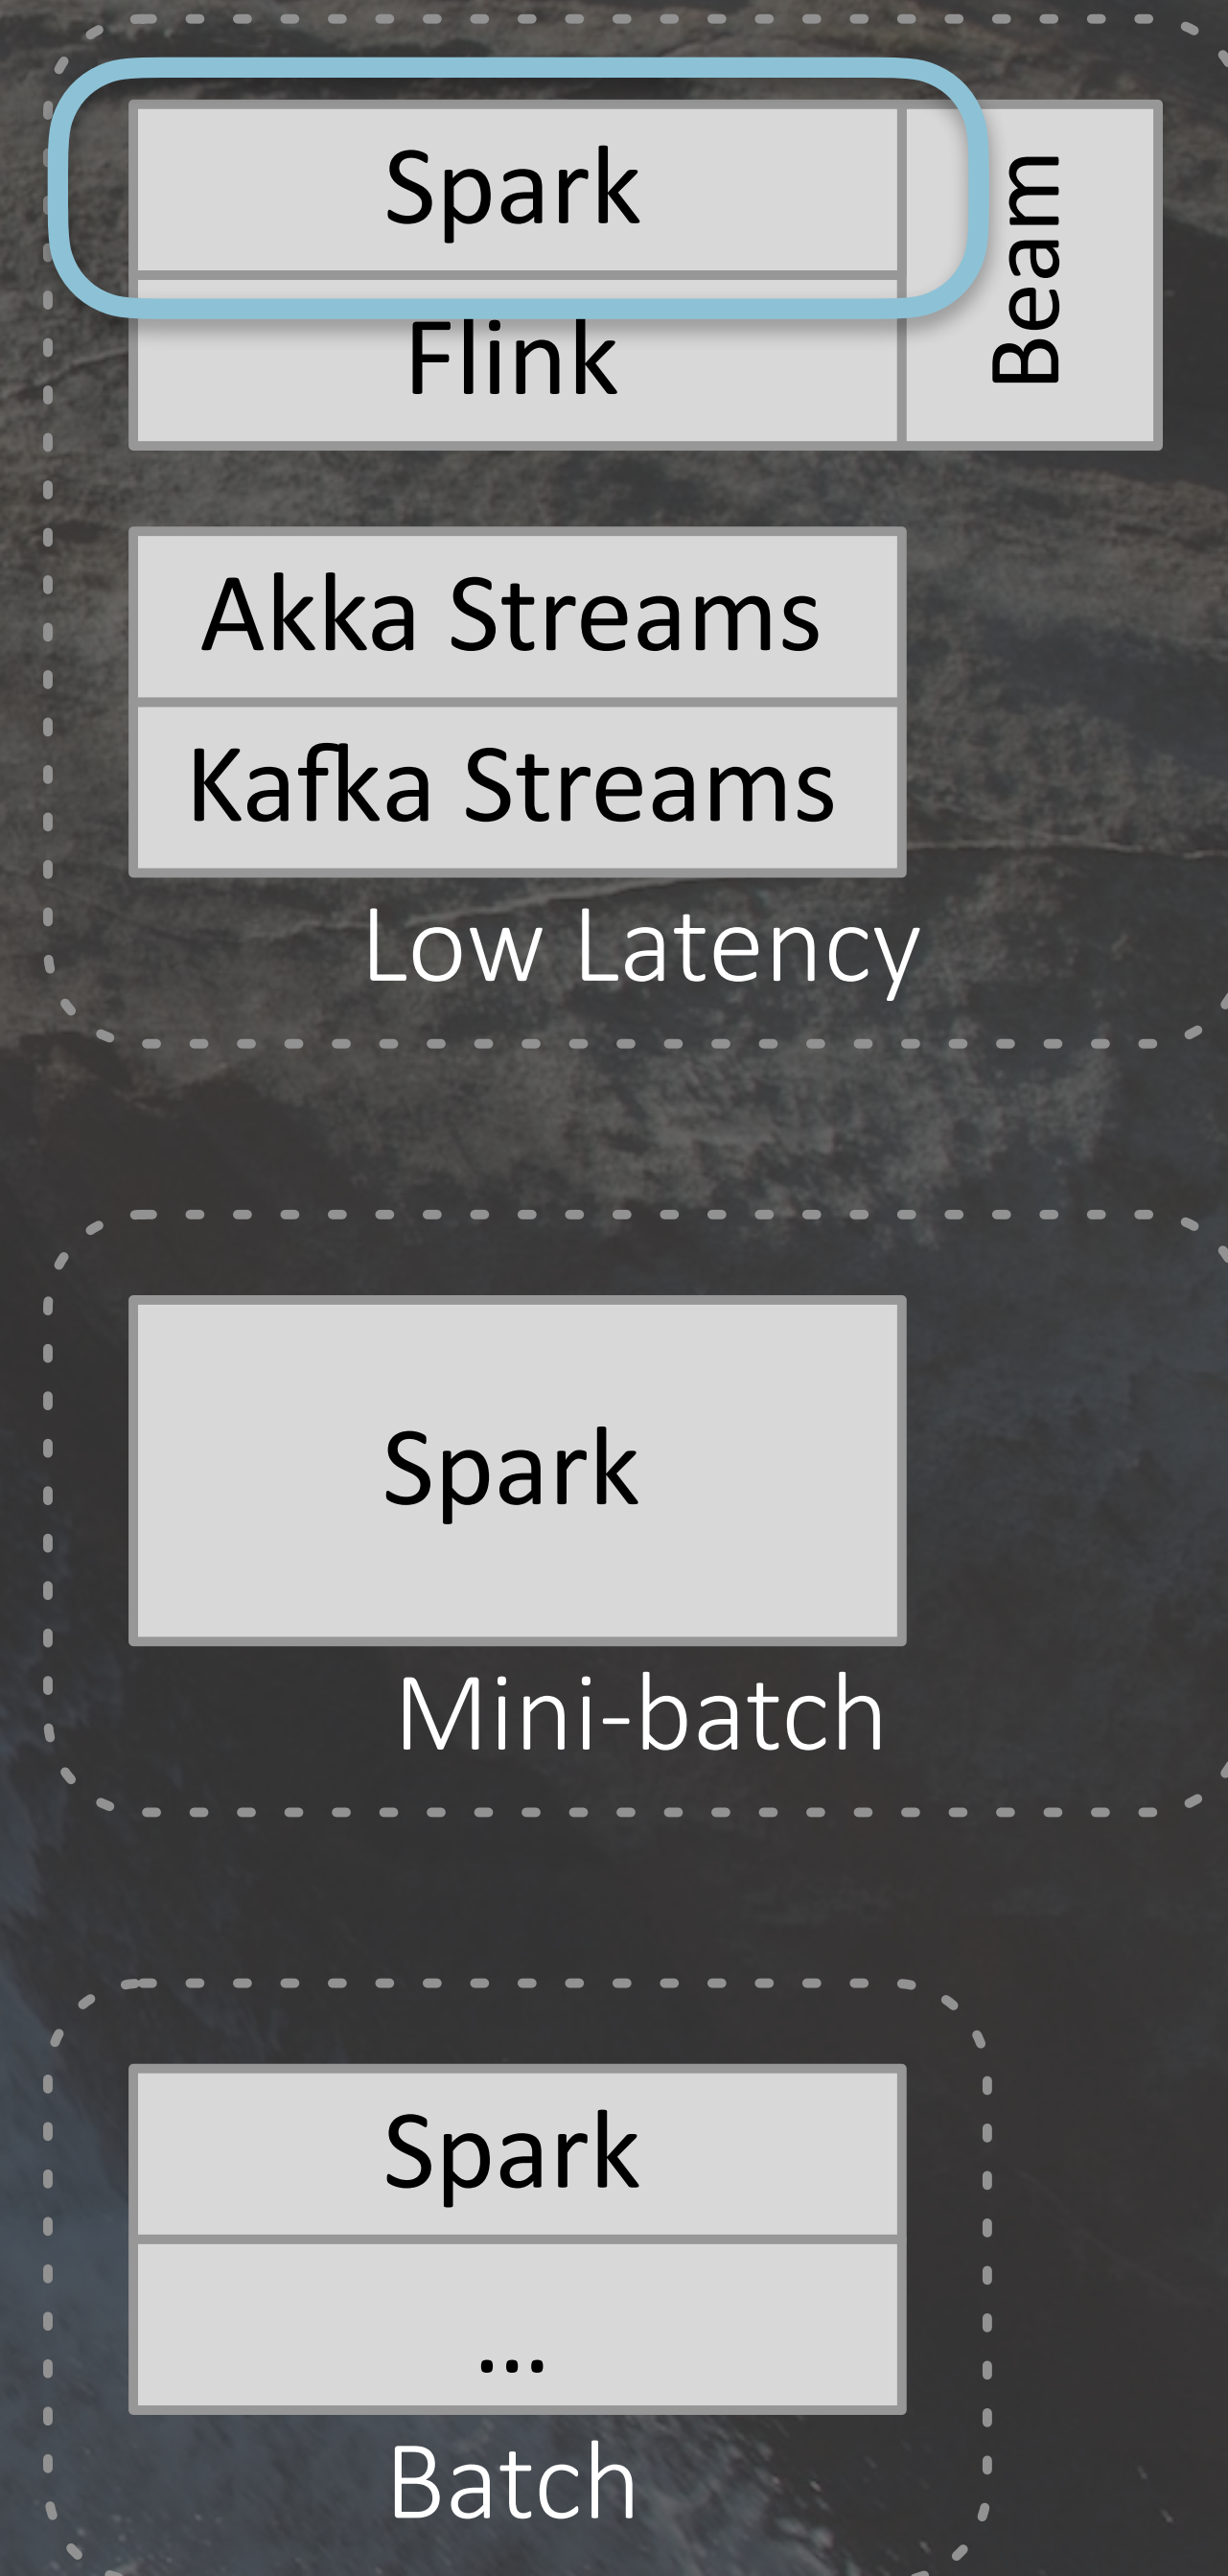


See these blog posts: <https://www.oreilly.com/people/09f01-tyler-akidau>

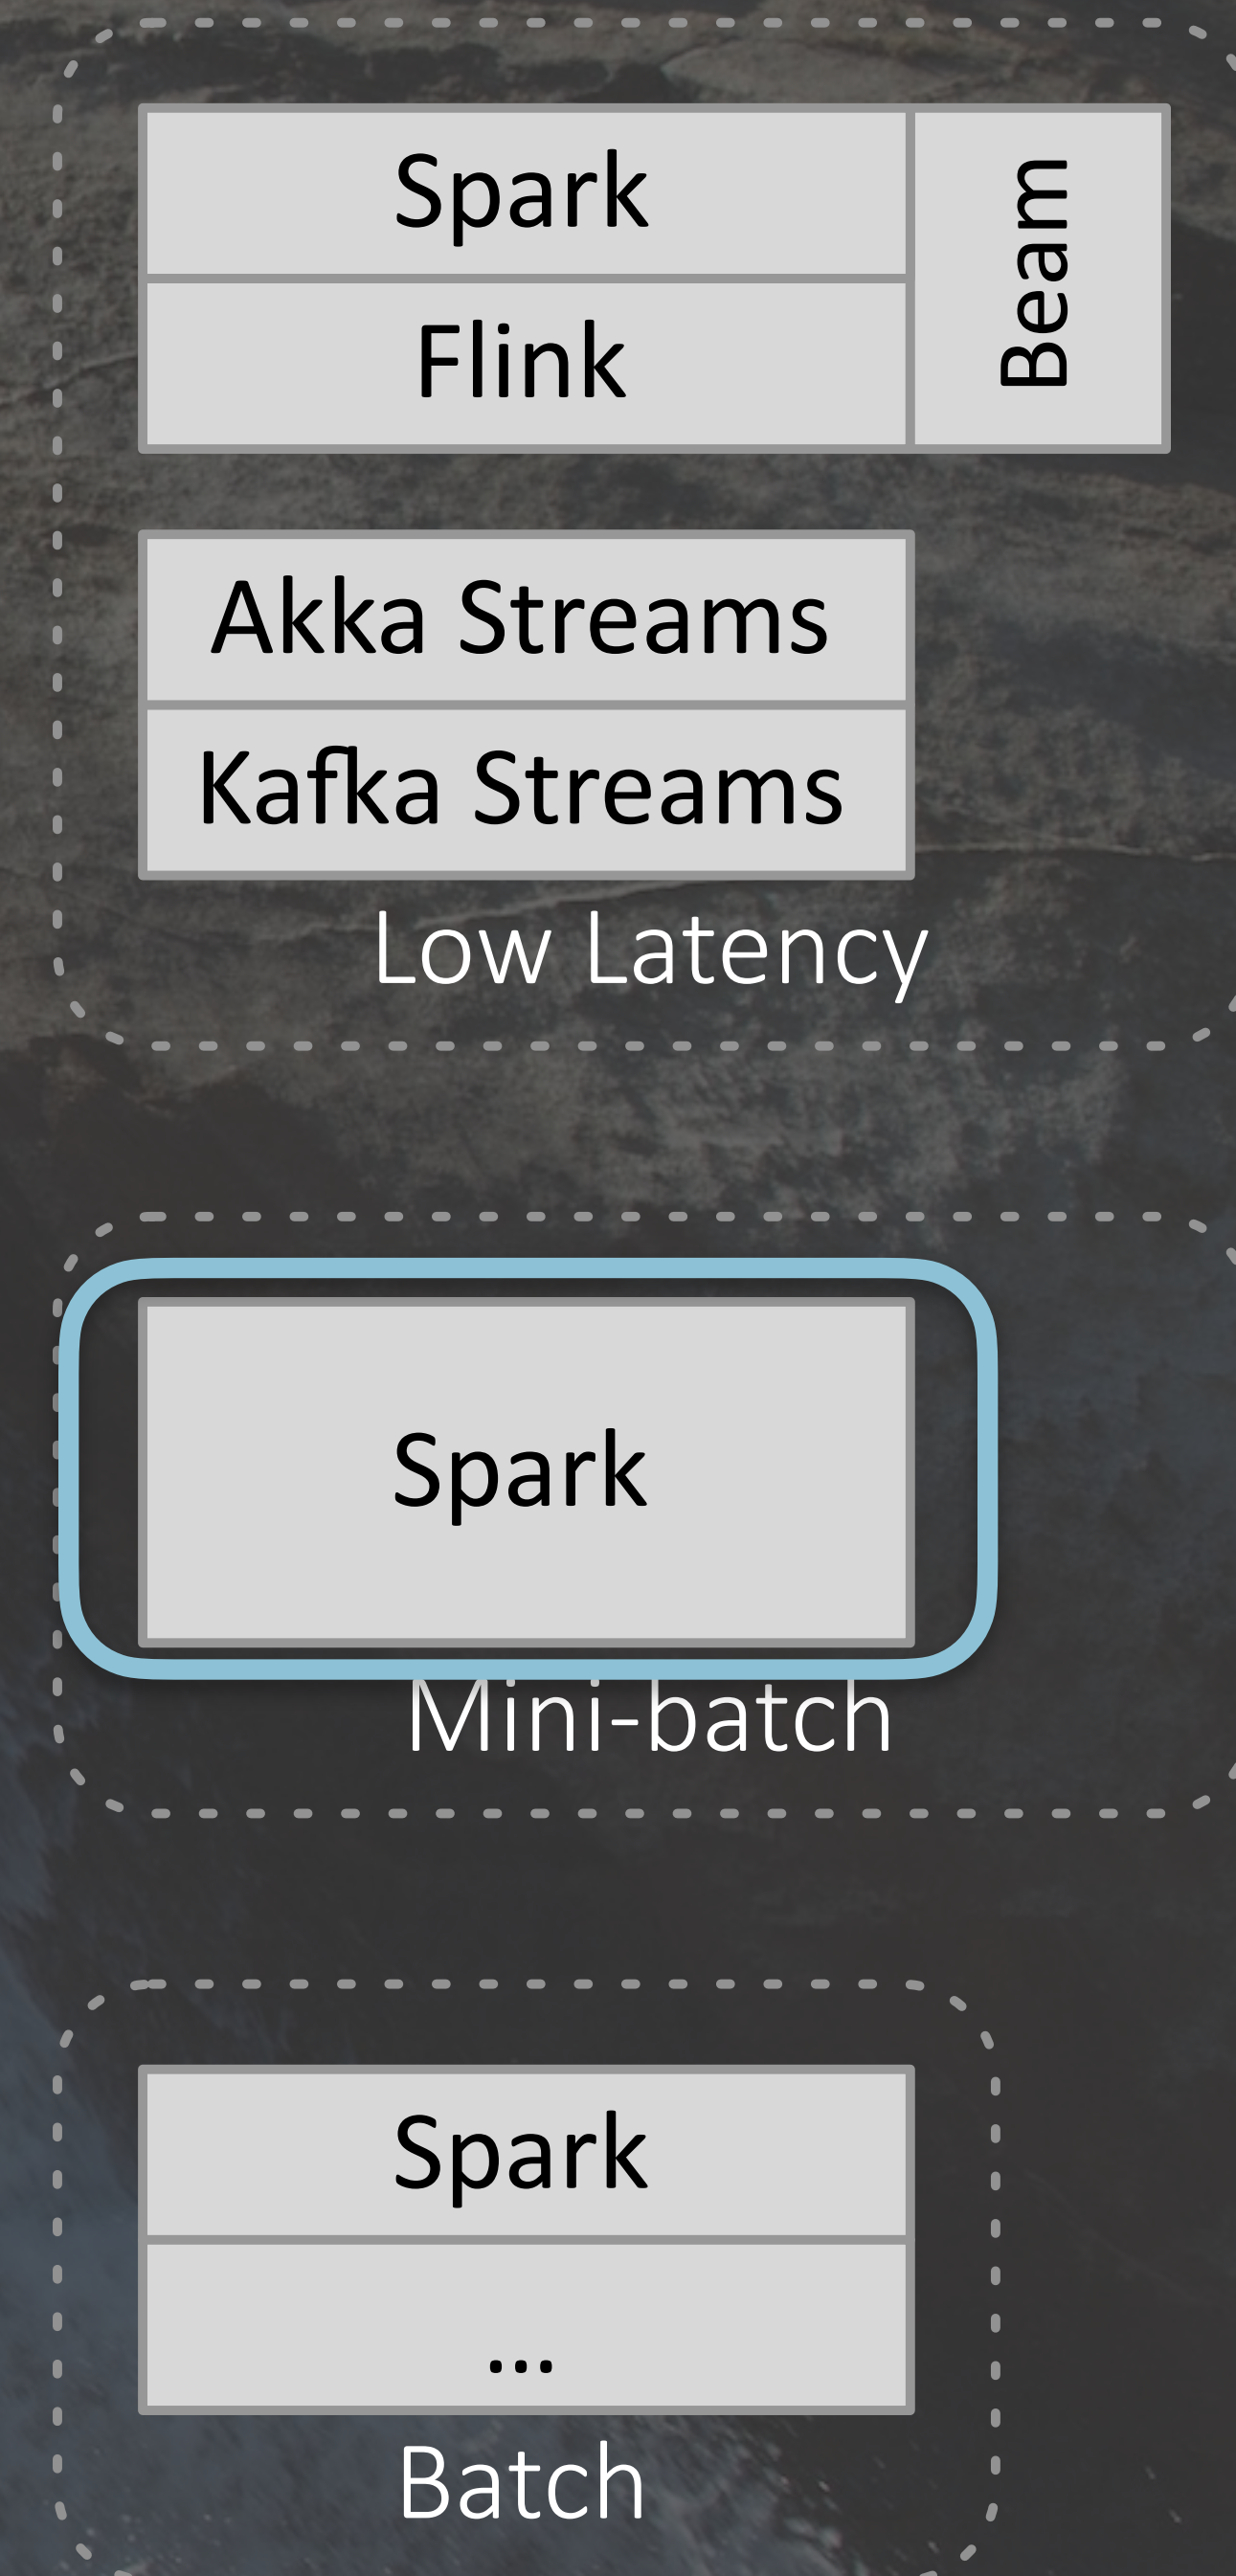




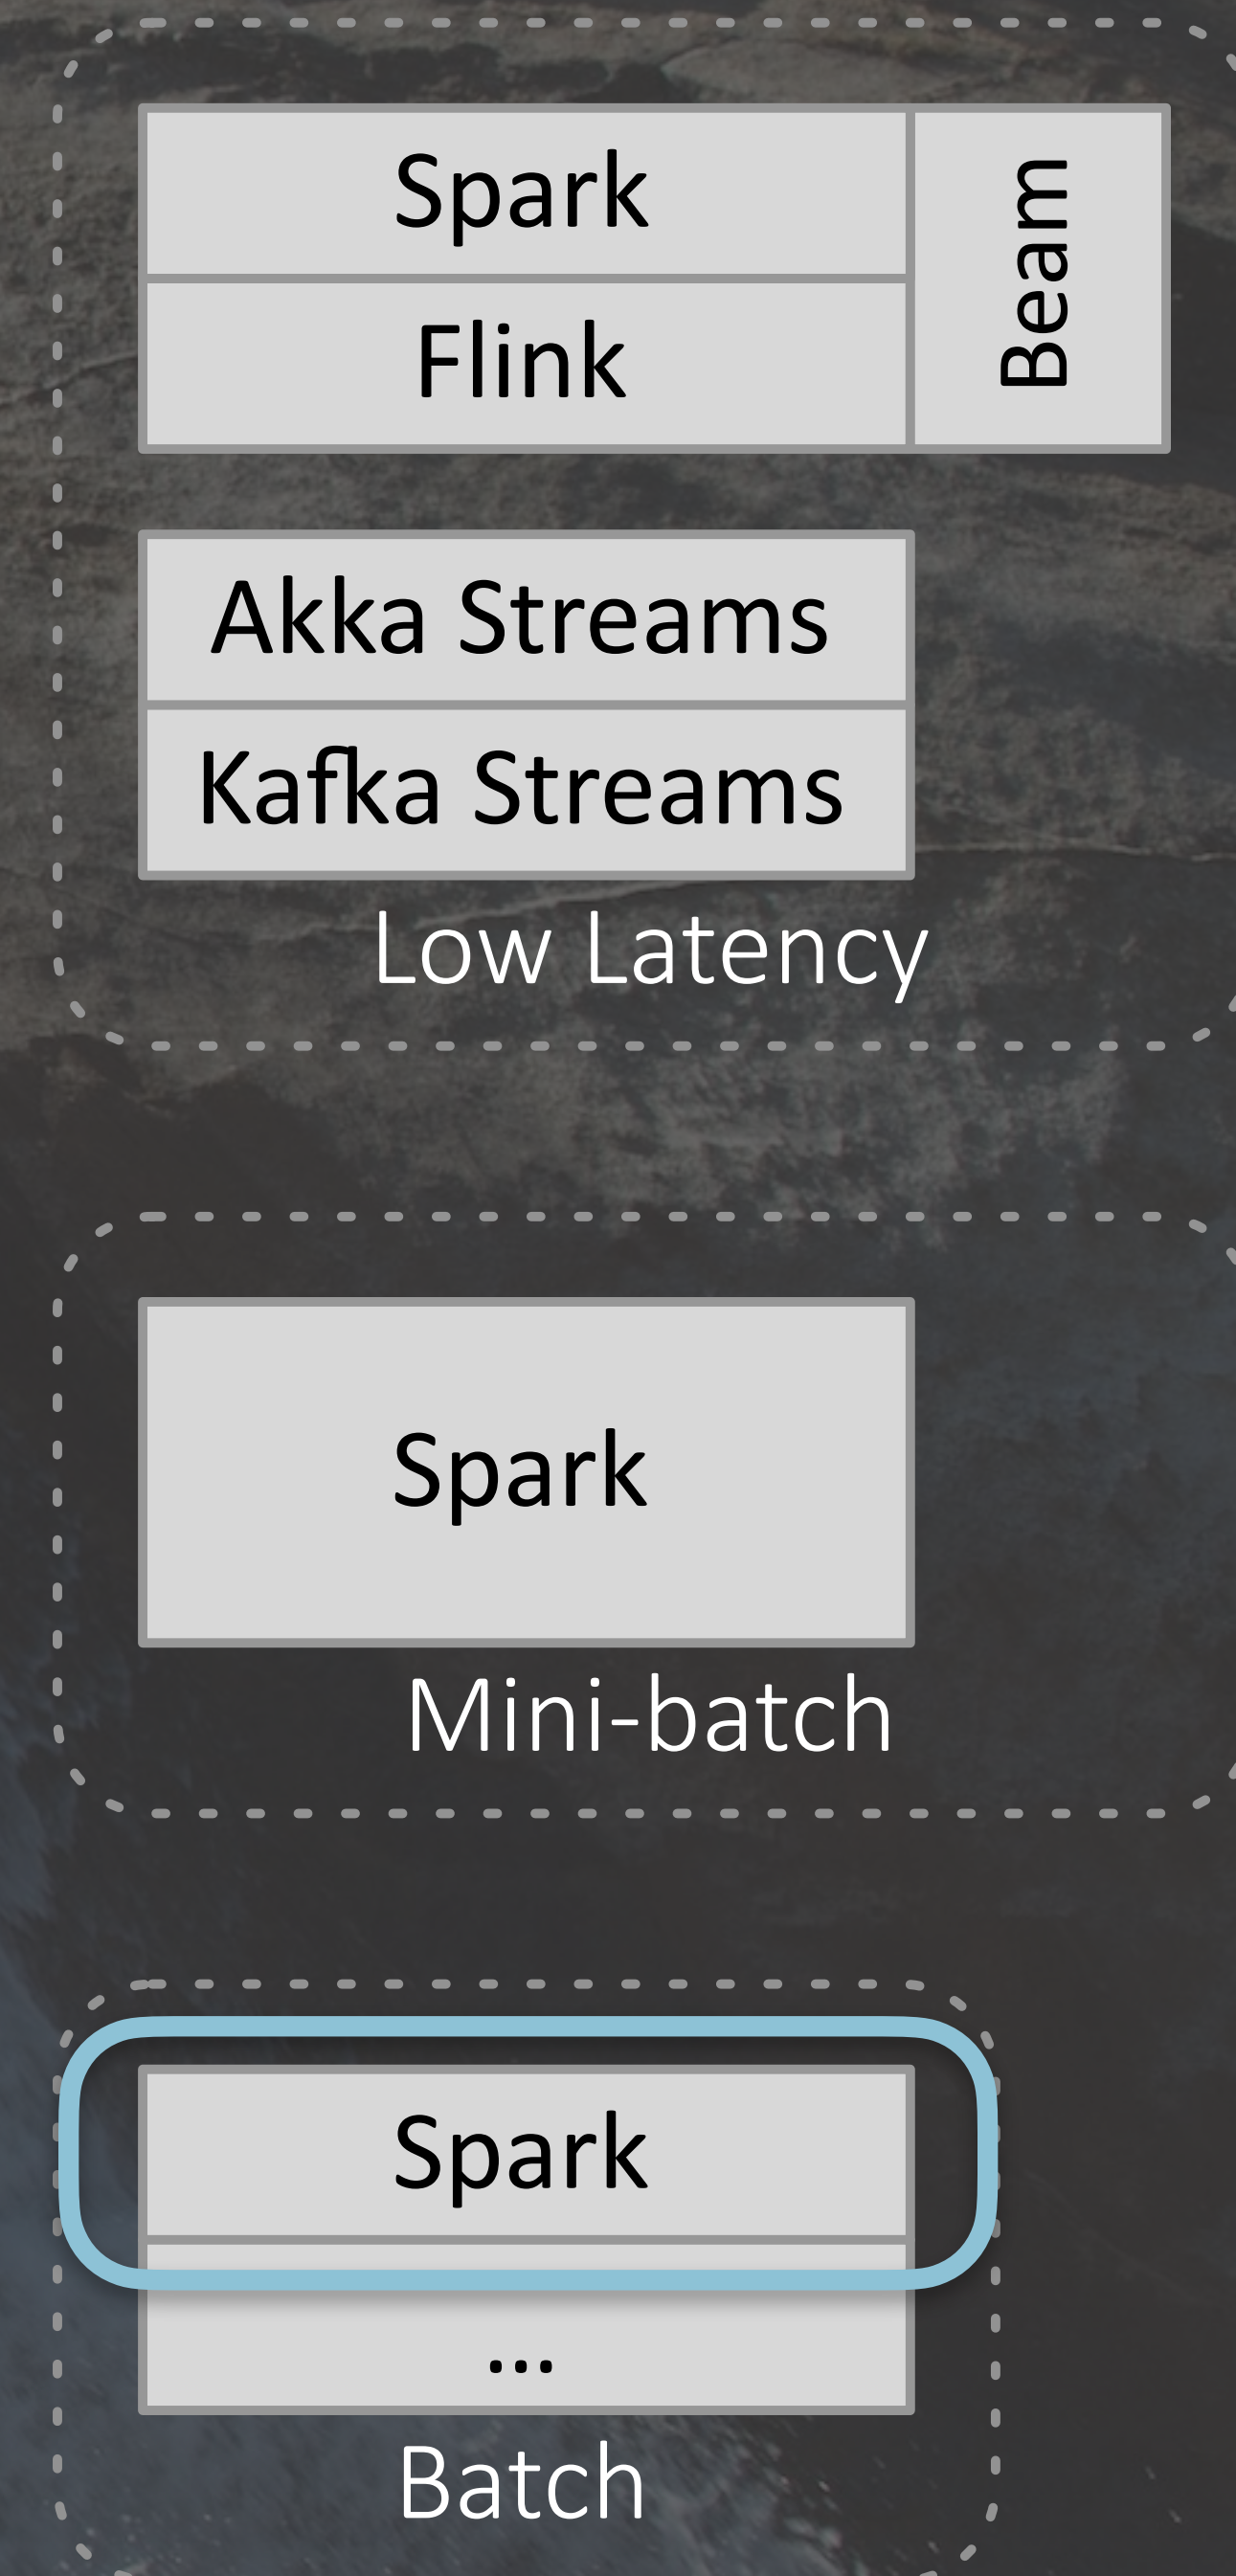
- Spark Structured Streaming
- “Dataset” - SQL
- Millisecond latency
- Ideal for Rich SQL, ML.



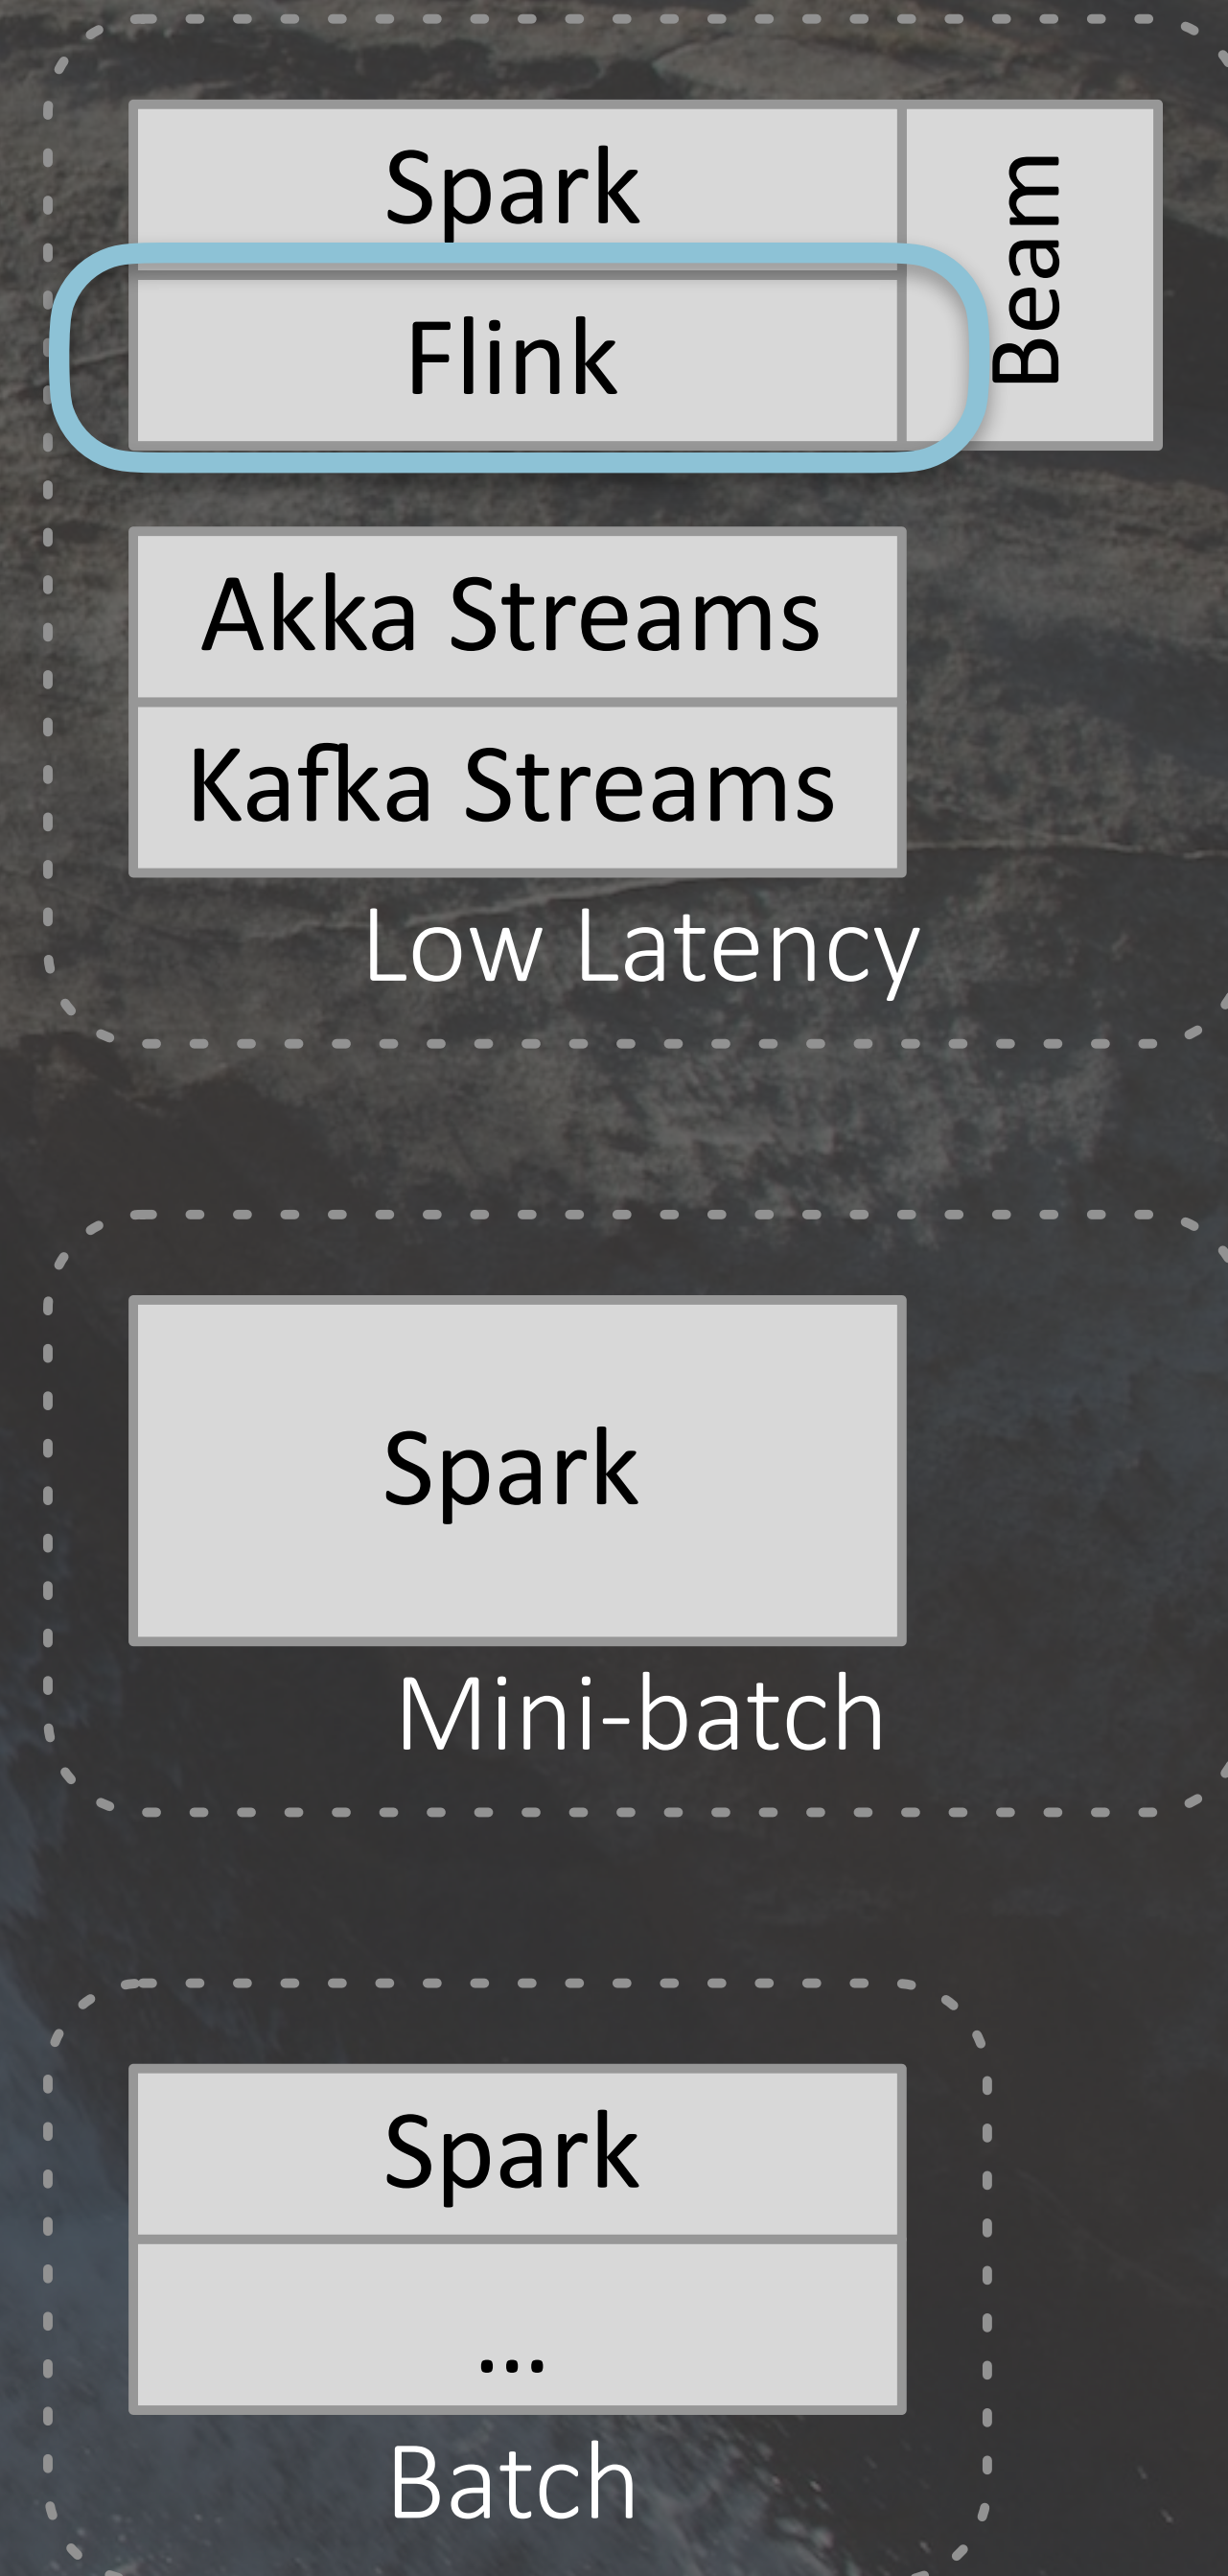
- Spark Streaming
