

# Leveraging Stateful Functions to Power the Next Generation of Event-Driven Applications

---

Seth Wiesman

@sjwiesman on most platforms

# About Ververica



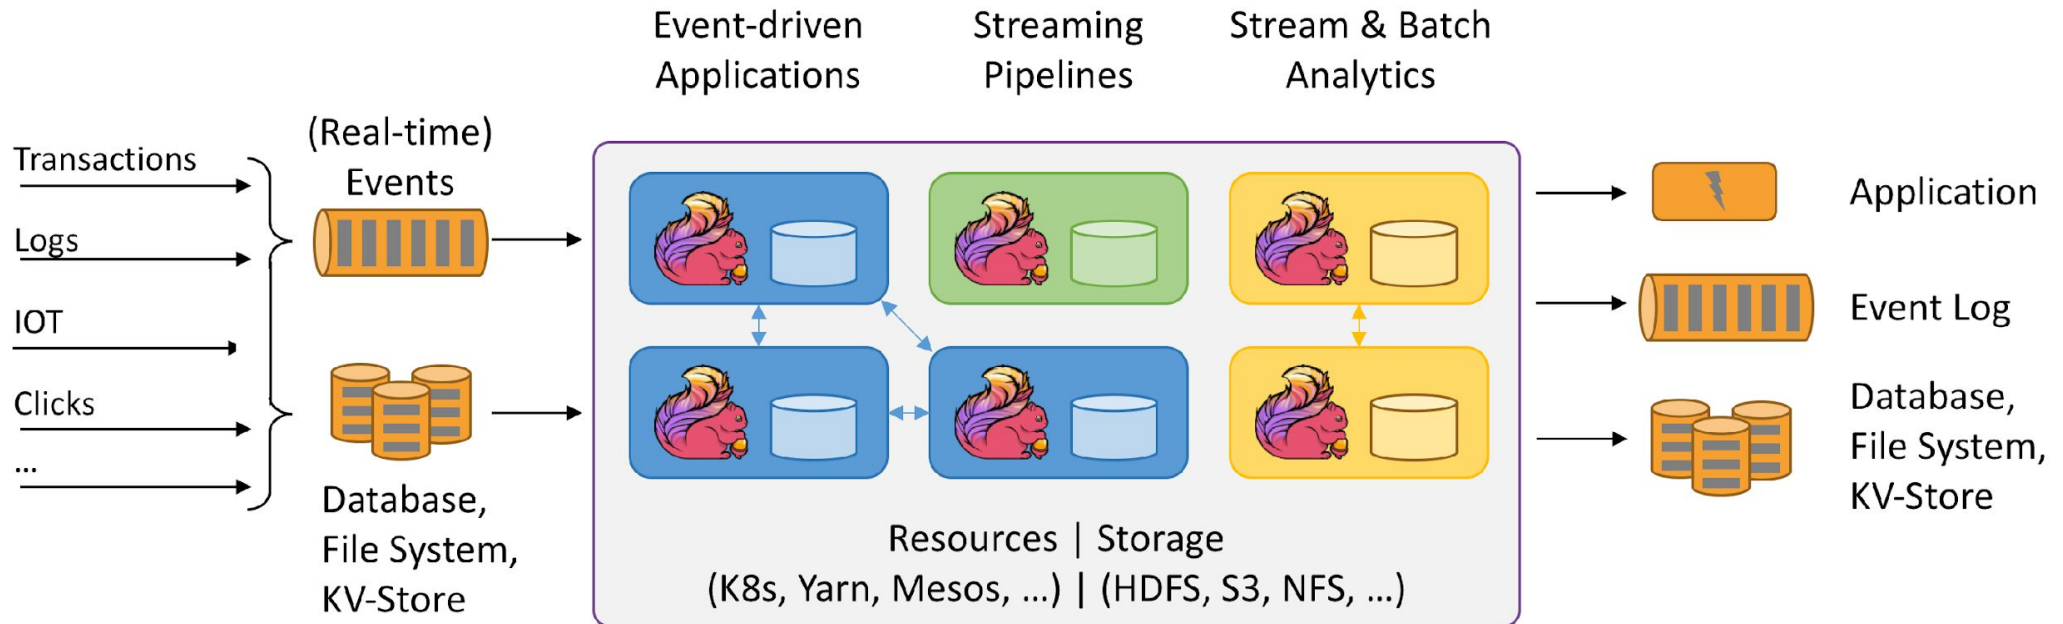
Original creators of  
Apache Flink®



Enterprise  
Stream Processing



# Stateful Computations over Data Streams



# Some Apache Flink Users



Sources: [Powered by Flink](#), [Speakers – Flink Forward San Francisco 2019](#), [Speakers – Flink Forward Europe 2019](#)



# Apache Flink at



## The "Singles Day" (11/11/2018)

machines

queries

throughput

latency

state size



10K



10K



1.7B

events / sec



Sub-Second



100TB



Let's look at building Applications

# Building an Application Today



# Building an Application Today




The big trend: Serverless





# Functions as a Service


an event-driven function



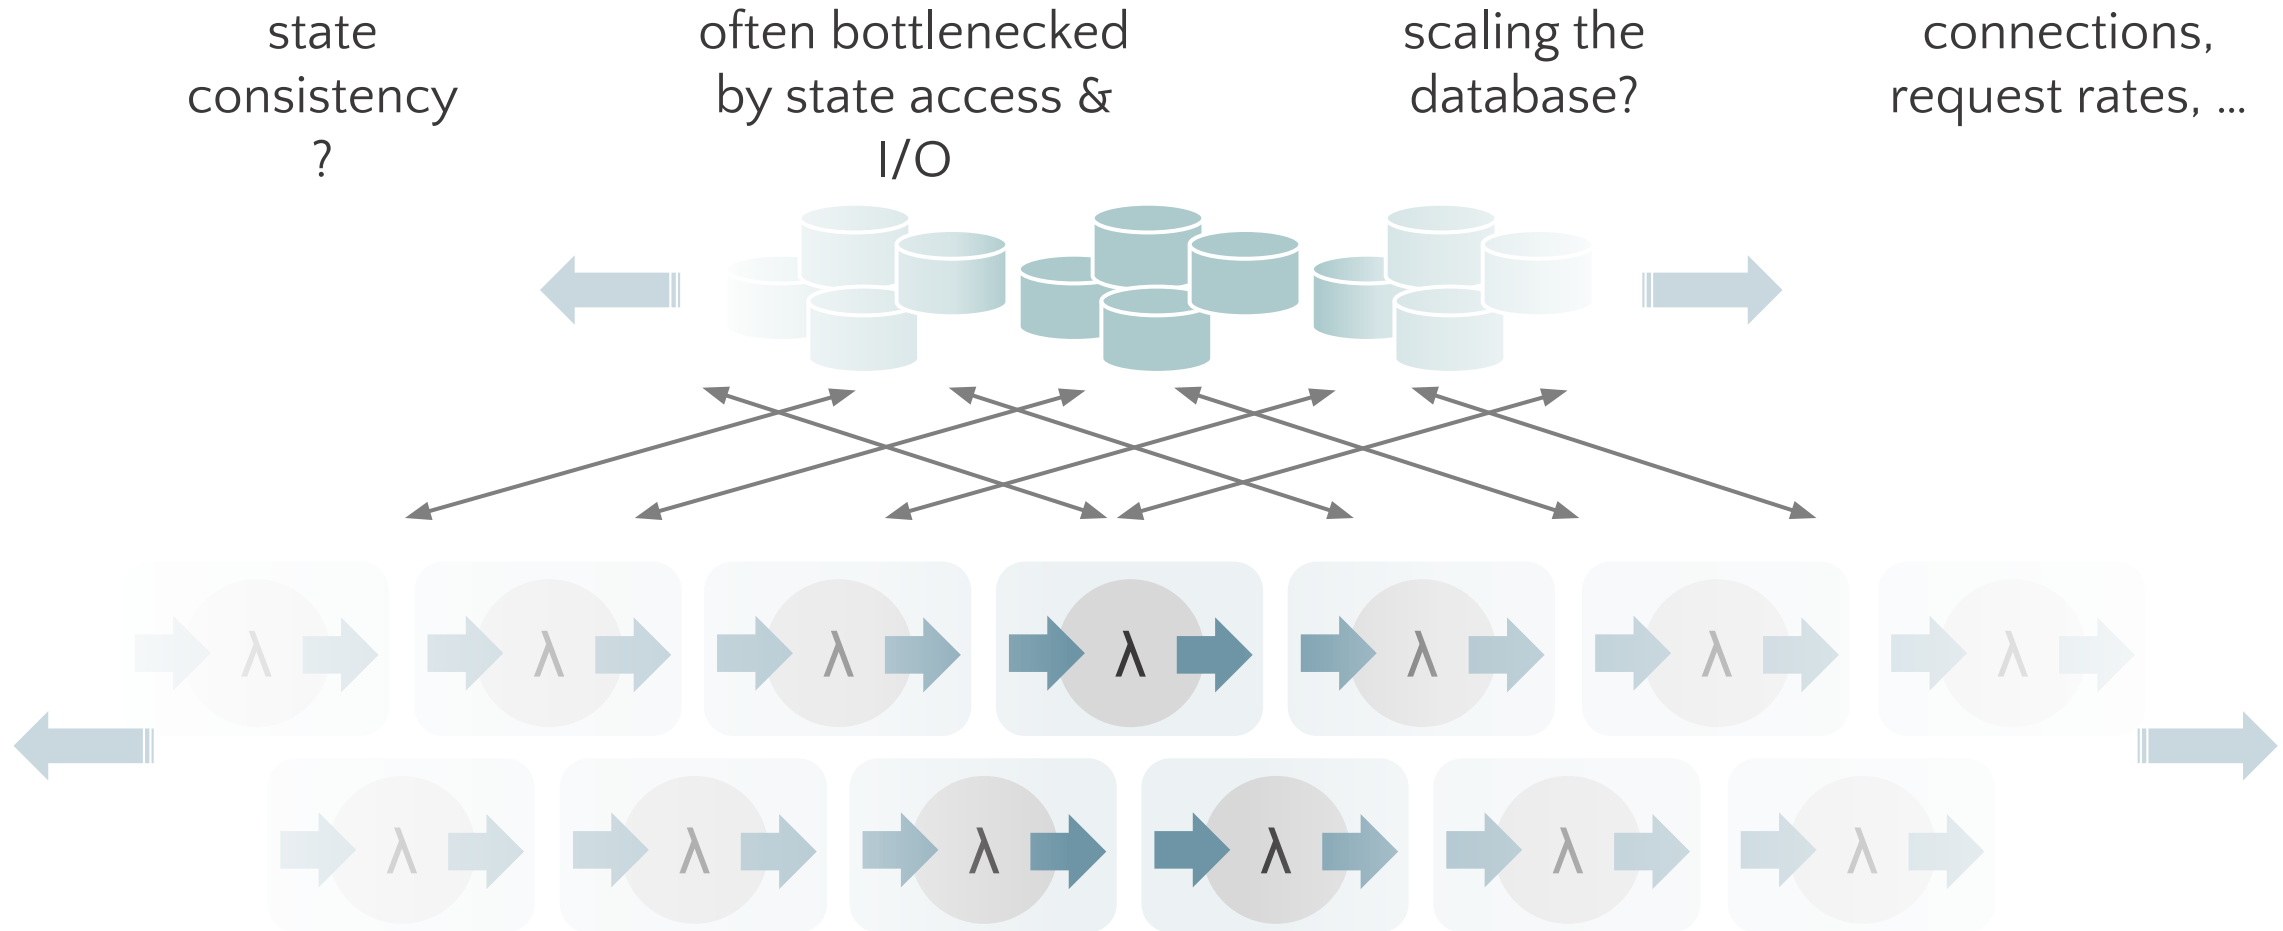
# Functions as a Service

elastically  
scalable

“lightweight resource footprint”



# Functions as a Service – Handling State in Applications



# Functions as a Service – Handling State in Applications

Following ▾

Hardest part of building your services architecture is still your data and stateful services. [#serverless](#) does not solve that for [#btw](#)

5:01 AM - 9 Nov 2018

Following ▾

Storage is the single hardest problem in our domain. Storage related tradeoffs are sometimes the hardest tradeoffs to tackle. Storage decisions often impact every other design decision. I don't know why we are acting like it ain't so.