- Mini-batch model
- “RDD” (dataflow) based
- ~0.5 sec latency
- Original model - obsolete



- Spark Batch
- Same Dataset and RDD features as streaming.
- Massive scalability
- Excellent performance



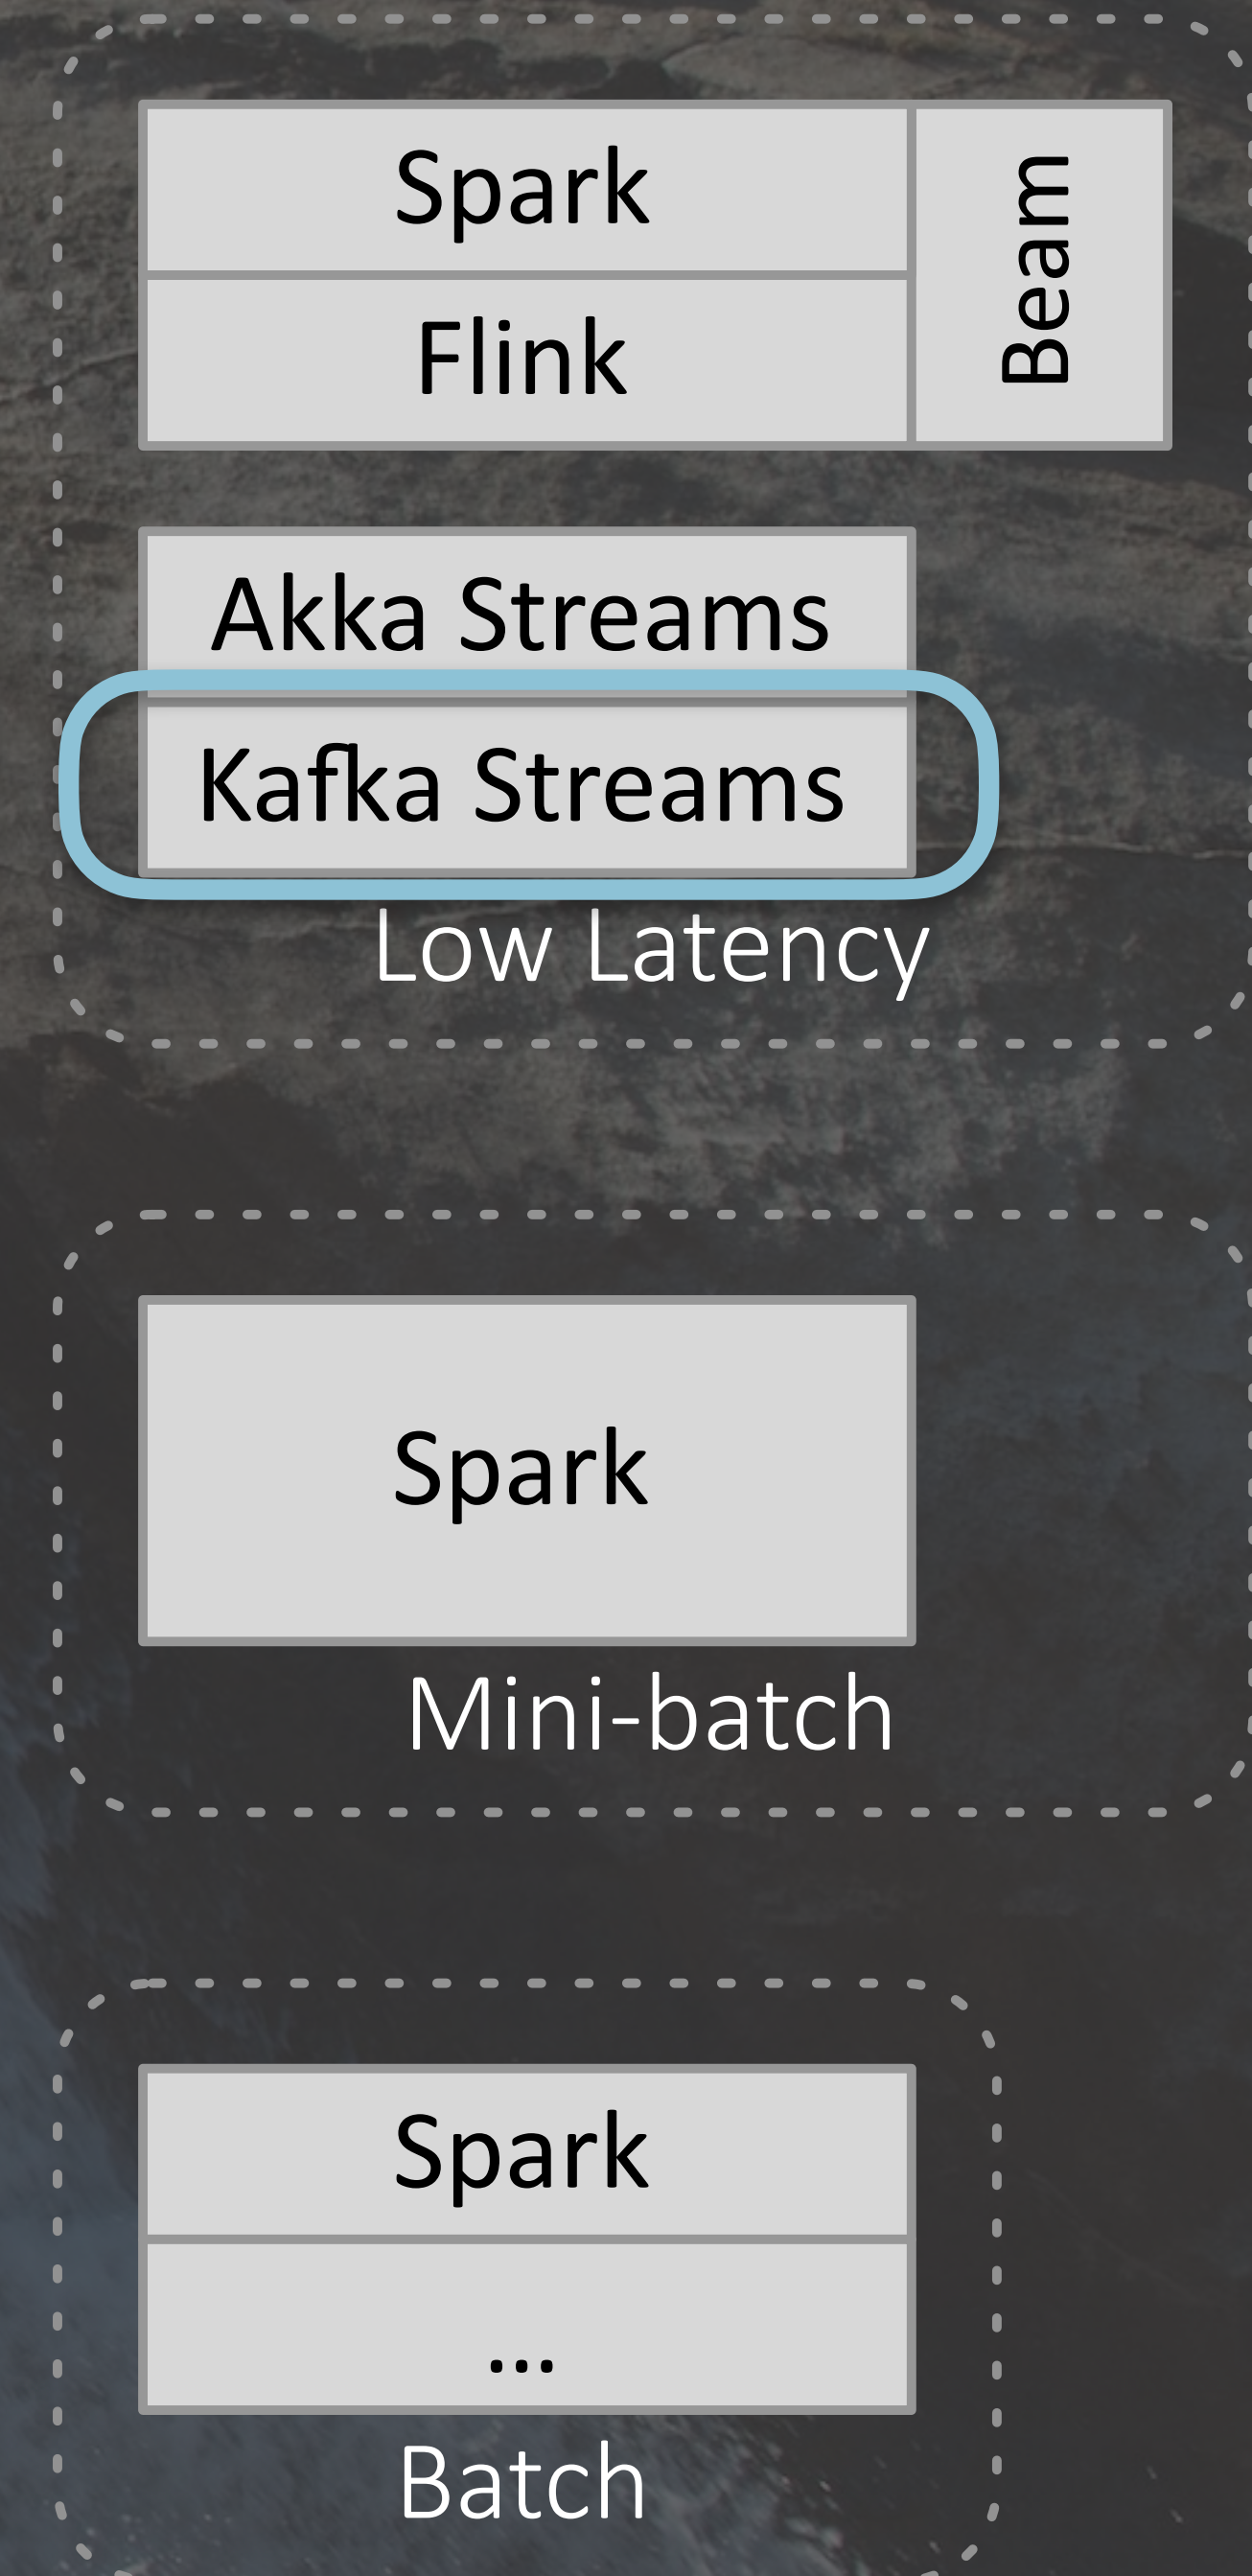
- Apache Flink
- High volume, low latency
- Sophisticated streaming (Beam) semantics
- SQL, evolving ML support




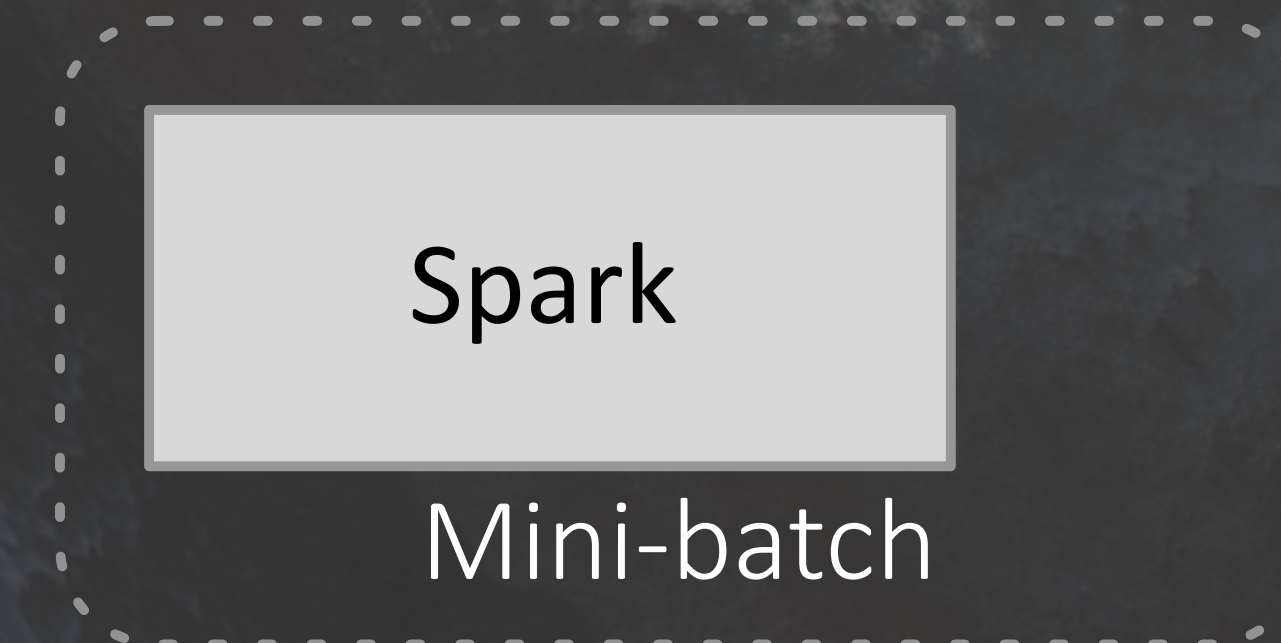
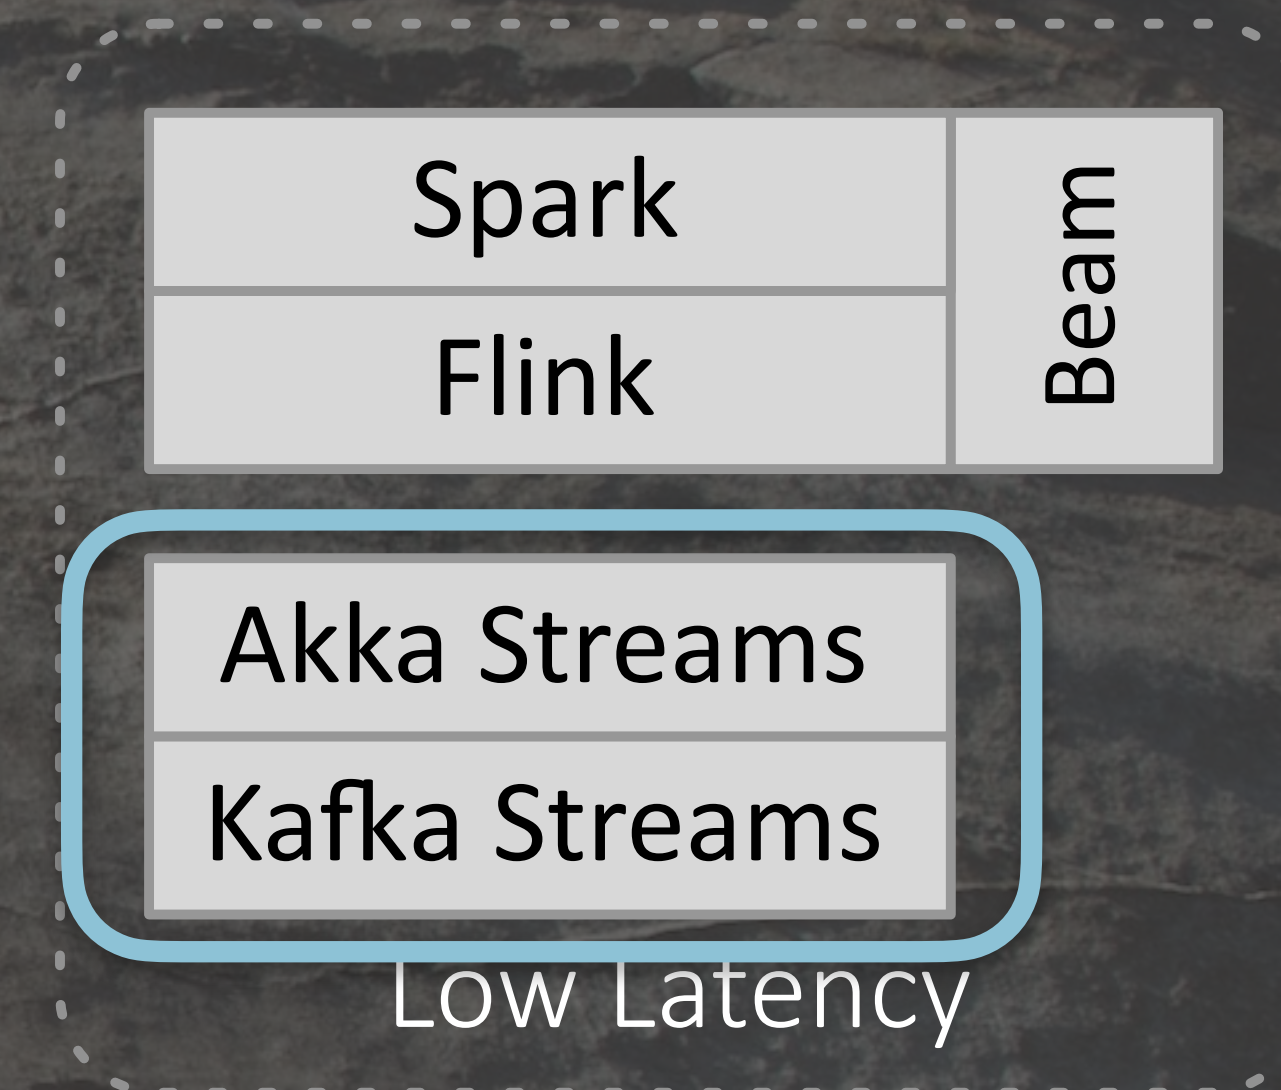
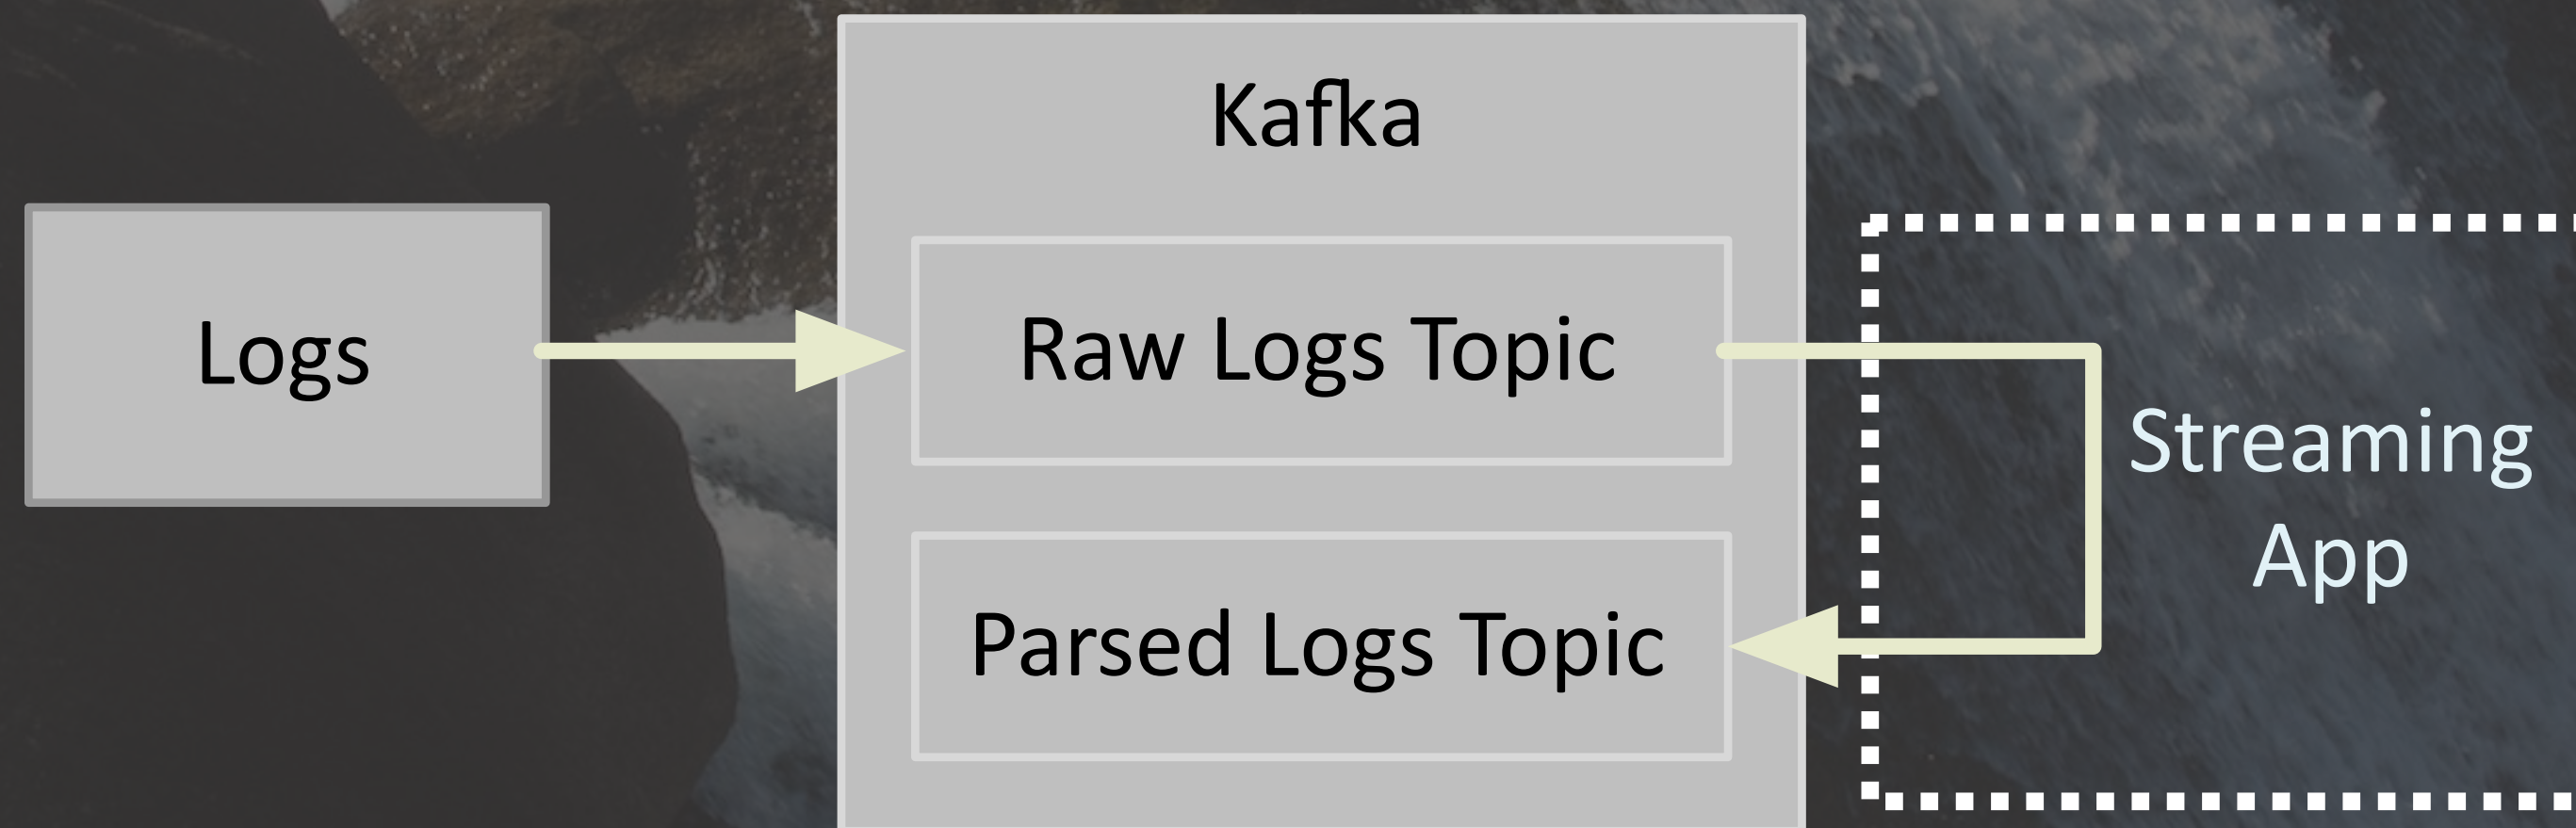
- Akka Streams
- Low latency
- Complex Event Processing
- Efficient, per event
- Mid-volume pipelines




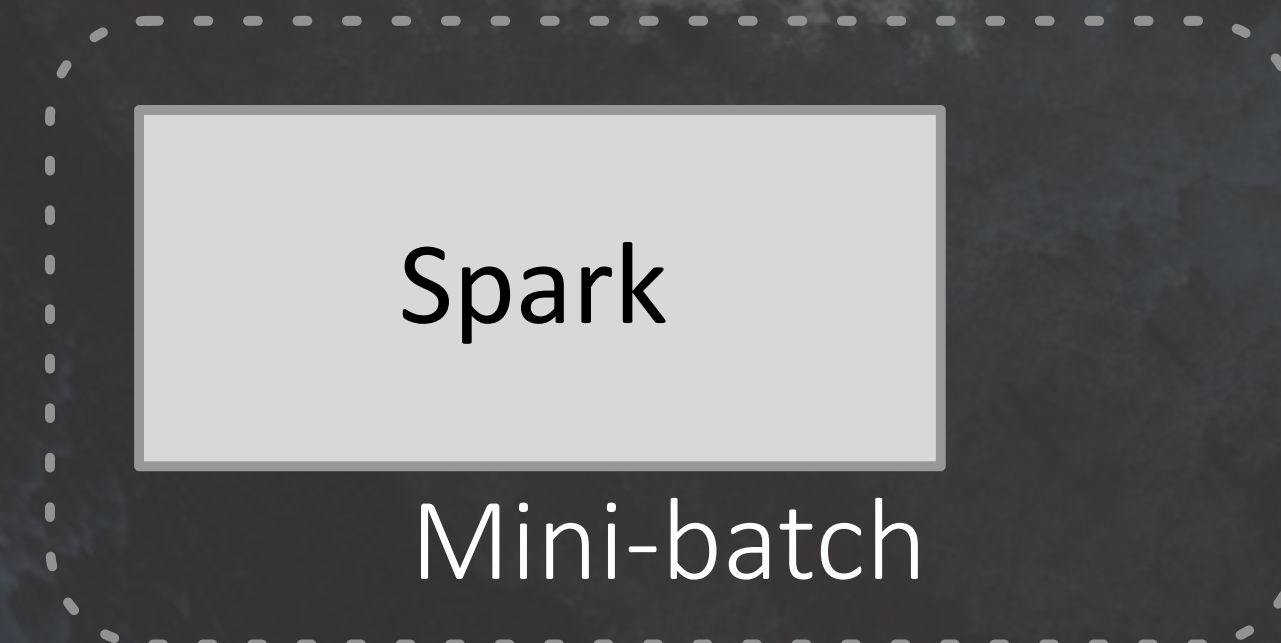
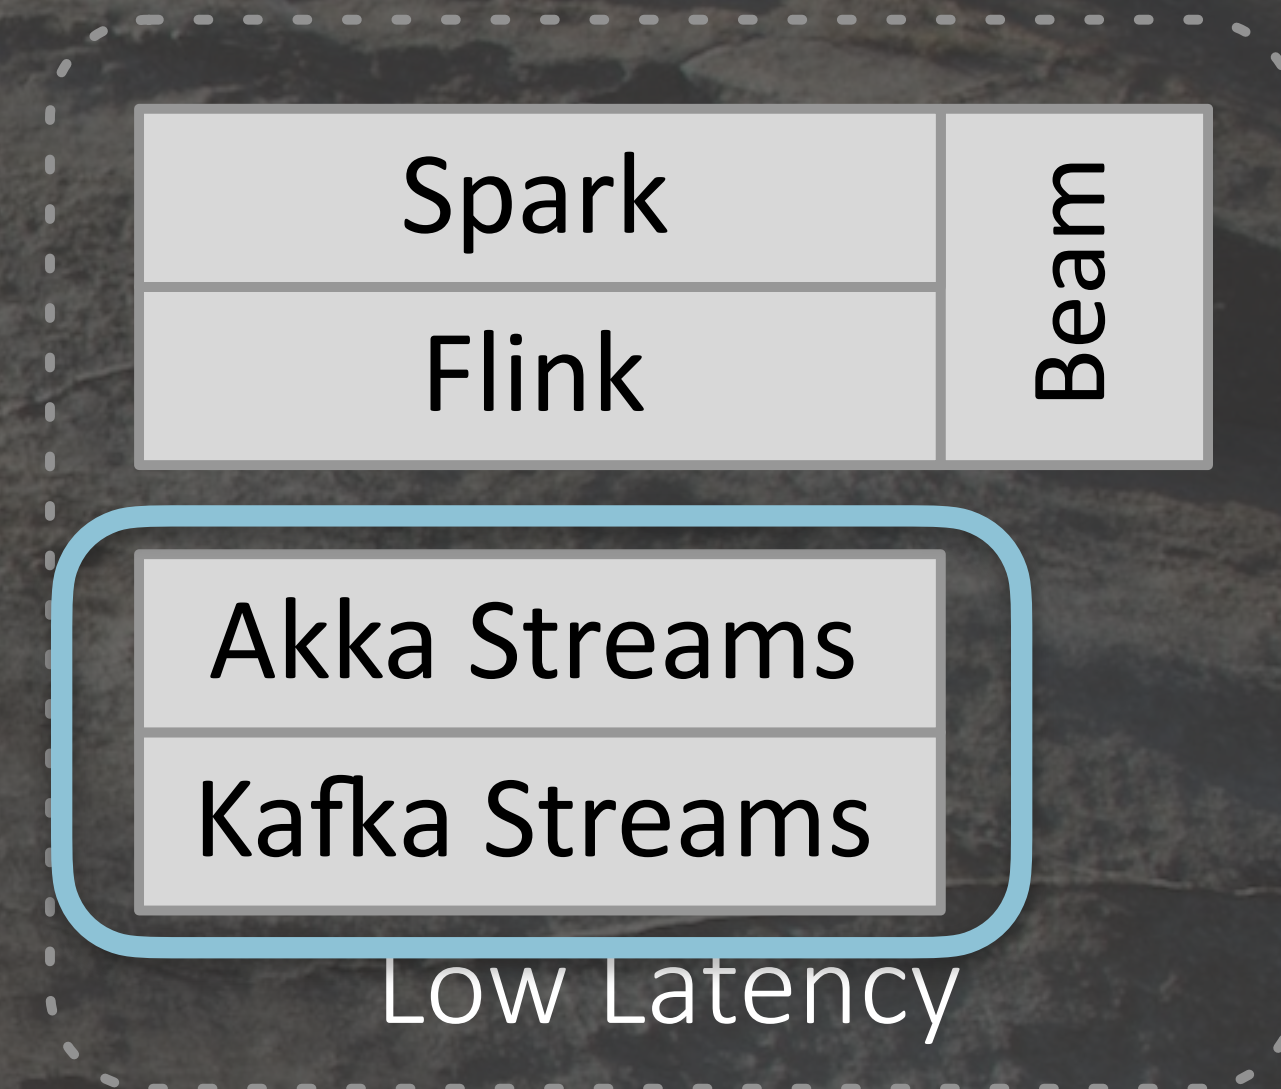
- Kafka Streams
- Low overhead Kafka topic processing
- Ideal for ETL and aggregations