12:57 AM - 12 Jan 2019

Following ▾

Nobody talking about data consistency issues in stateful microservices and I'm angry about it.

4:51 PM - 13 Mar 2019

Follow ▾

OMG yes! So much energy being poured into orchestrating stateless applications. That isn't *\*totally\** trivial but it's pretty damn close relative to state and storage. And application devs too often pretend selecting a RDBMS means they don't have to worry about state consistency

**Jaana B. Dogan** @rakyll  
Storage is the single hardest problem in our domain. Storage related tradeoffs are sometimes the hardest tradeoffs to tackle. Storage decisions often impact every other design decision. I don't know why we are acting like it ain't so.


4:12 PM - 12 Jan 2019



# Composition of Functions

Not straightforward to build more complex applications

Lack of messaging / composition primitives



workflows of functions as a workaround, but not a general solution




event-driven

composable

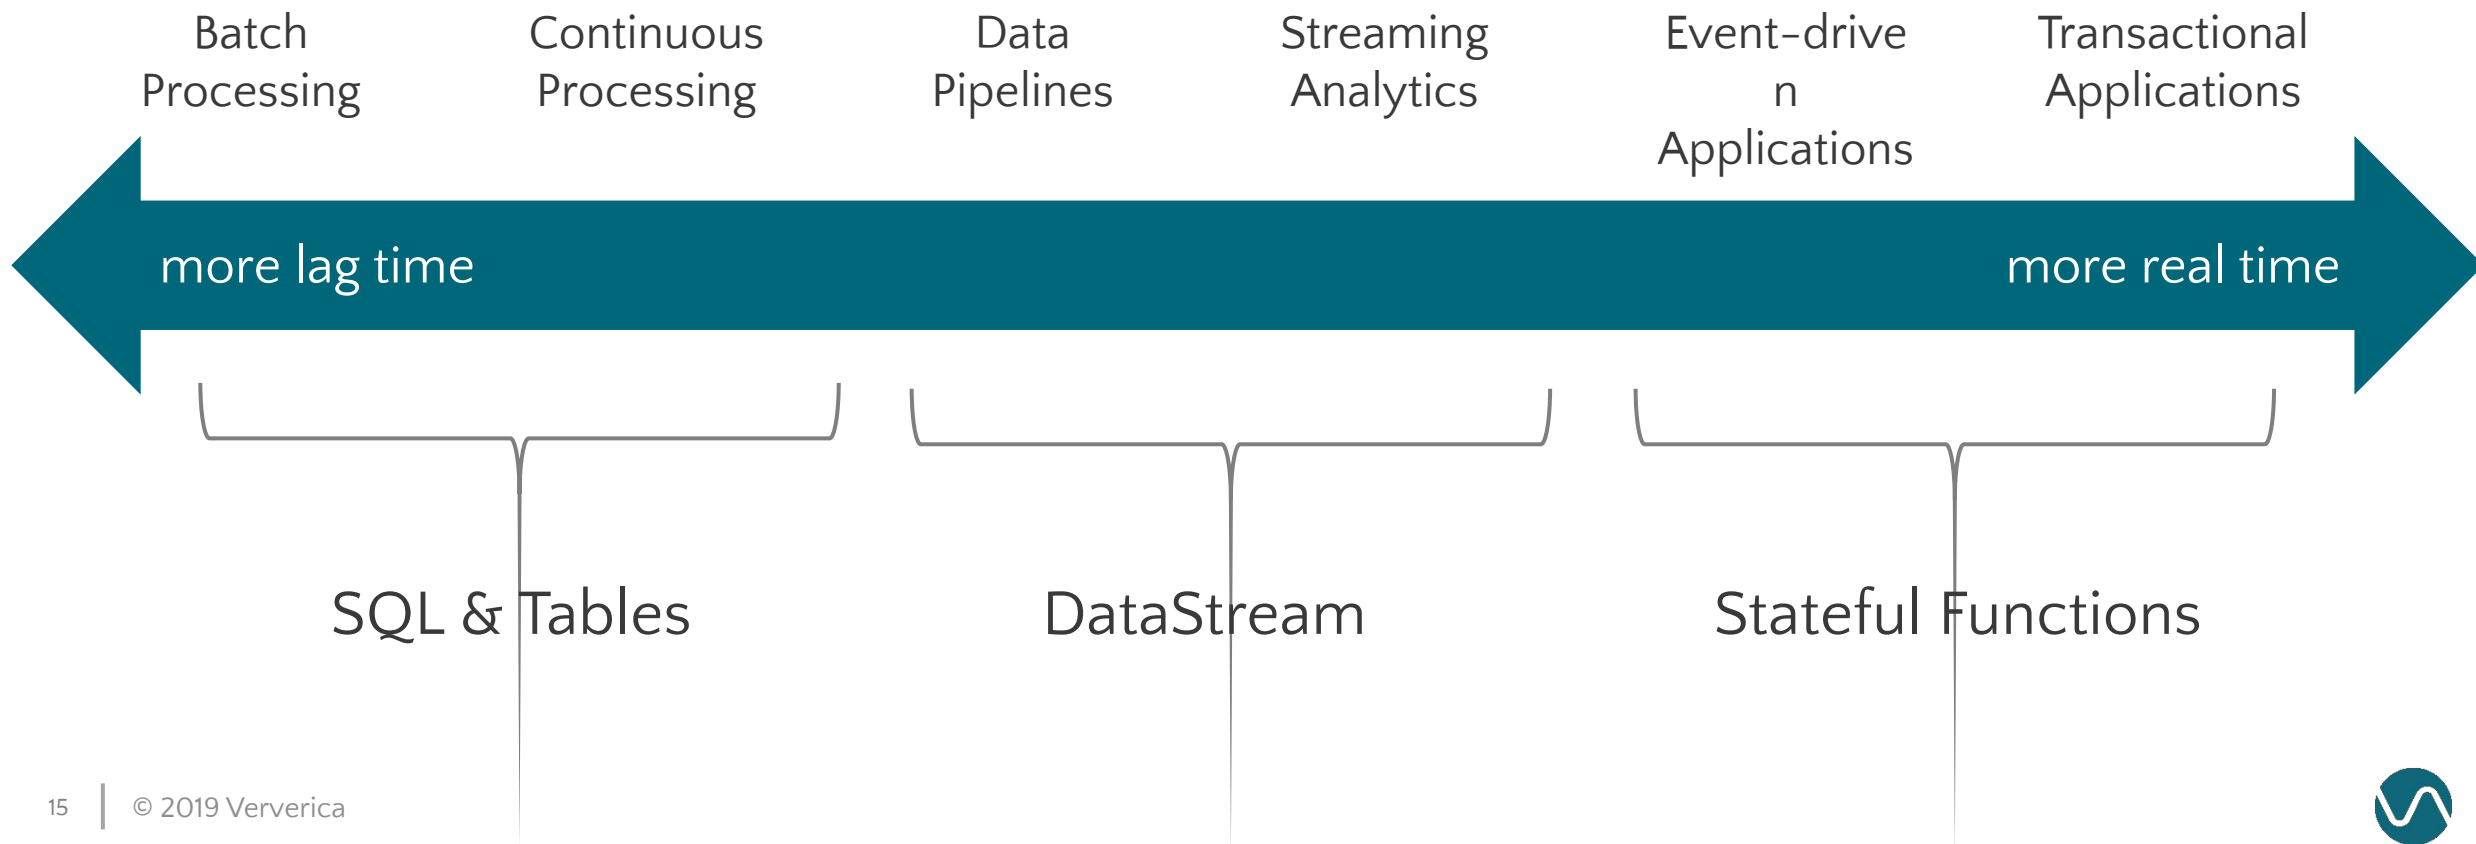
state management

...that sound like...