- Akka and Kafka Streams
- “Exactly once” with transactions



- Akka and Kafka Streams
- Neither have built-in support for state checkpointing



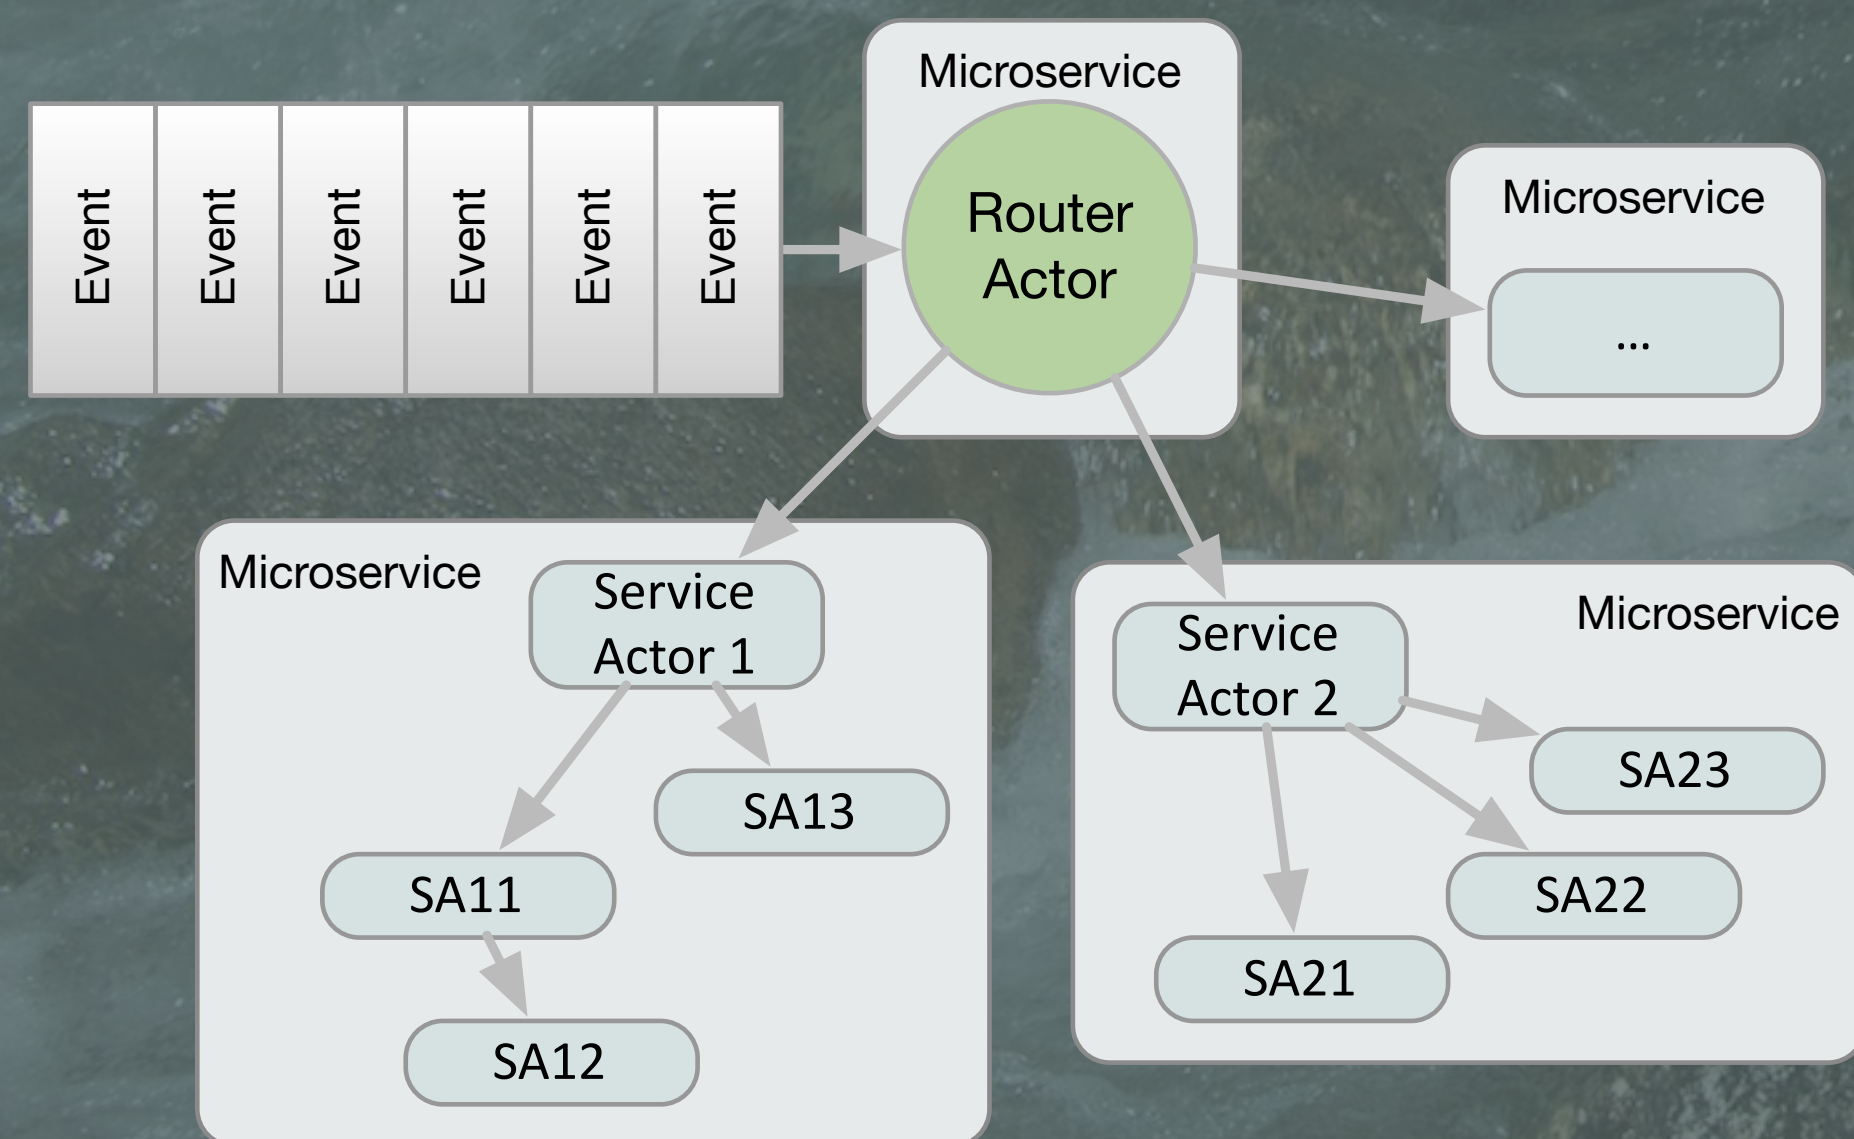
• Process data individually or in bulk?



Each grew out of one end of this spectrum...




Event-driven μ -services



• “Record-centric” μ -services

```
SELECT COUNT(*)  
FROM my-iot-data  
GROUP BY zip-code
```


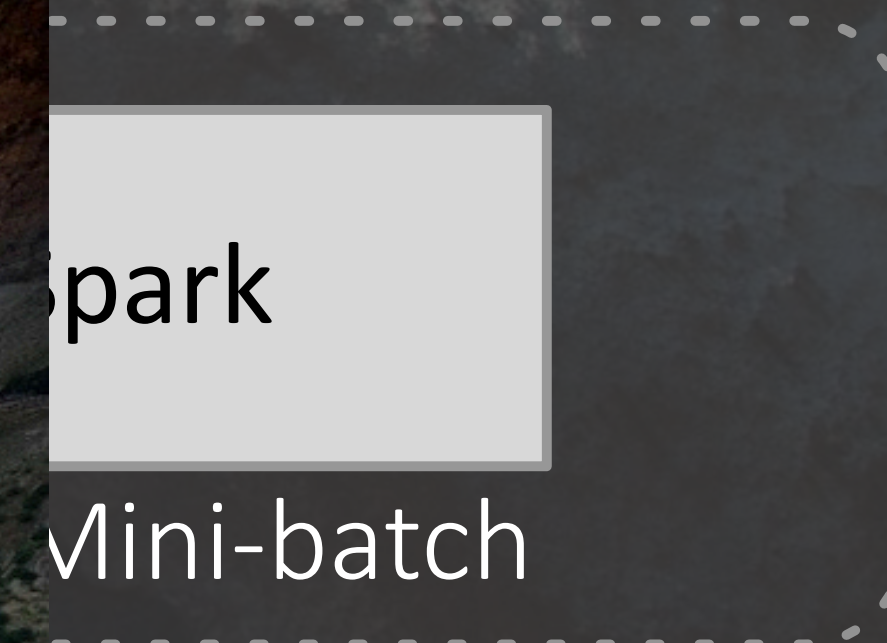
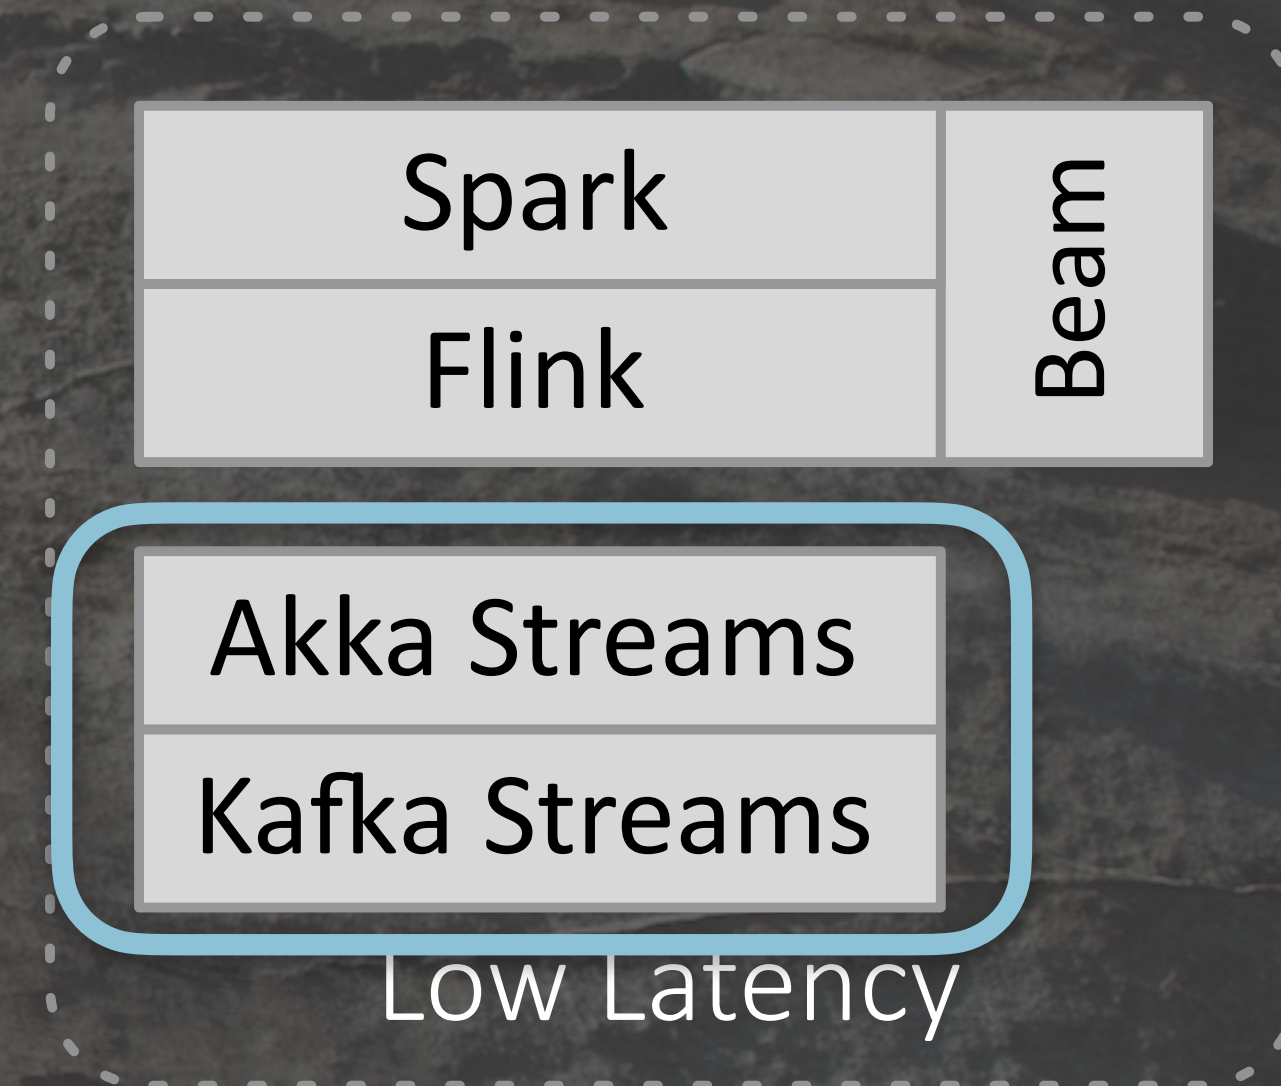