Stream  
Processing



# The Spectrum of Streaming Data Use Cases



# Disclaimer

Stateful Functions is currently a standalone project

<https://statefun.io/>

<https://github.com/ververica/stateful-functions>


The Apache Flink community has voted to adopt

The project is still new and dynamic.  
A good time to get involved to get traction ;-)






# Stateful Functions



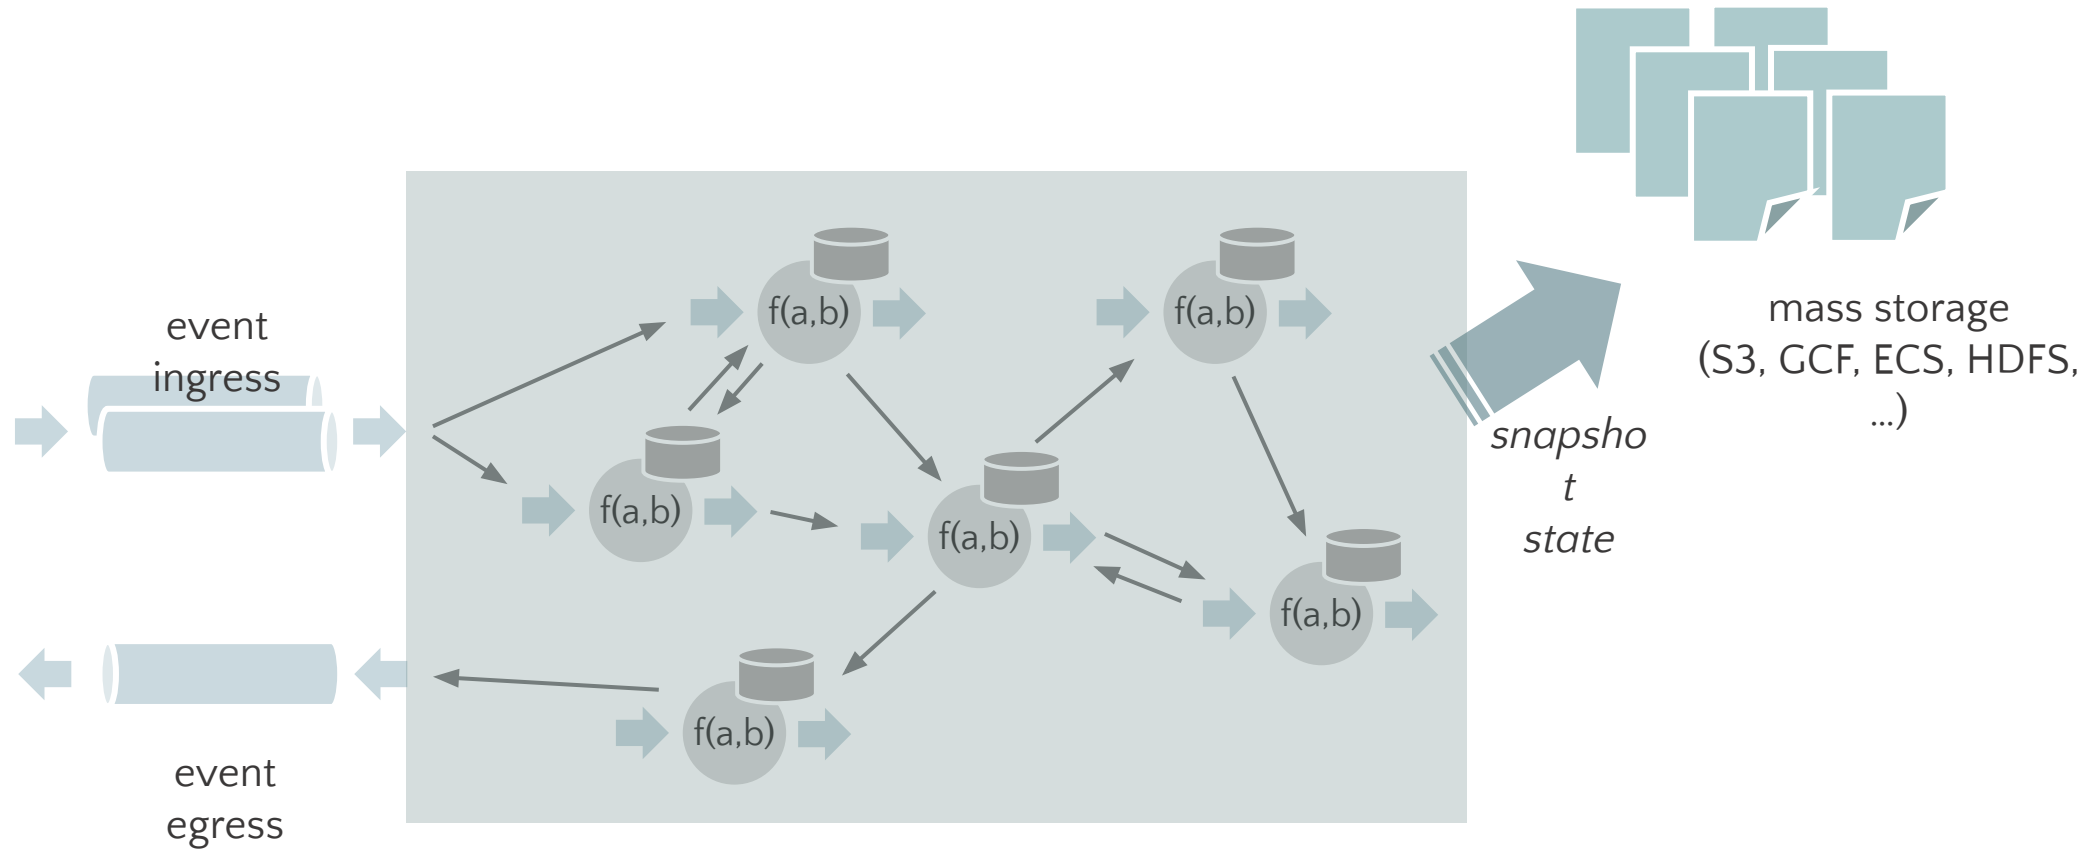
# Stateful Functions



Event ingresses supply events that trigger functions




# Stateful Functions



Multiple functions send event to each other  
Arbitrary addressing, no restriction to DAG