Events

Records

- Akka Streams vs. Kafka Streams talk

- Also at polyglotprogramming.com/talks/



Streaming Microservices

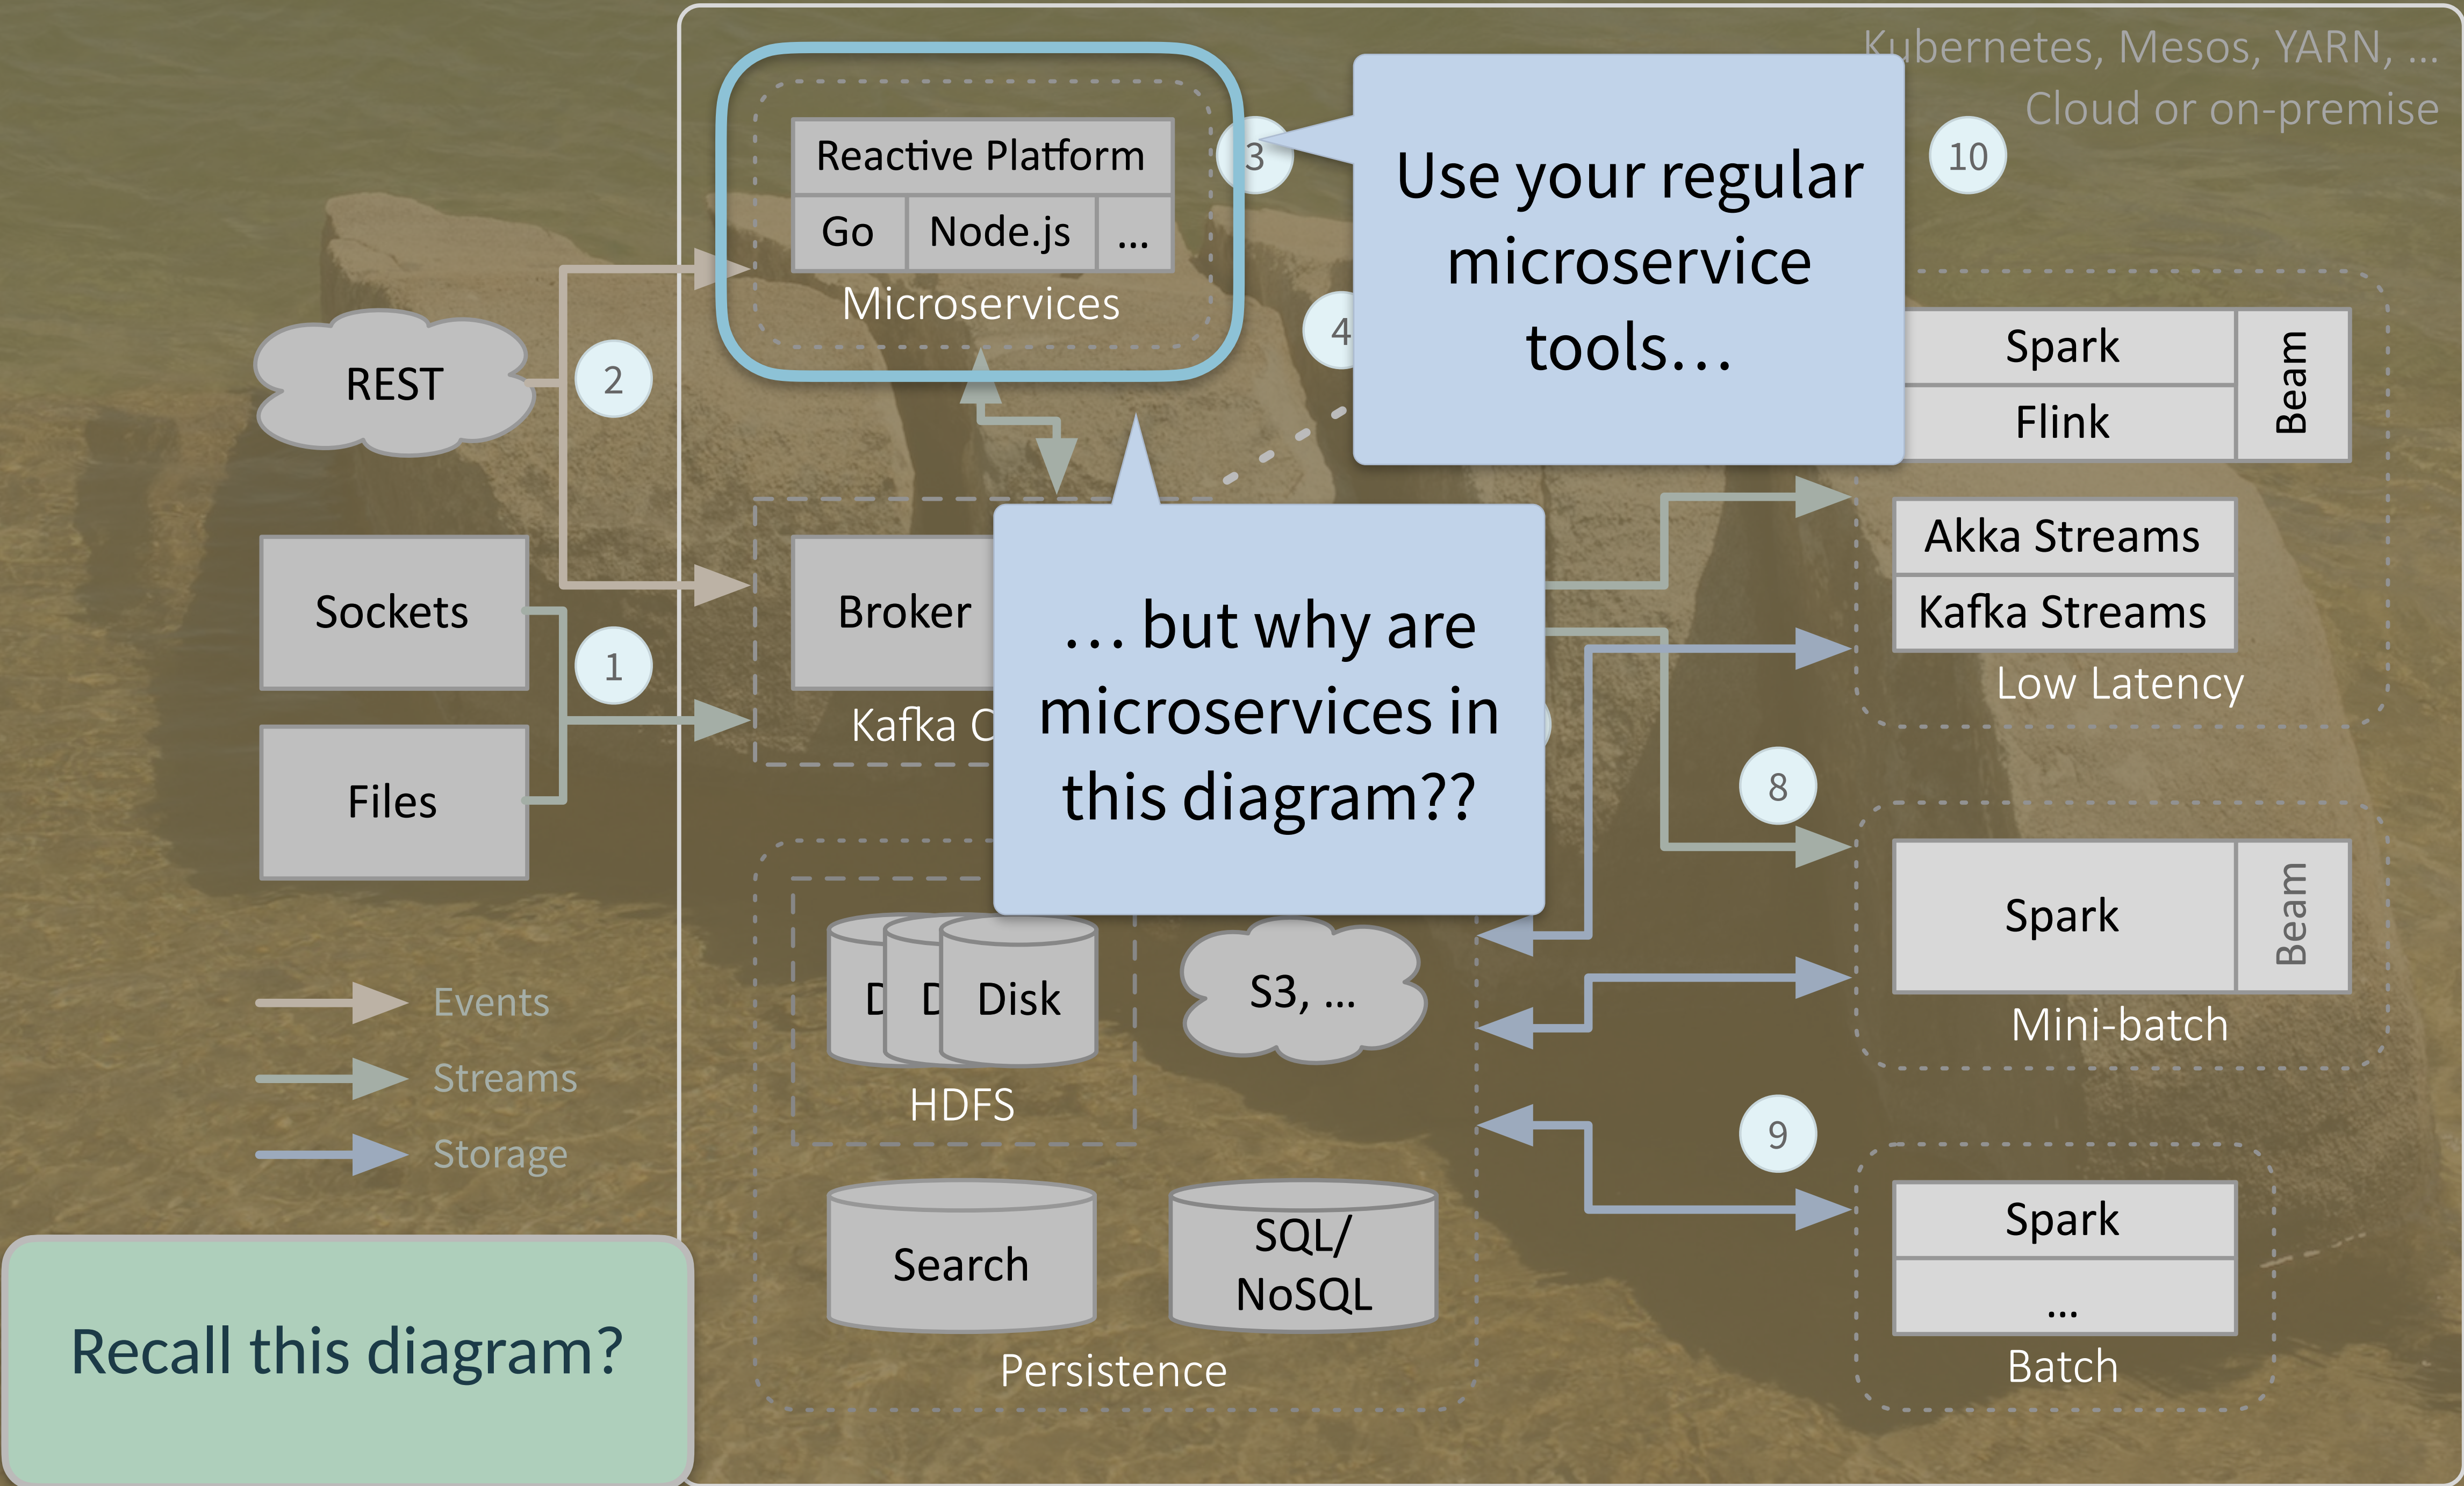
With Akka Streams and Kafka Streams

 Lightbend

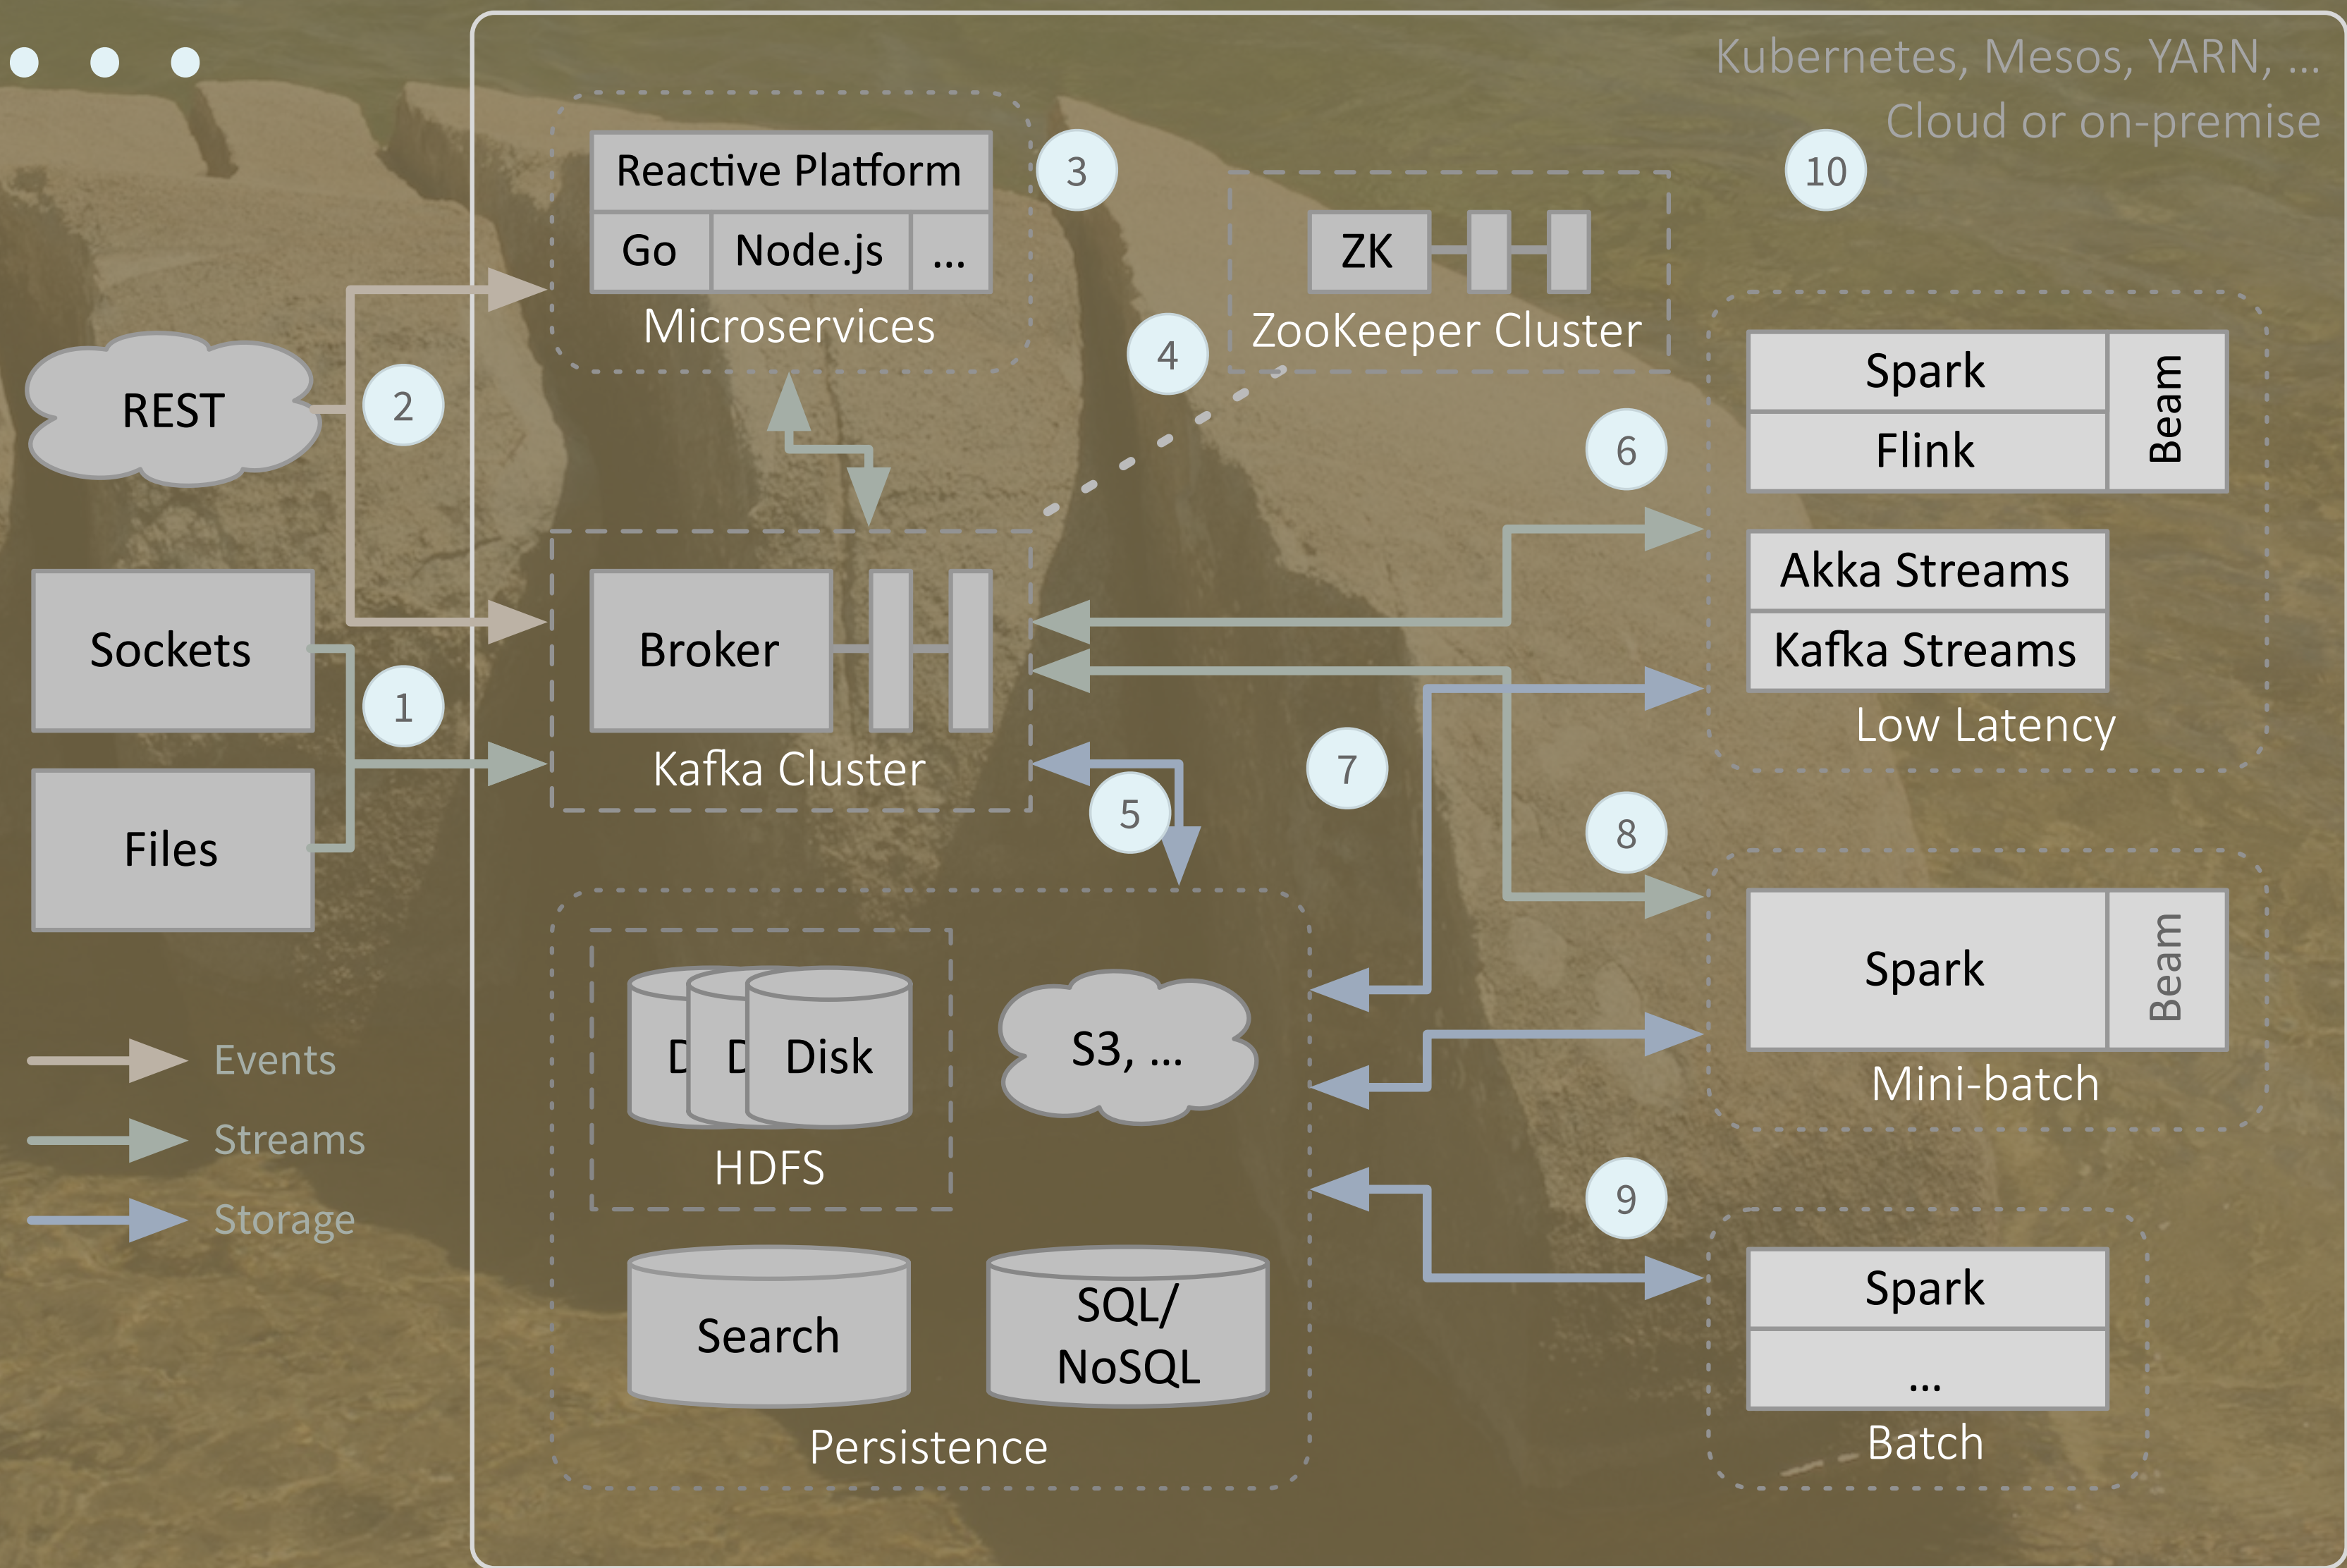
Dean Wampler, Ph.D.
dean@lightbend.com
@deanwampler



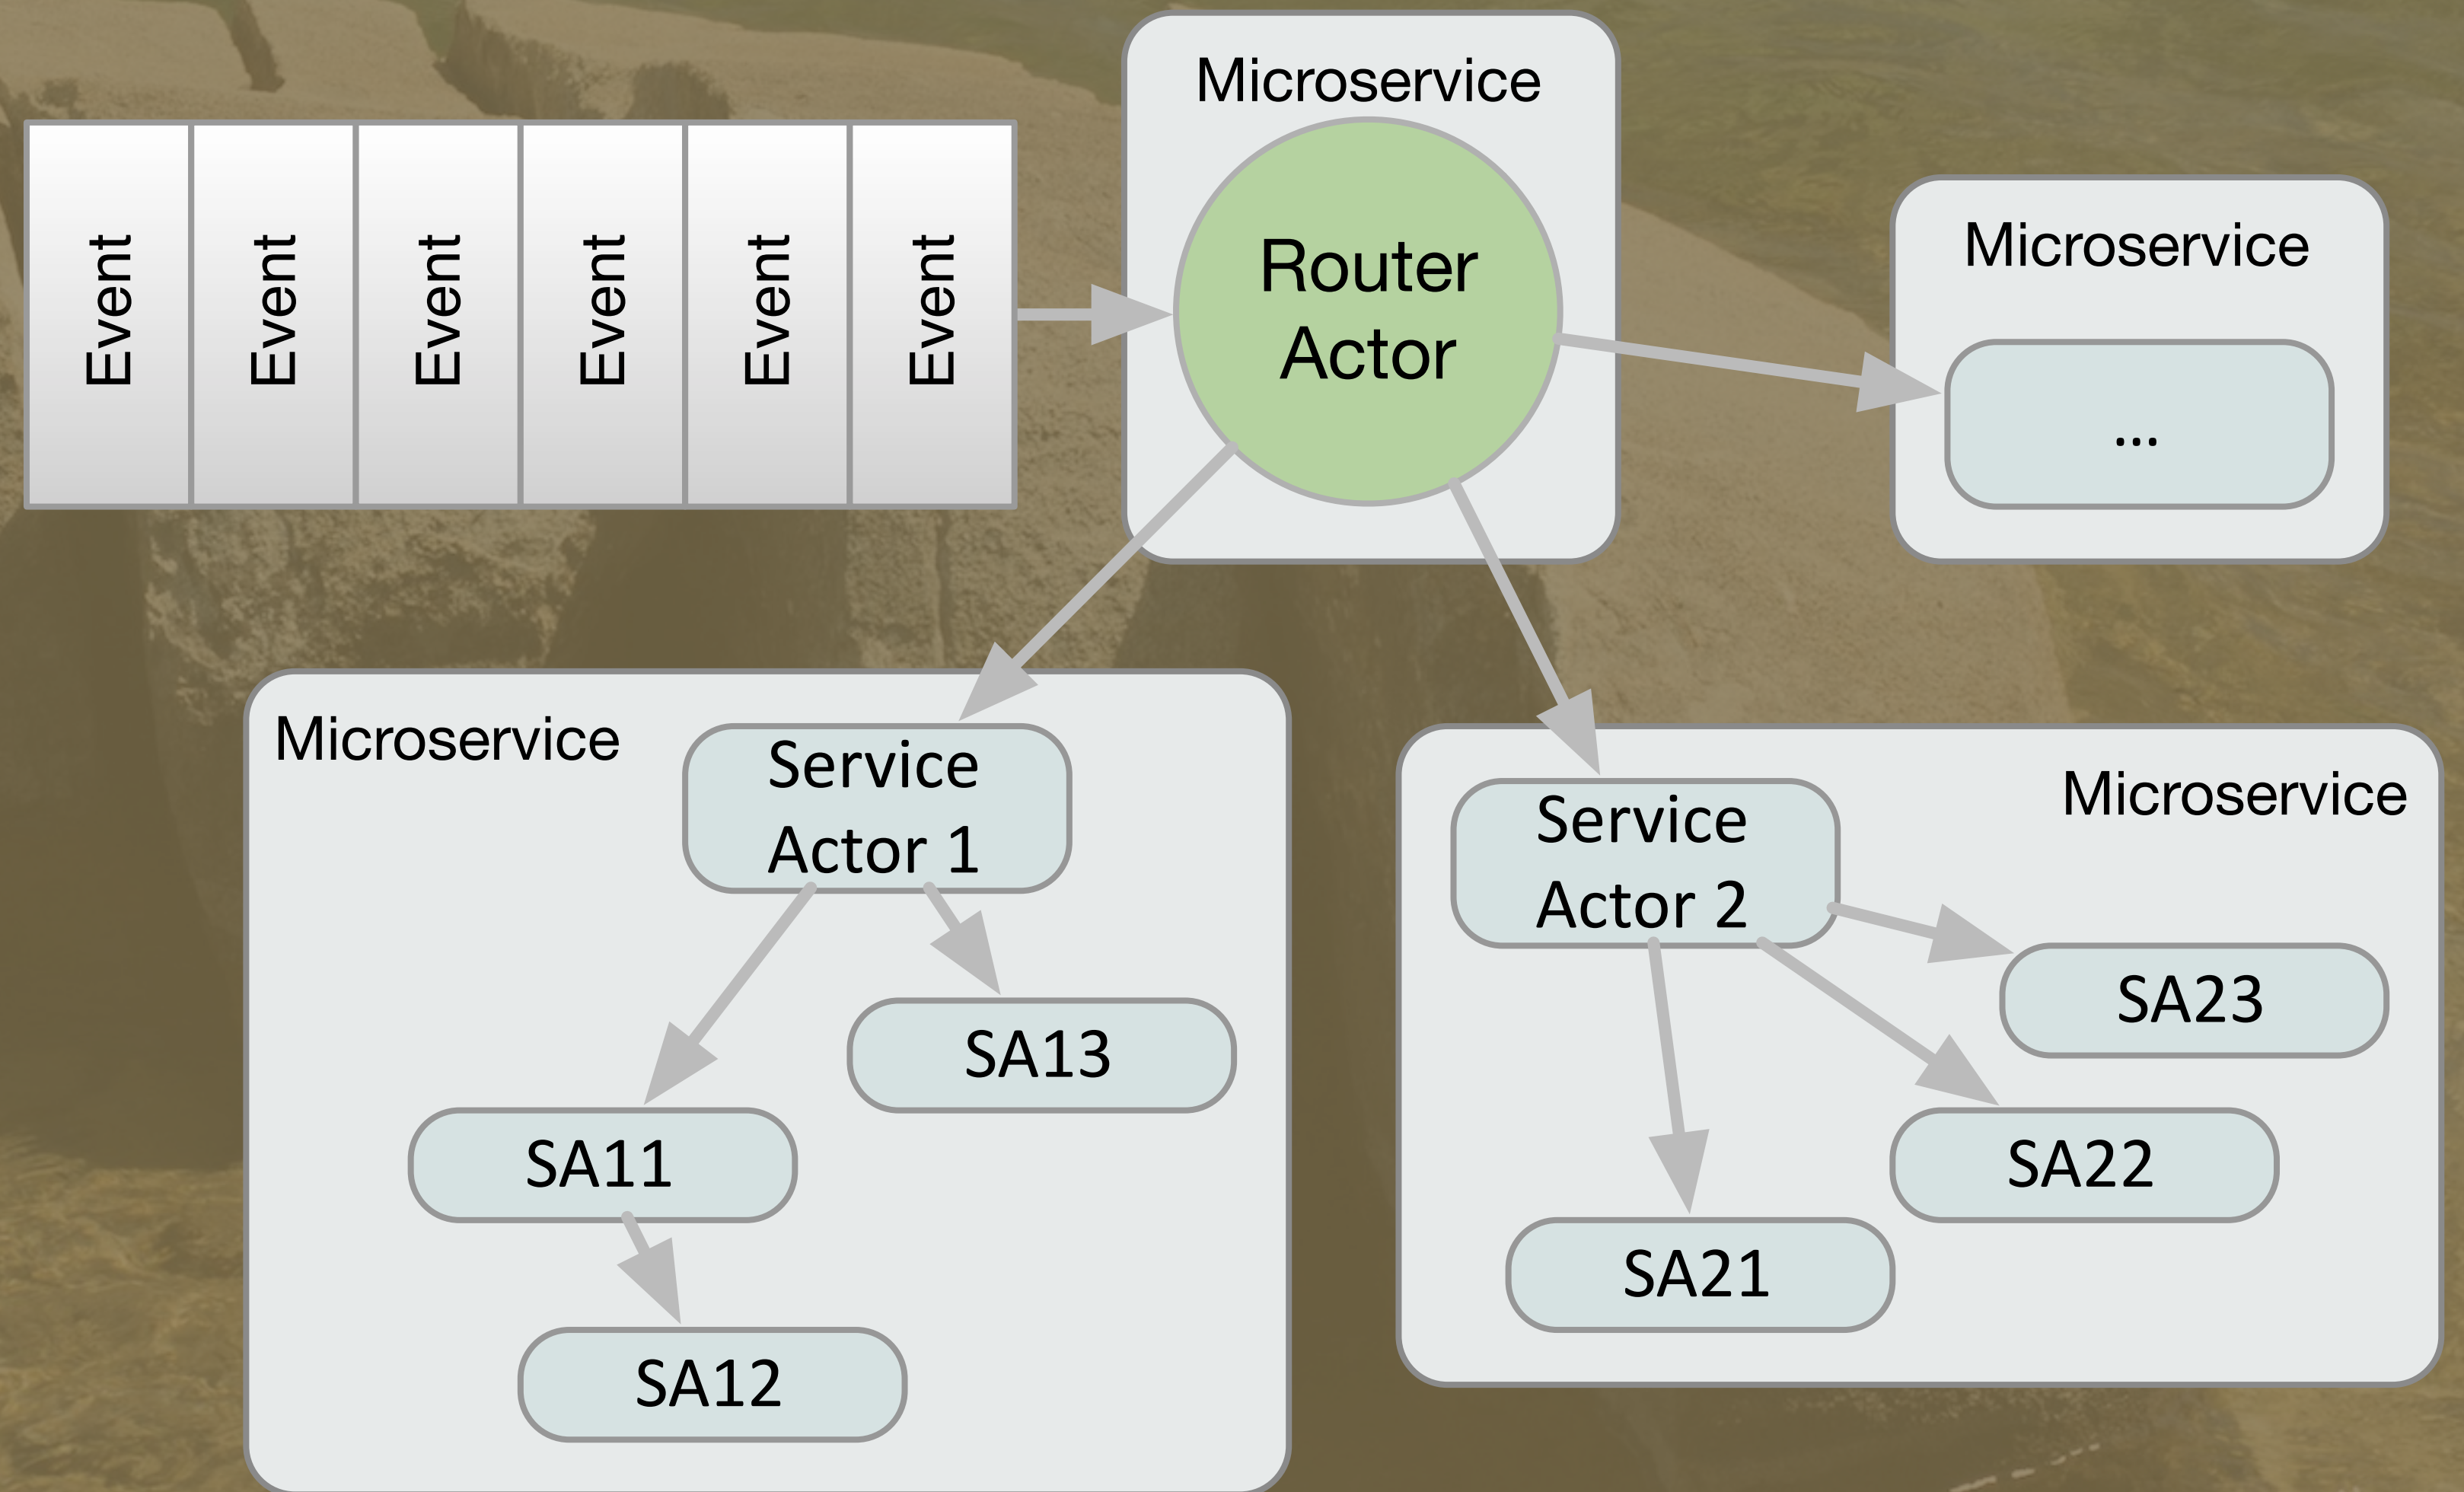
Microservices and Fast Data



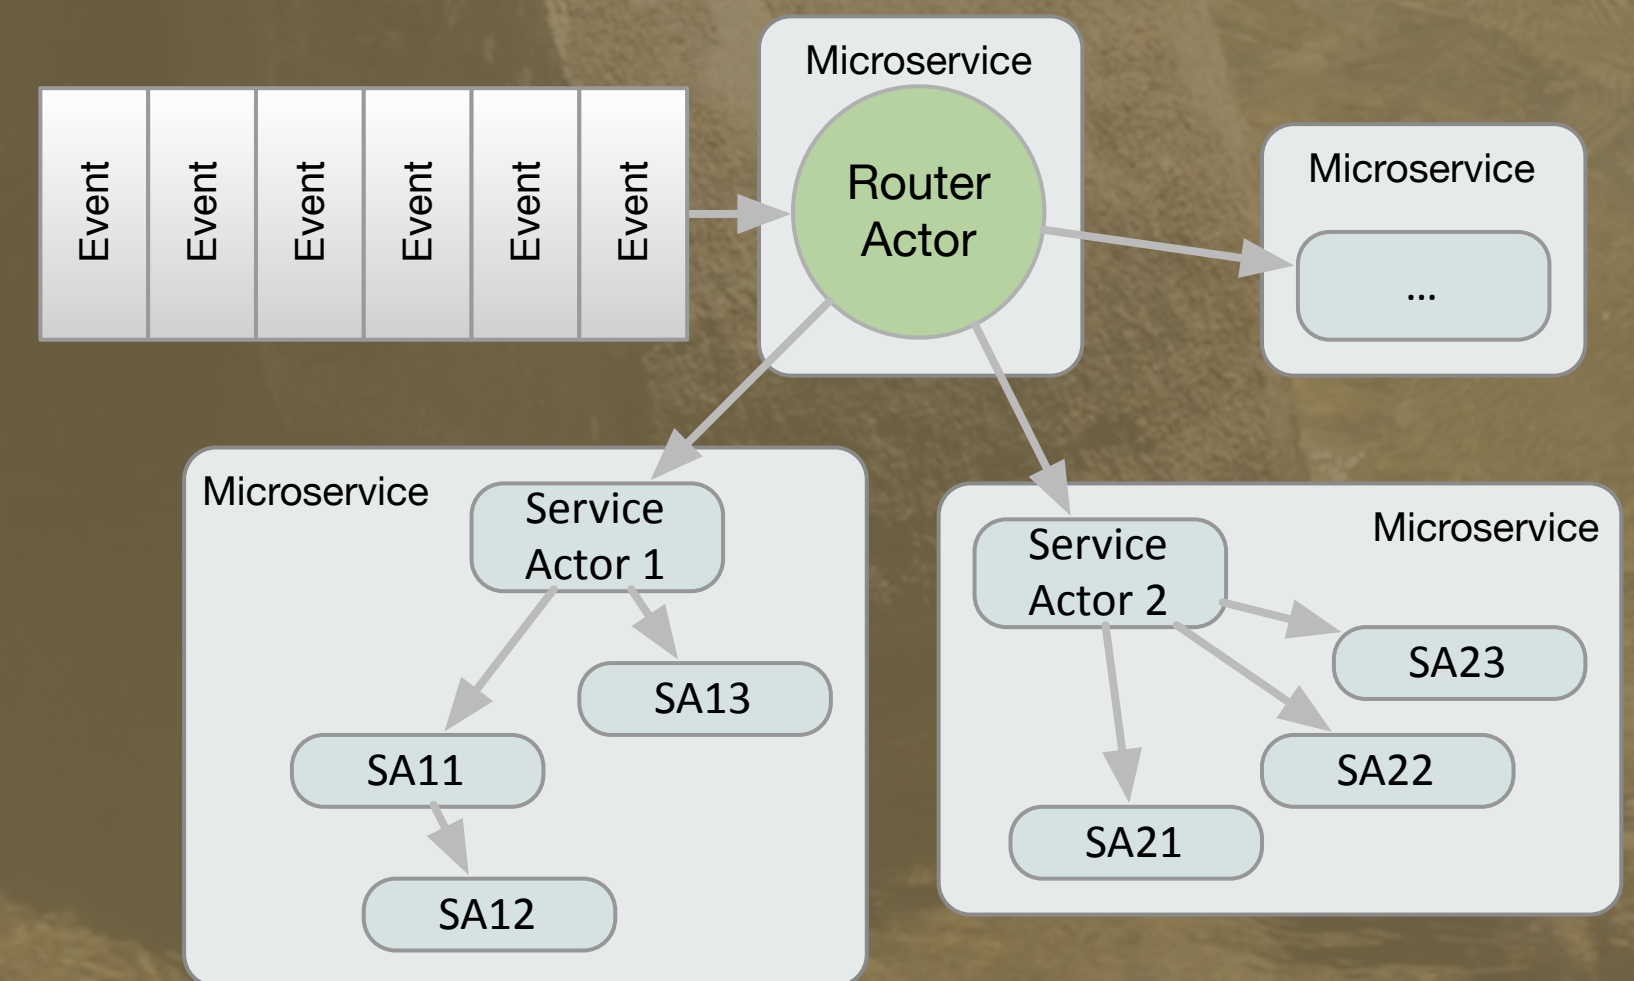
How is this...



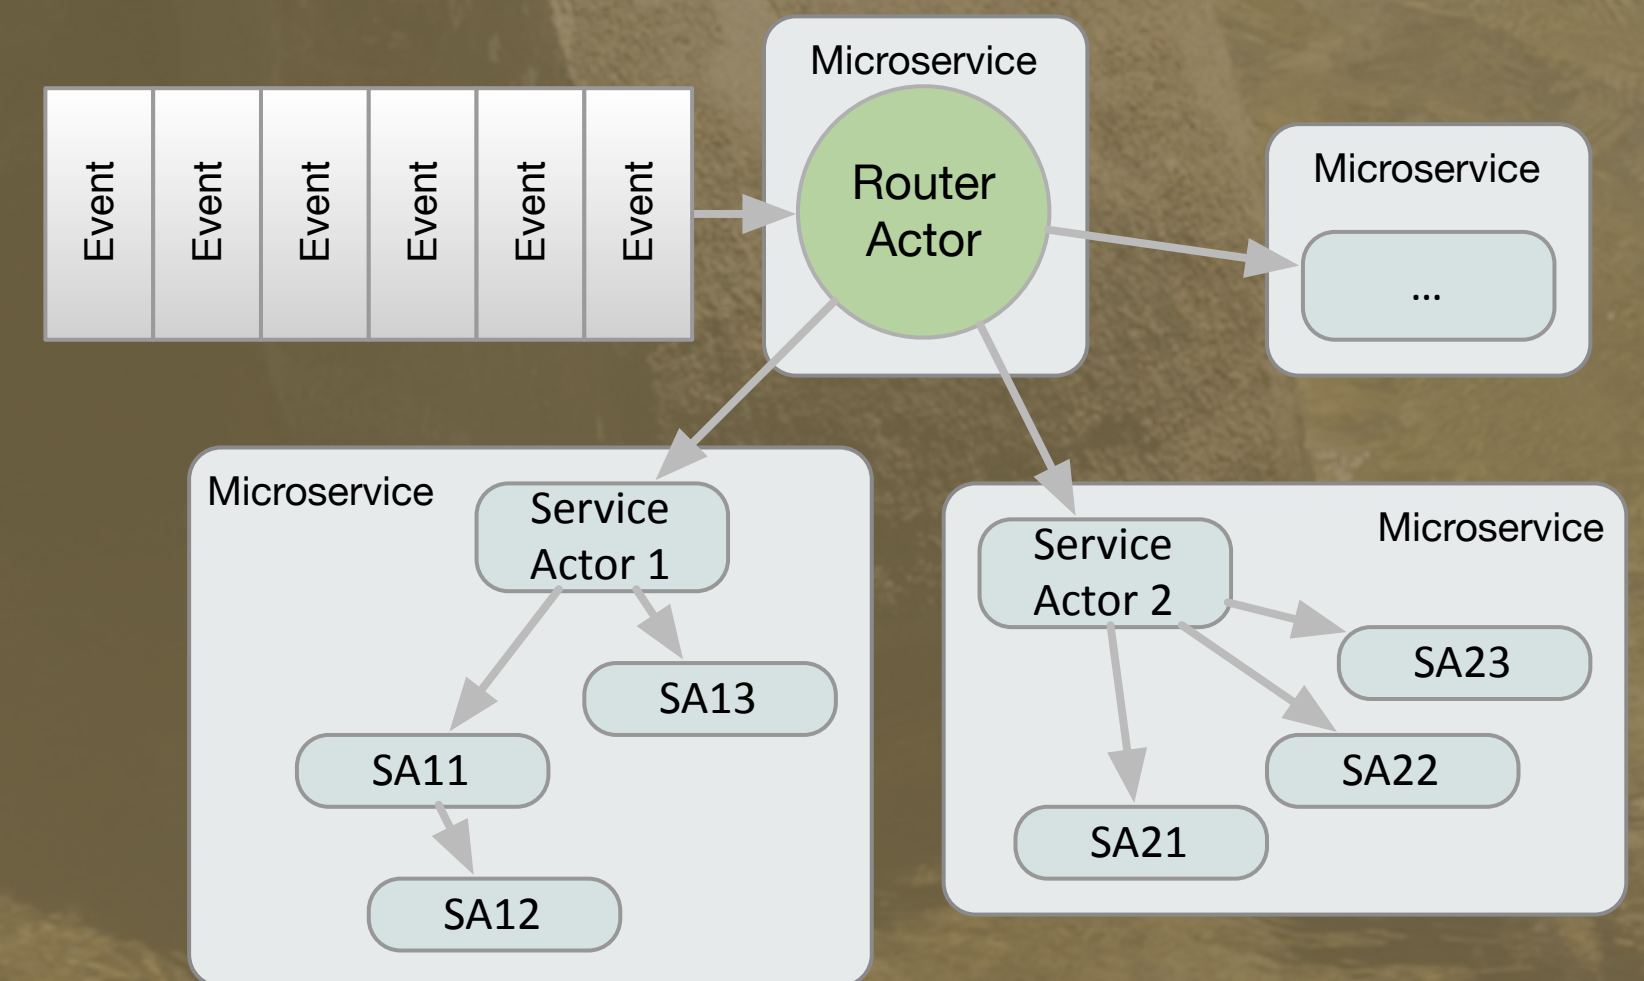
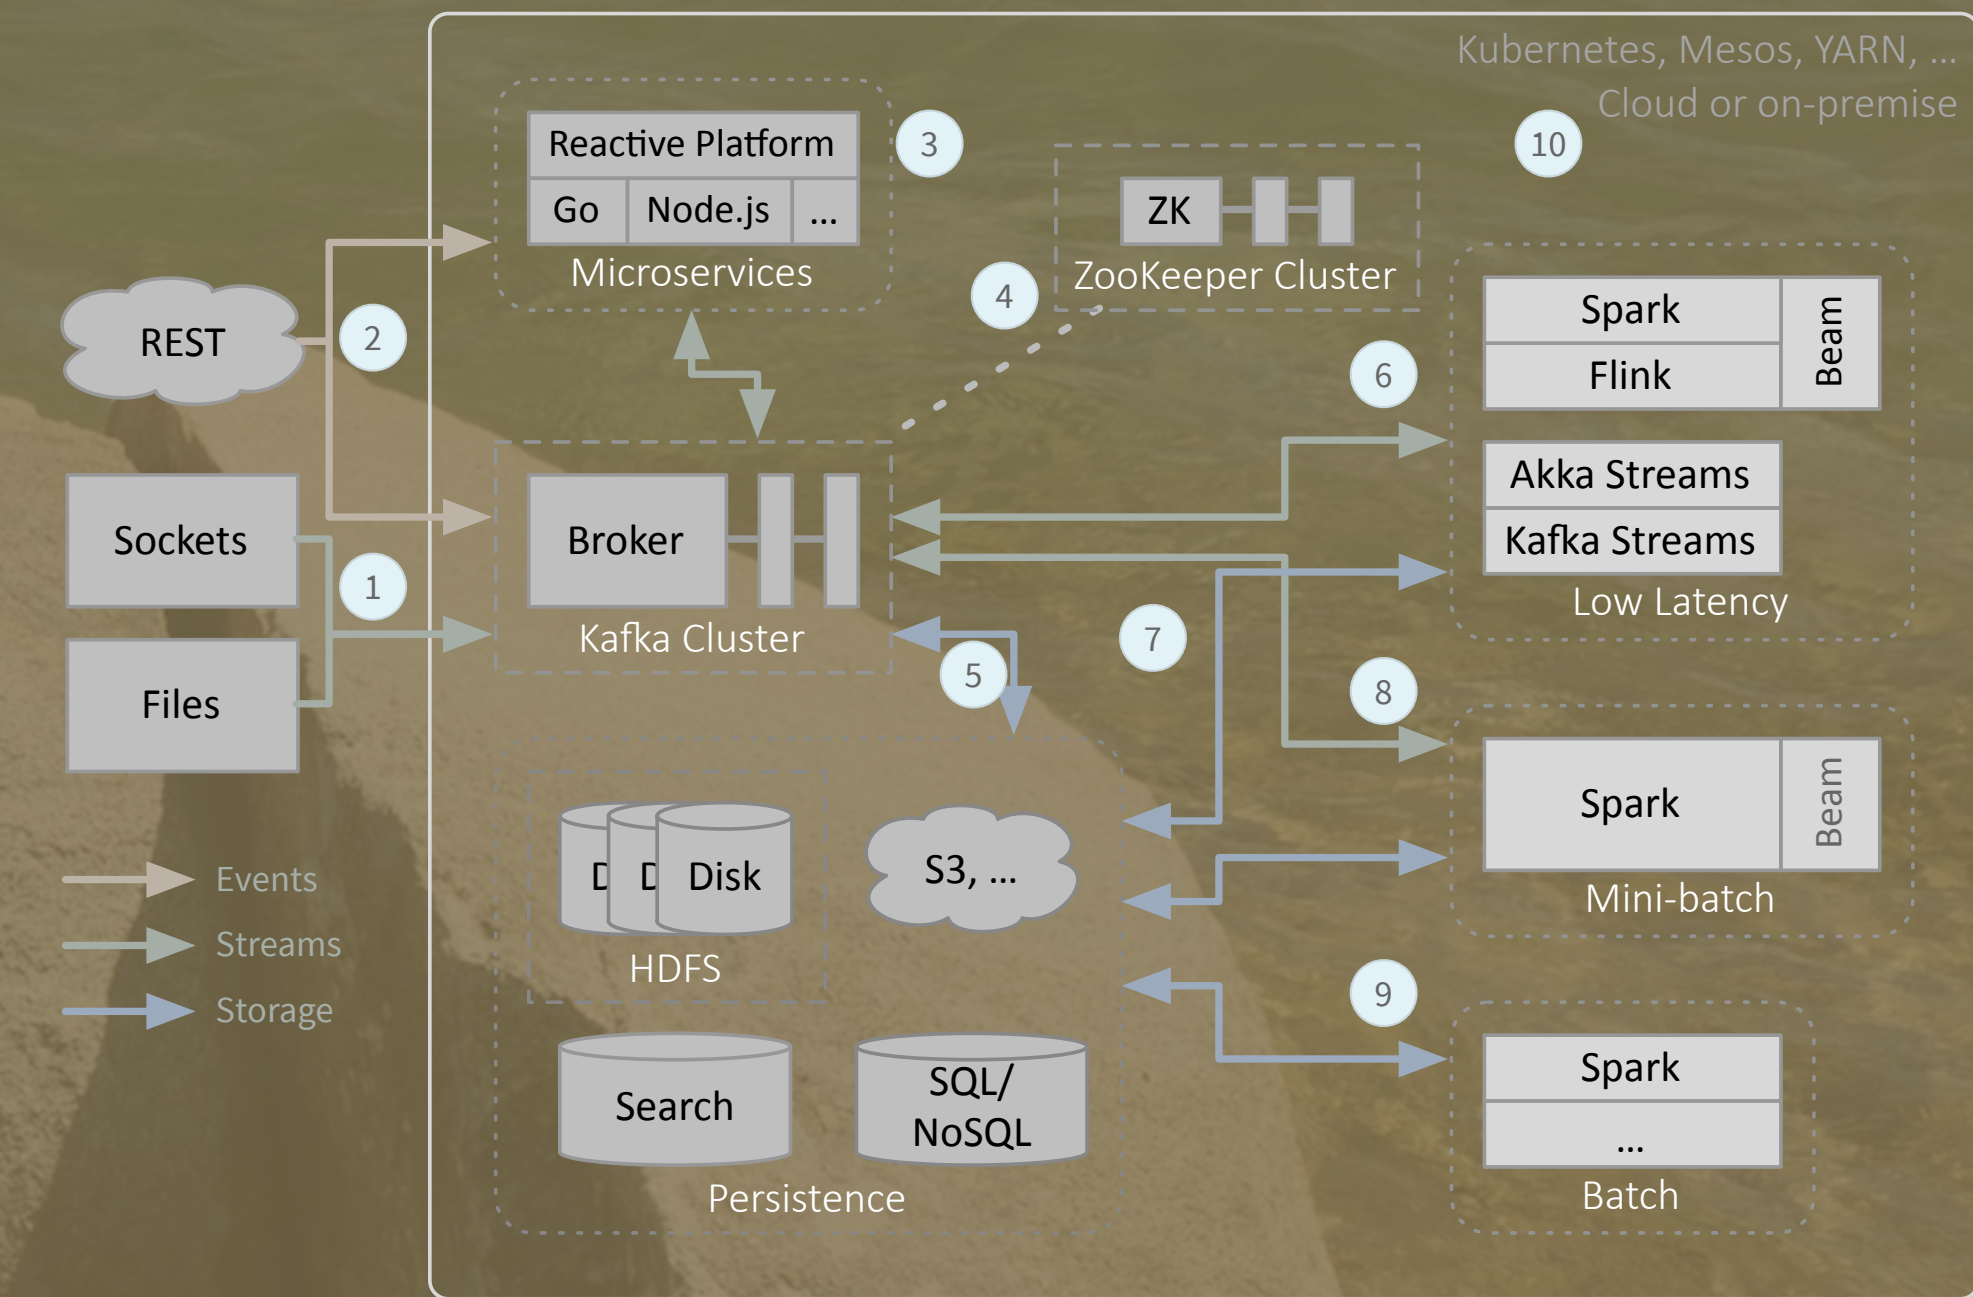
... like this?



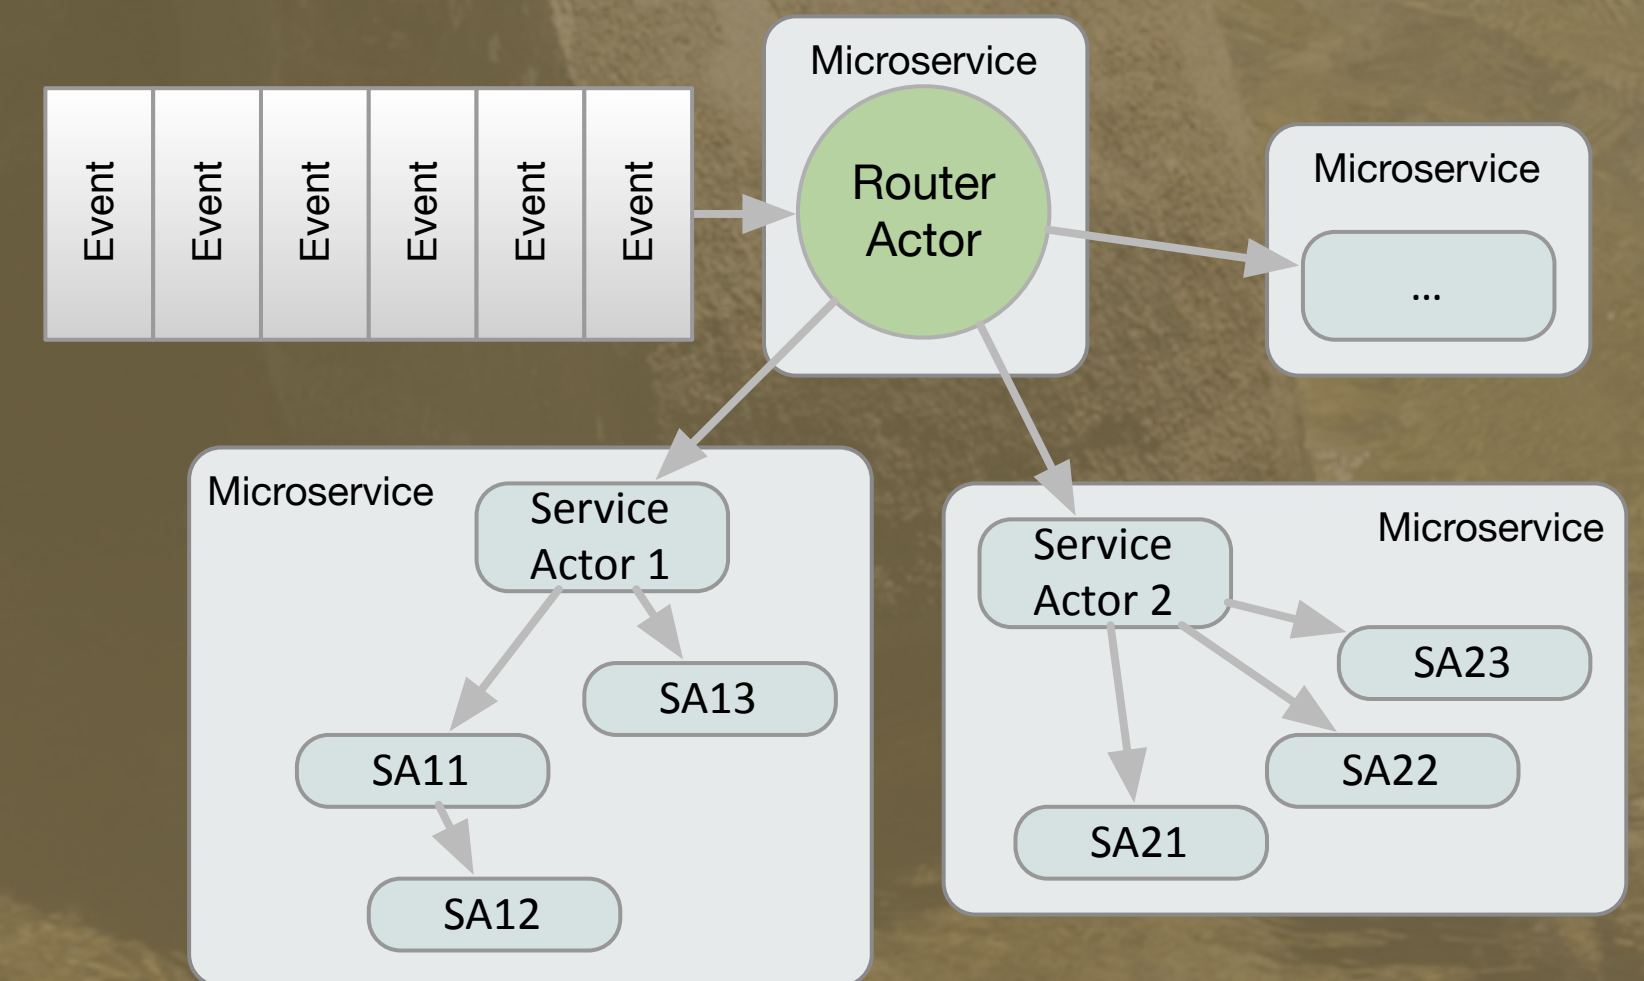
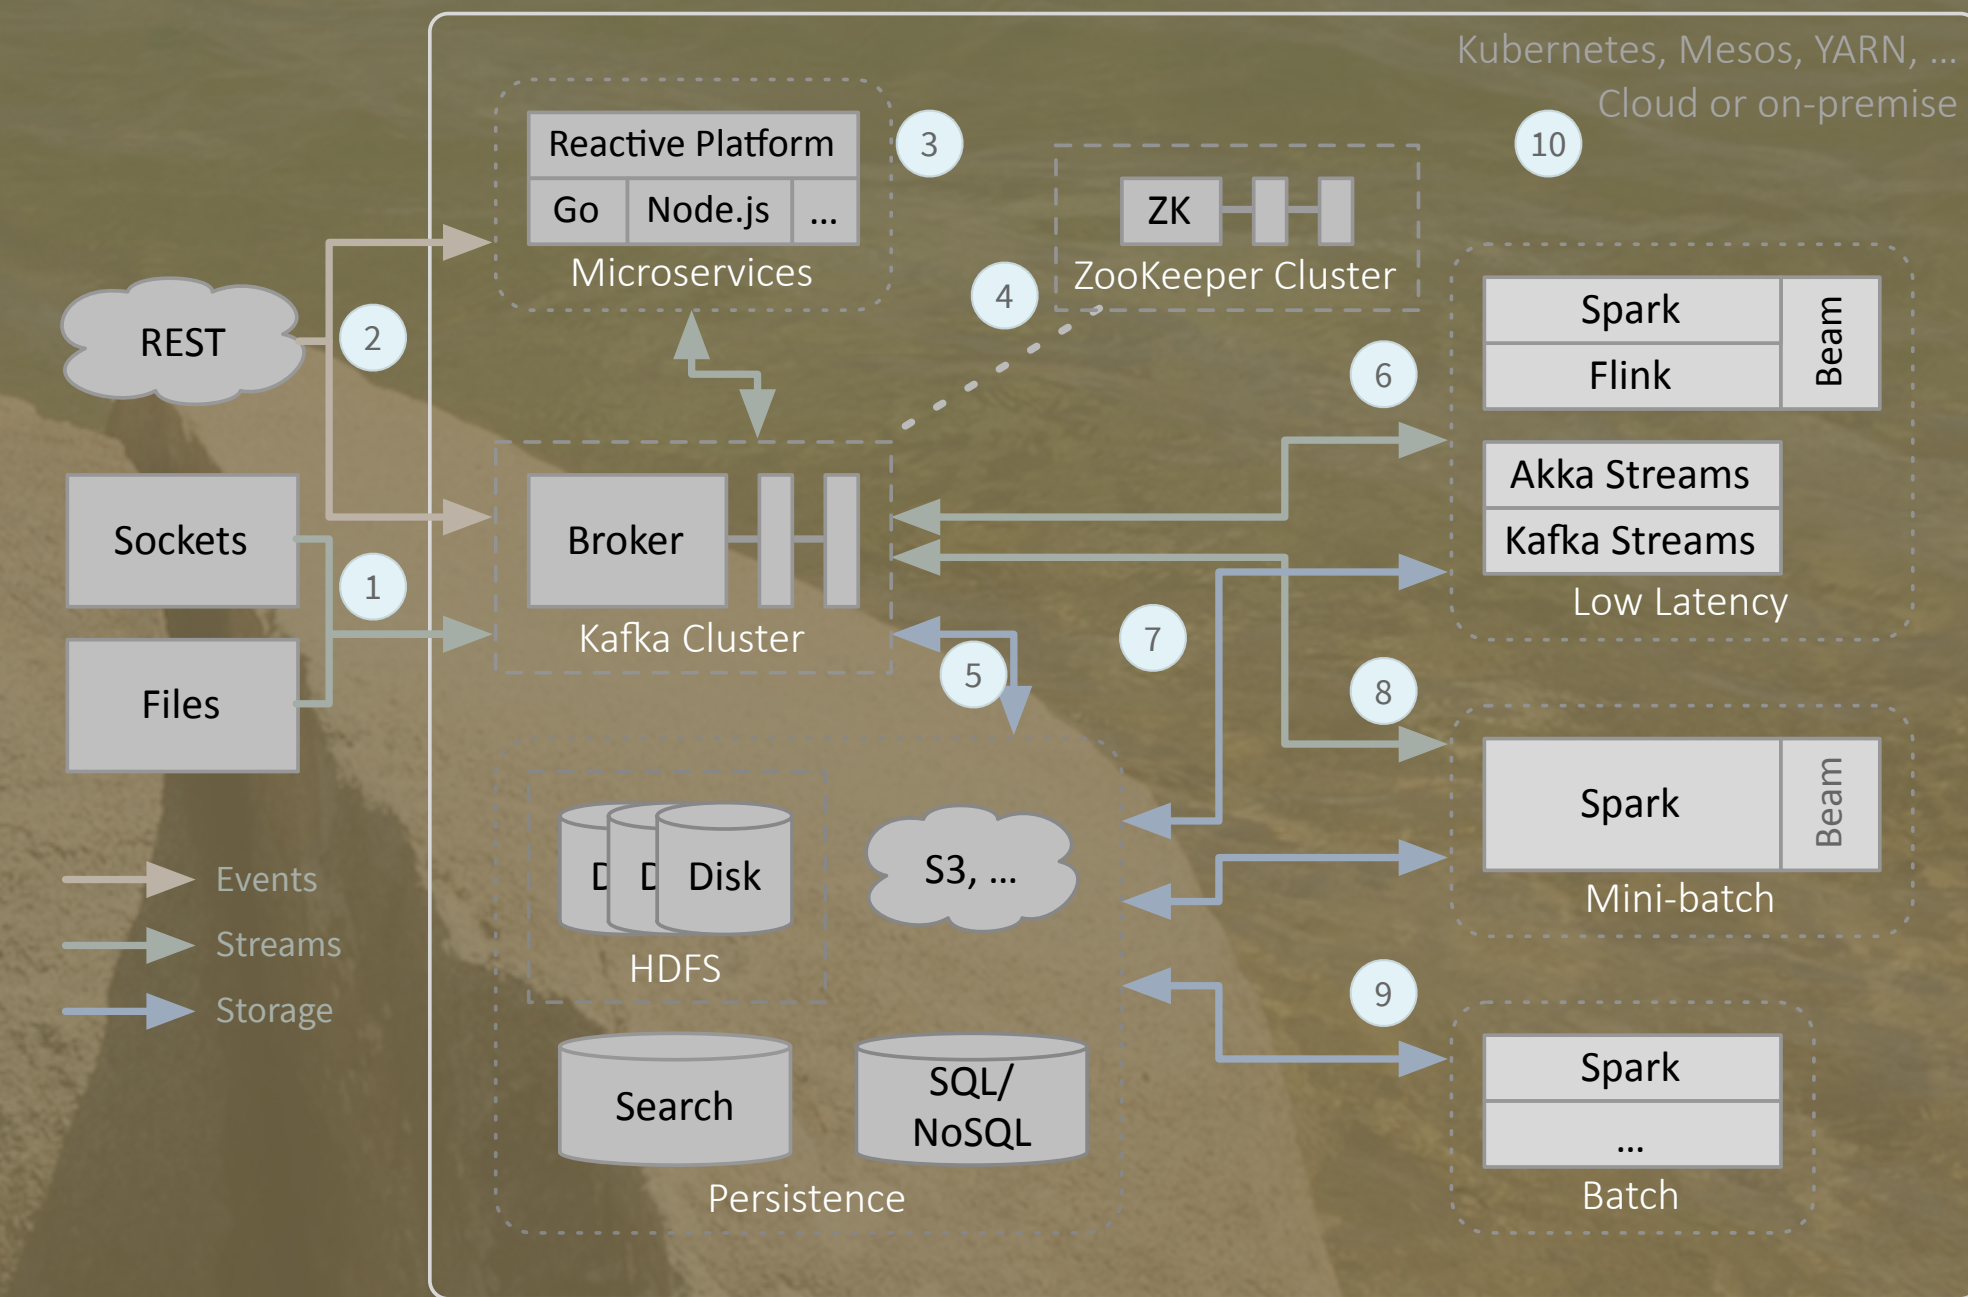
-



- A data app / microservice:
- A single responsibility.
- The input never ends.

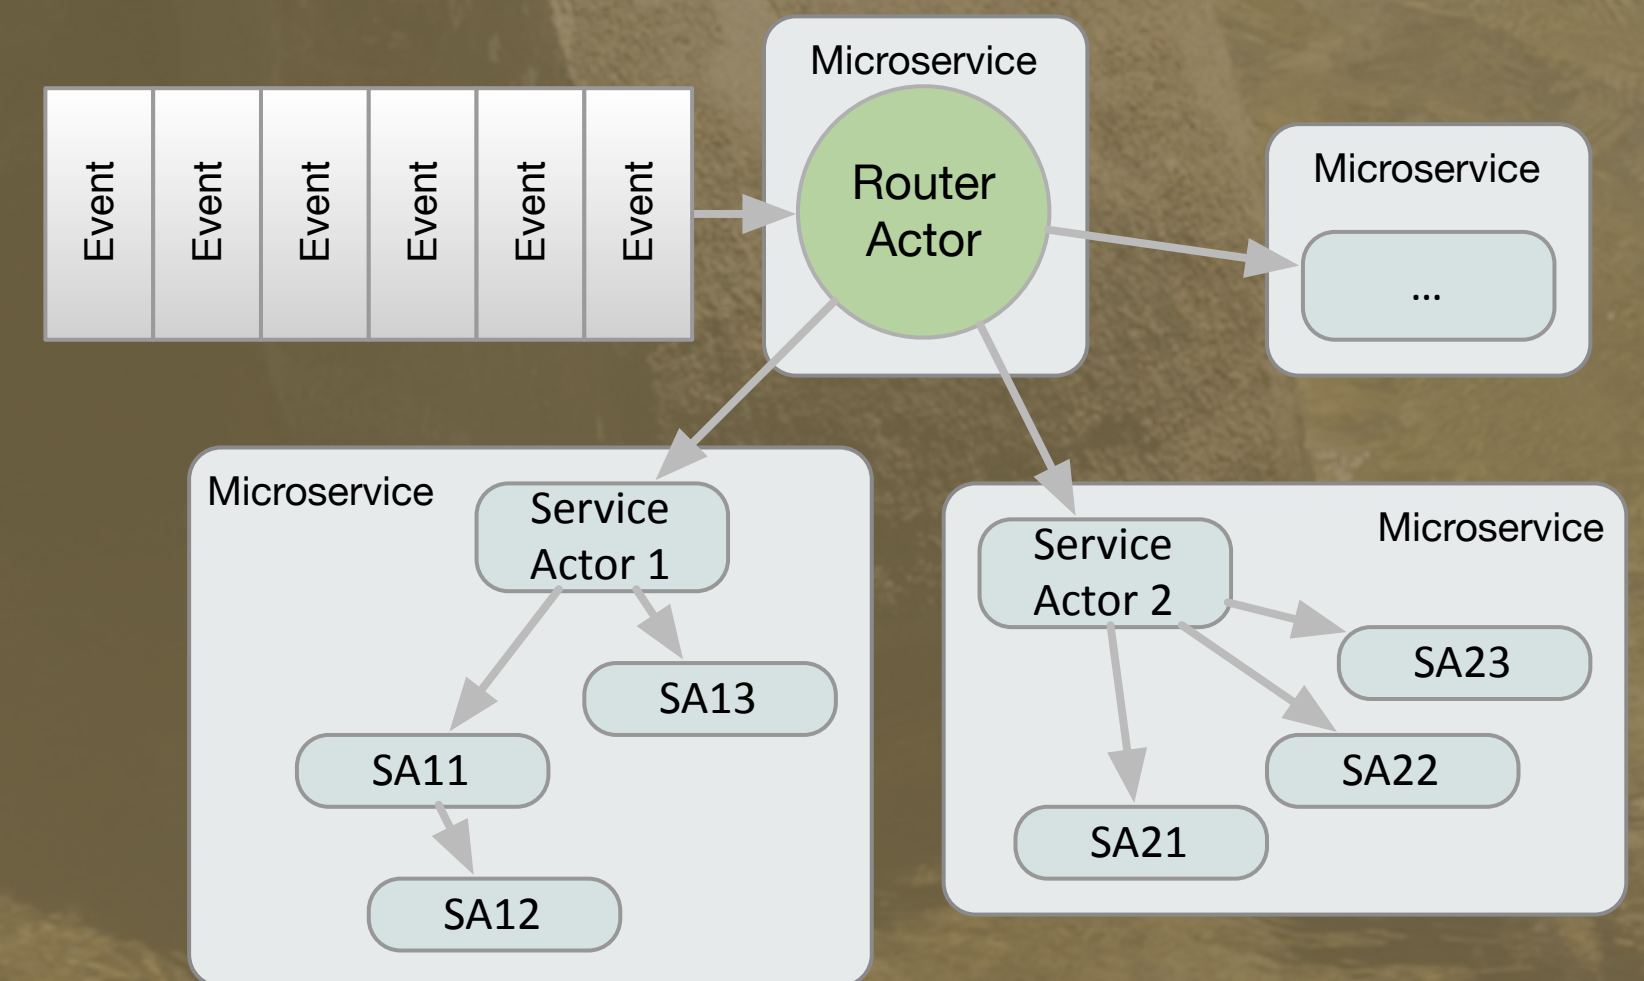
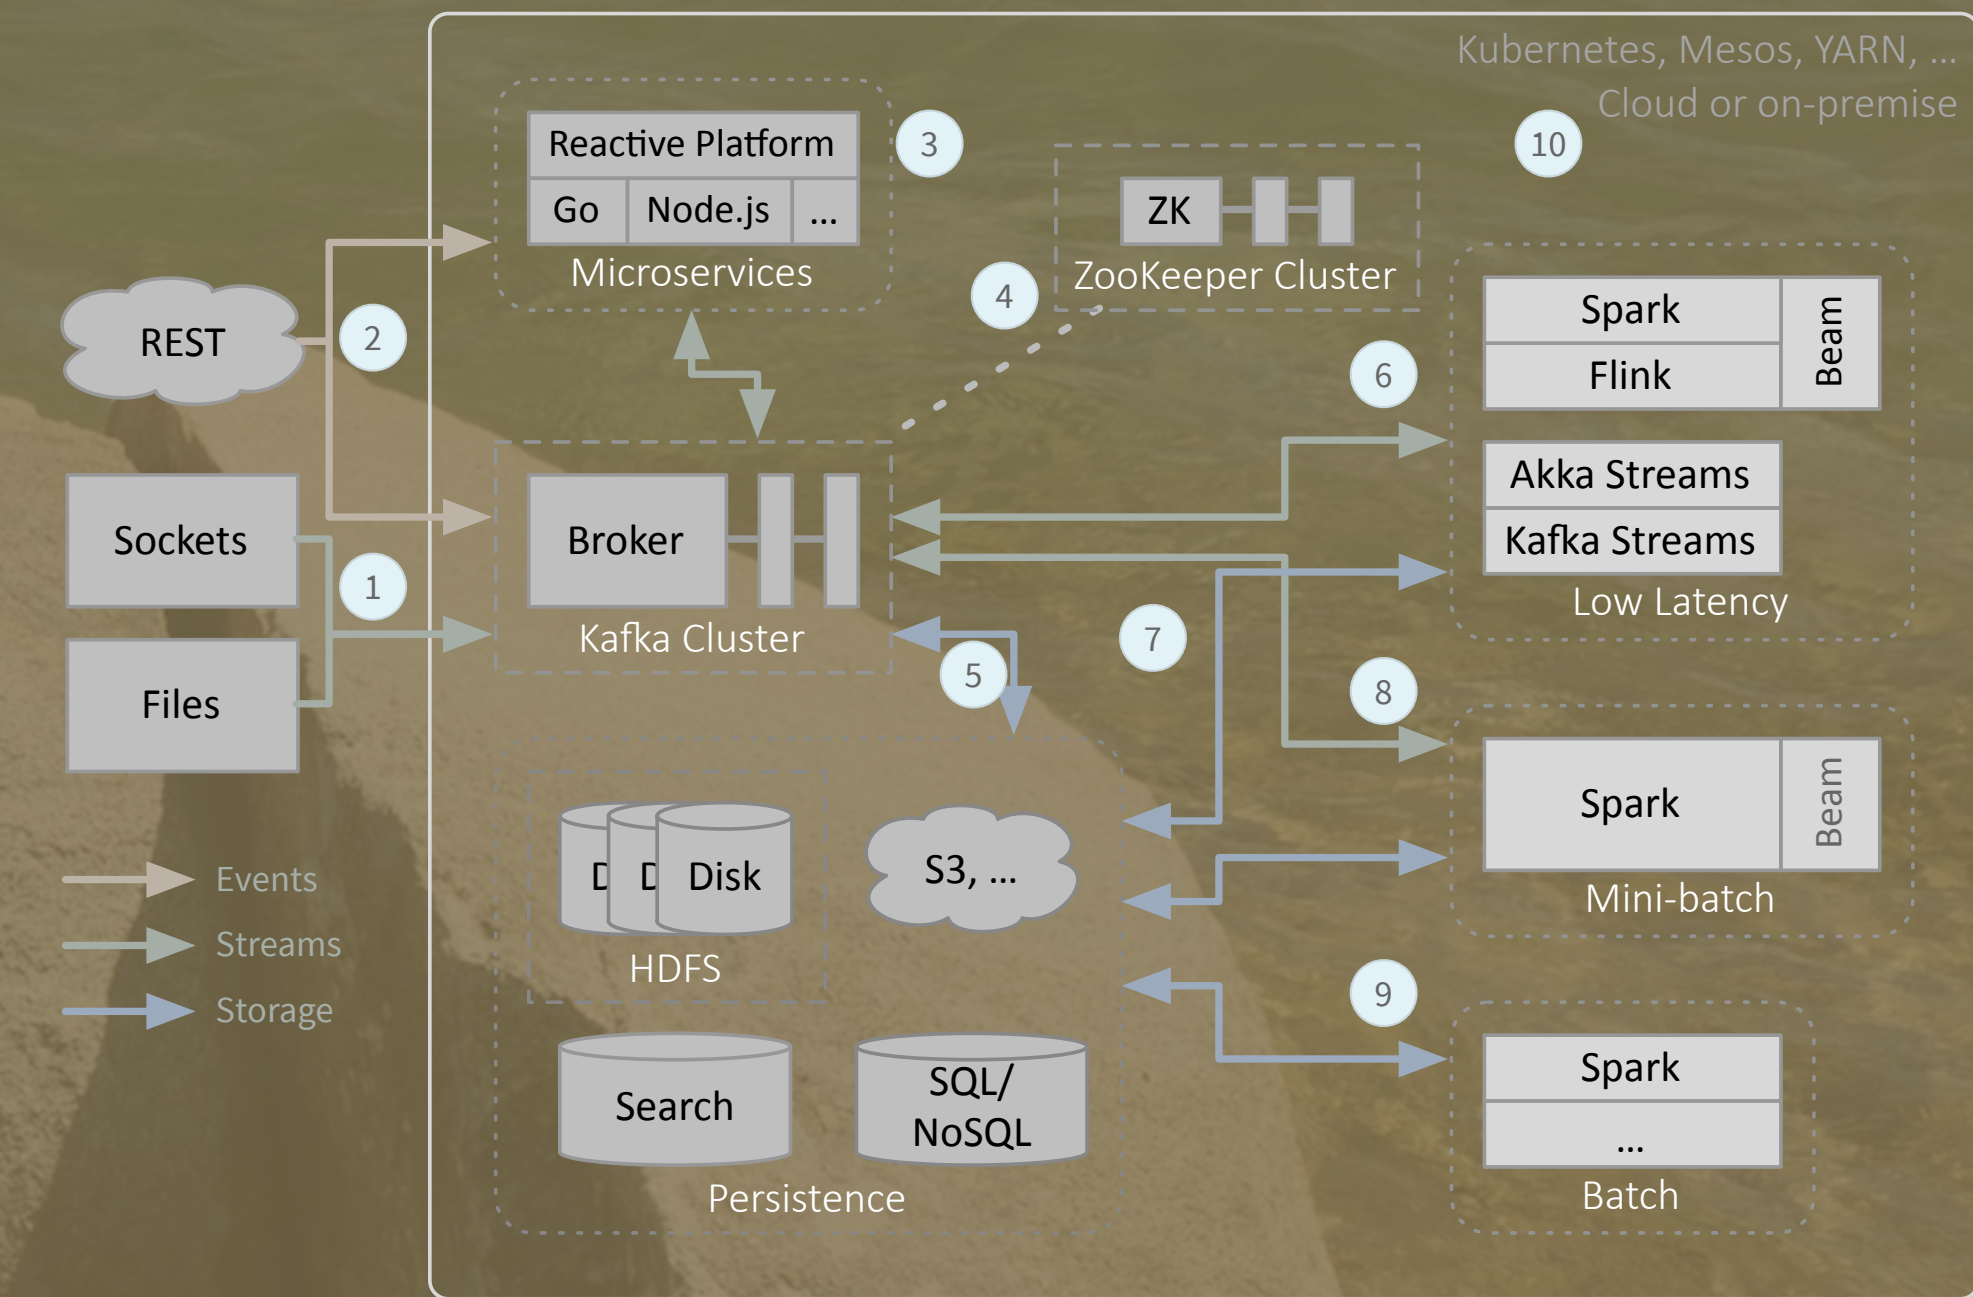


- A data app/microservice:
- A single responsibility.
- The input never ends.
- So, both must be available, responsive, resilient, & scalable. I.e., reactive




<http://www.reactivemanifesto.org/>

- Going the other way, “small” microservice architectures become data-centric, as the data grows.



The Recent Past



Some Overlap: Concerns, Architecture

The Present

Microservices
& *Fast* Data

Much More Overlap

The Future?

Why? Since streams process data incrementally, there is less need for large-scale tools like Spark, Flink

Unclear if this helps bridge the divide between data science and data engineering

... and using microservices for everything simplifies development, deployment, and operations

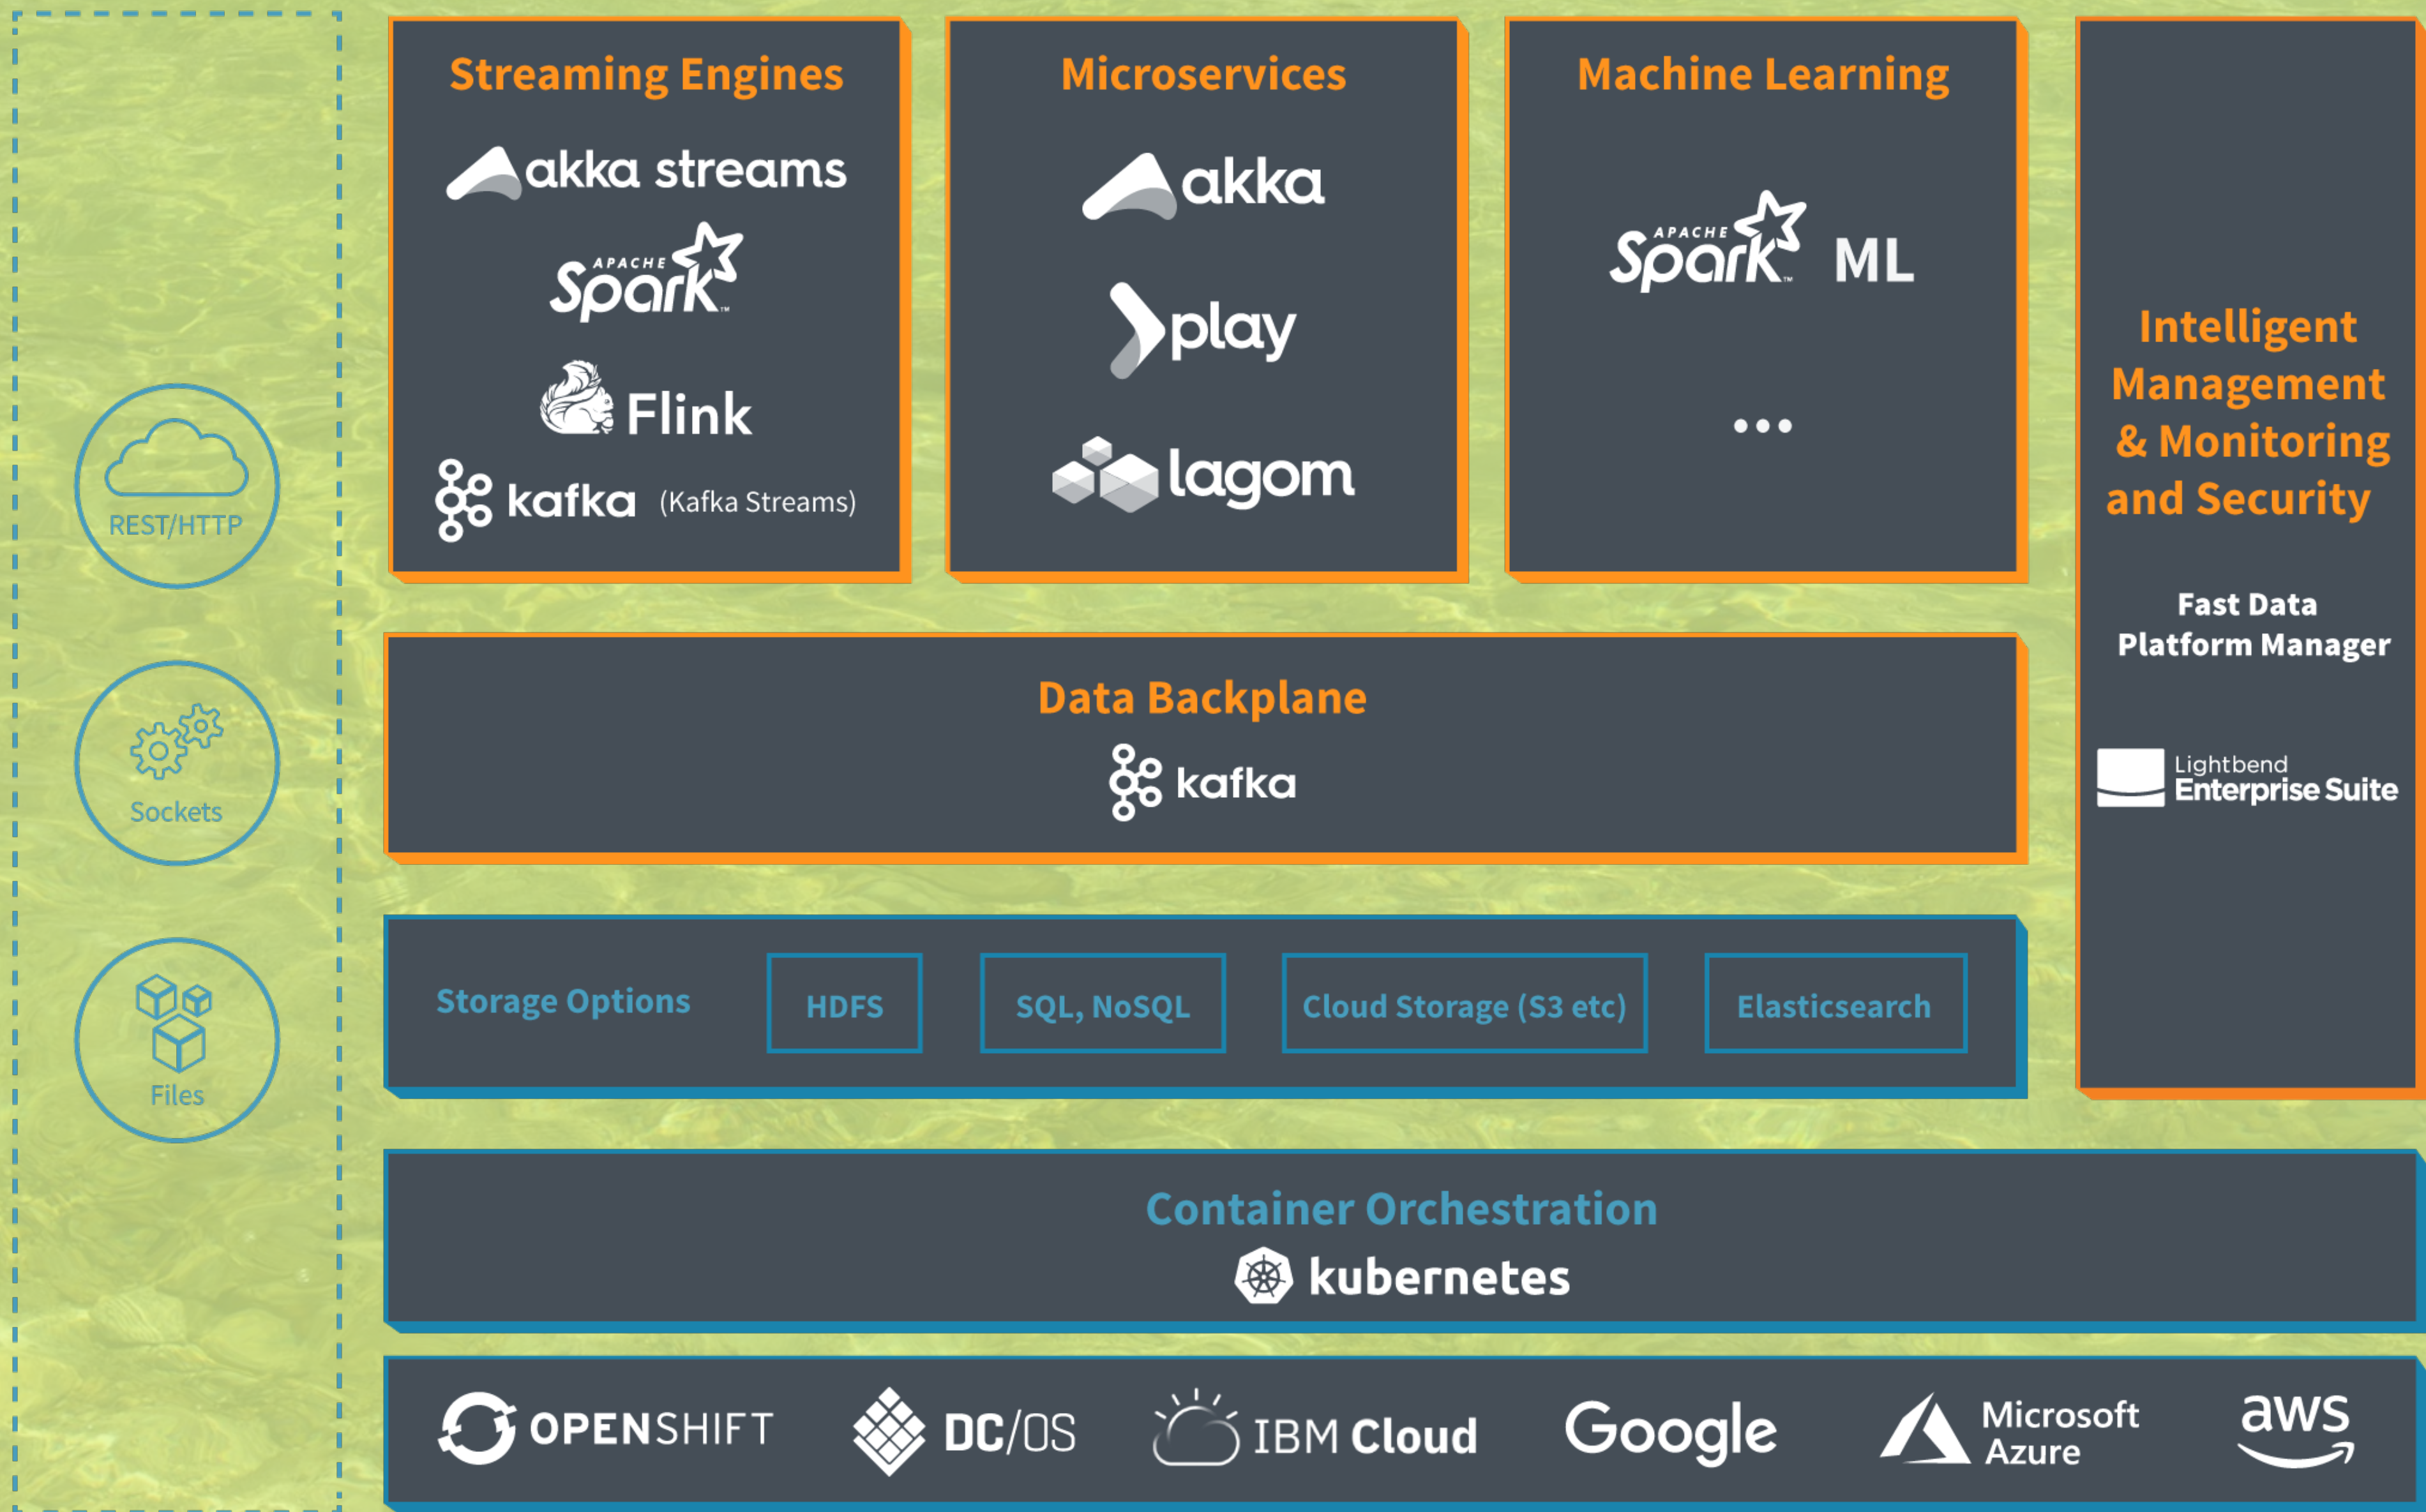
Microservices
for Fast Data

Much more microservice focused?

The background of the image is a close-up, high-resolution shot of water rippling over a rocky surface. The water is clear, revealing the texture and colors of the rocks beneath. The light reflects off the water's surface, creating a shimmering, iridescent effect with various shades of green, yellow, and brown. The overall texture is highly detailed and organic.

Lightbend Fast Data Platform

lightbend.com/fast-data-platform



lightbend.com/fast-data-platform

What we
discussed



Streaming Engines

akka streams

APACHE
Spark™

Flink

kafka (Kafka Streams)

Microservices

akka

play

lagom

Machine Learning

APACHE
Spark™ ML

...

Data Backplane

kafka

Storage Options

HDFS

SQL, NoSQL

Cloud Storage (S3 etc)

Elasticsearch

Container Orchestration

kubernetes

OPENSIFT

DC/OS

IBM Cloud

Google

Microsoft
Azure

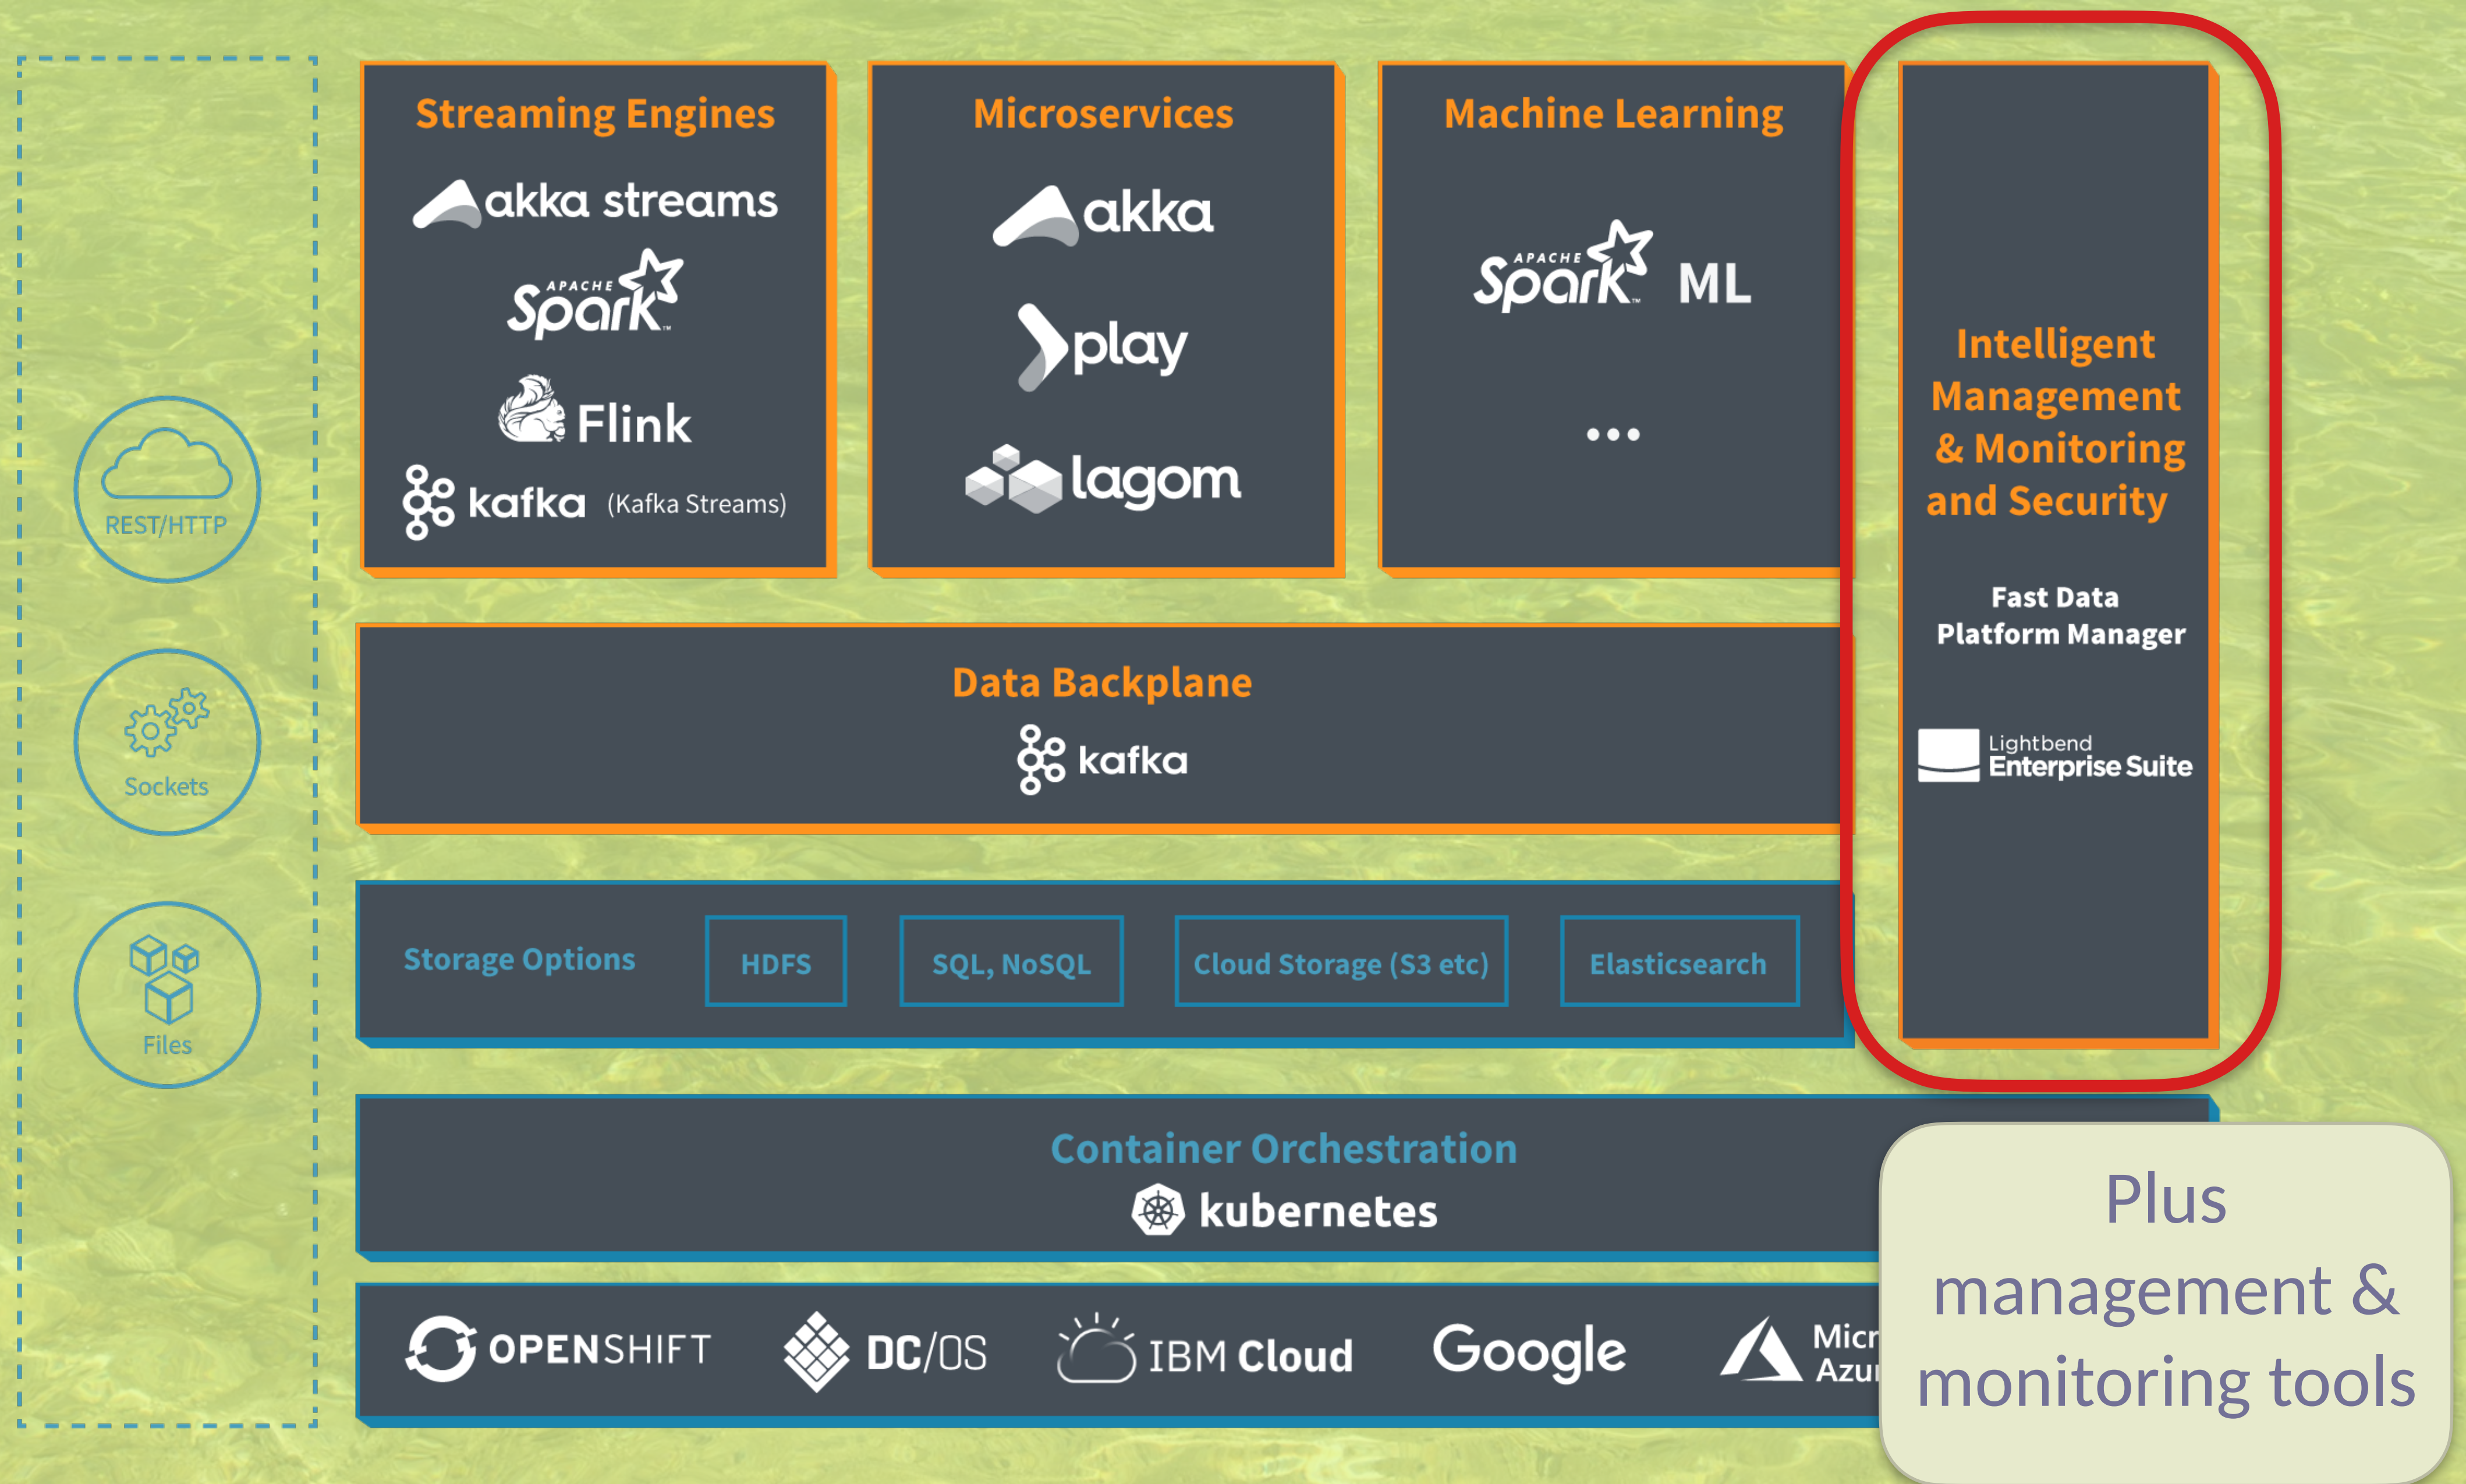
aws

**Intelligent
Management
& Monitoring
and Security**

**Fast Data
Platform Manager**

Lightbend
Enterprise Suite

lightbend.com/fast-data-platform



lightbend.com/fast-data-platform

Questions?

Dean Wampler, Ph.D.
dean@deanwampler.com
@deanwampler
polyglotprogramming.com/talks