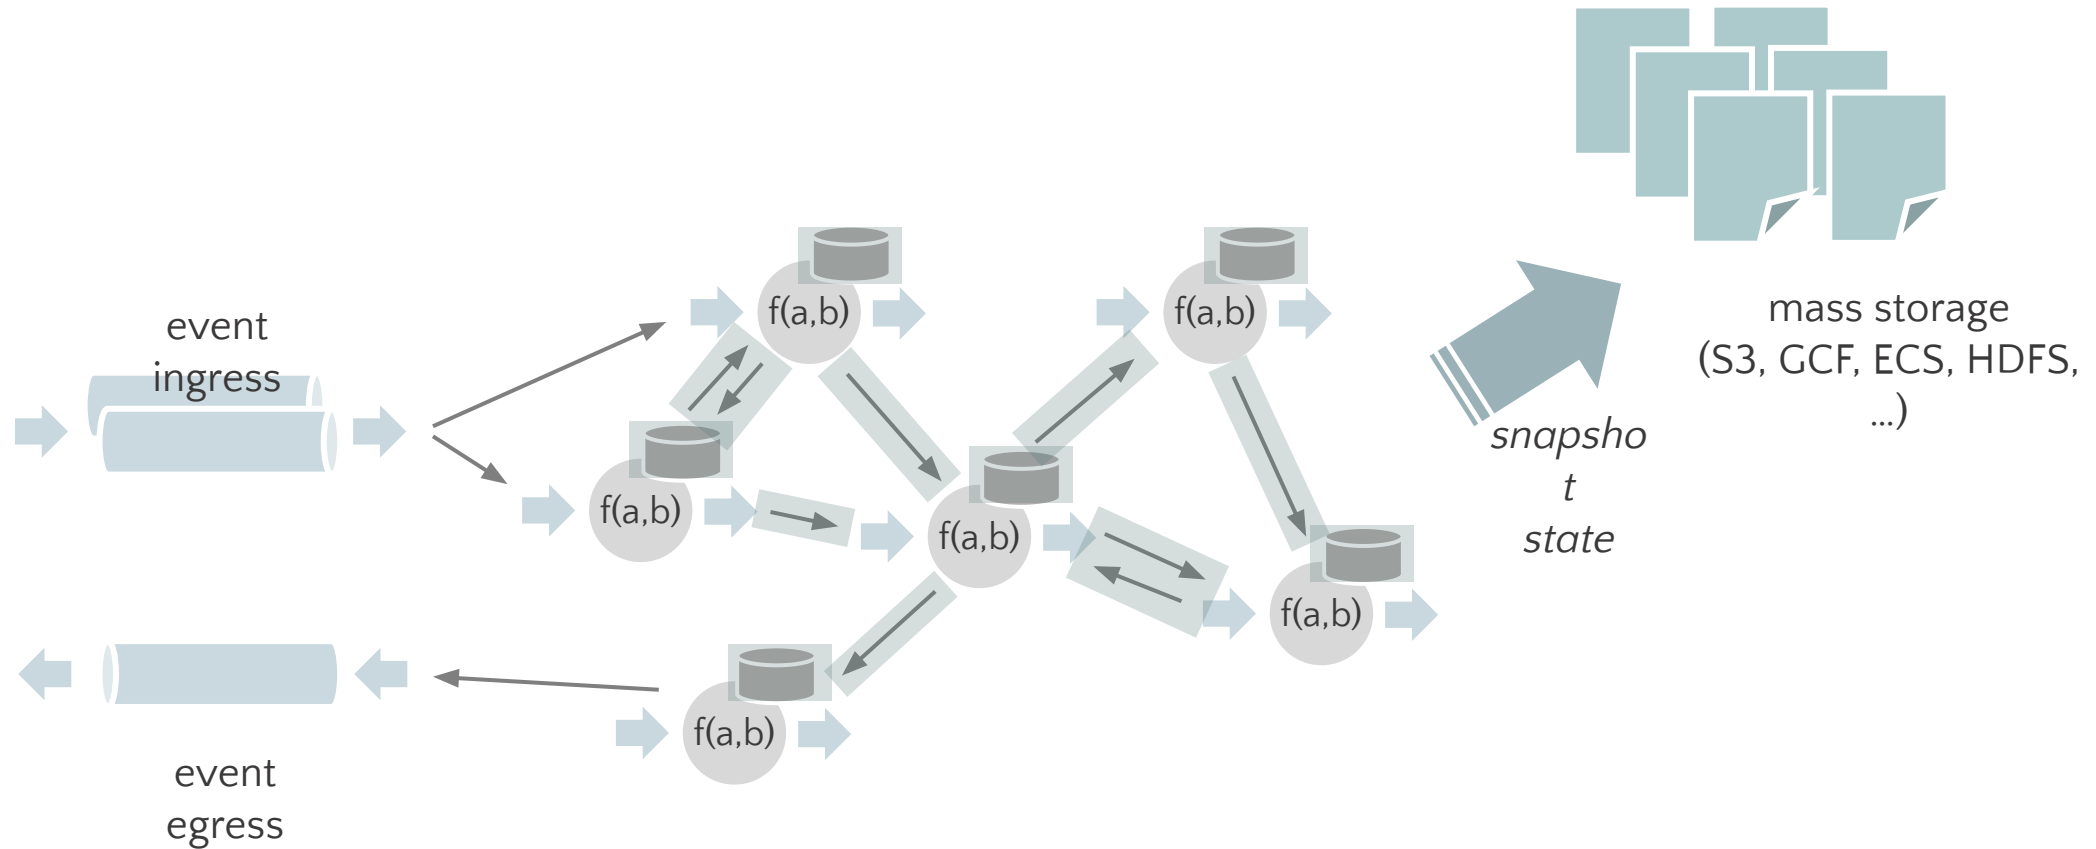
# Stateful Functions



Functions have locally embedded state



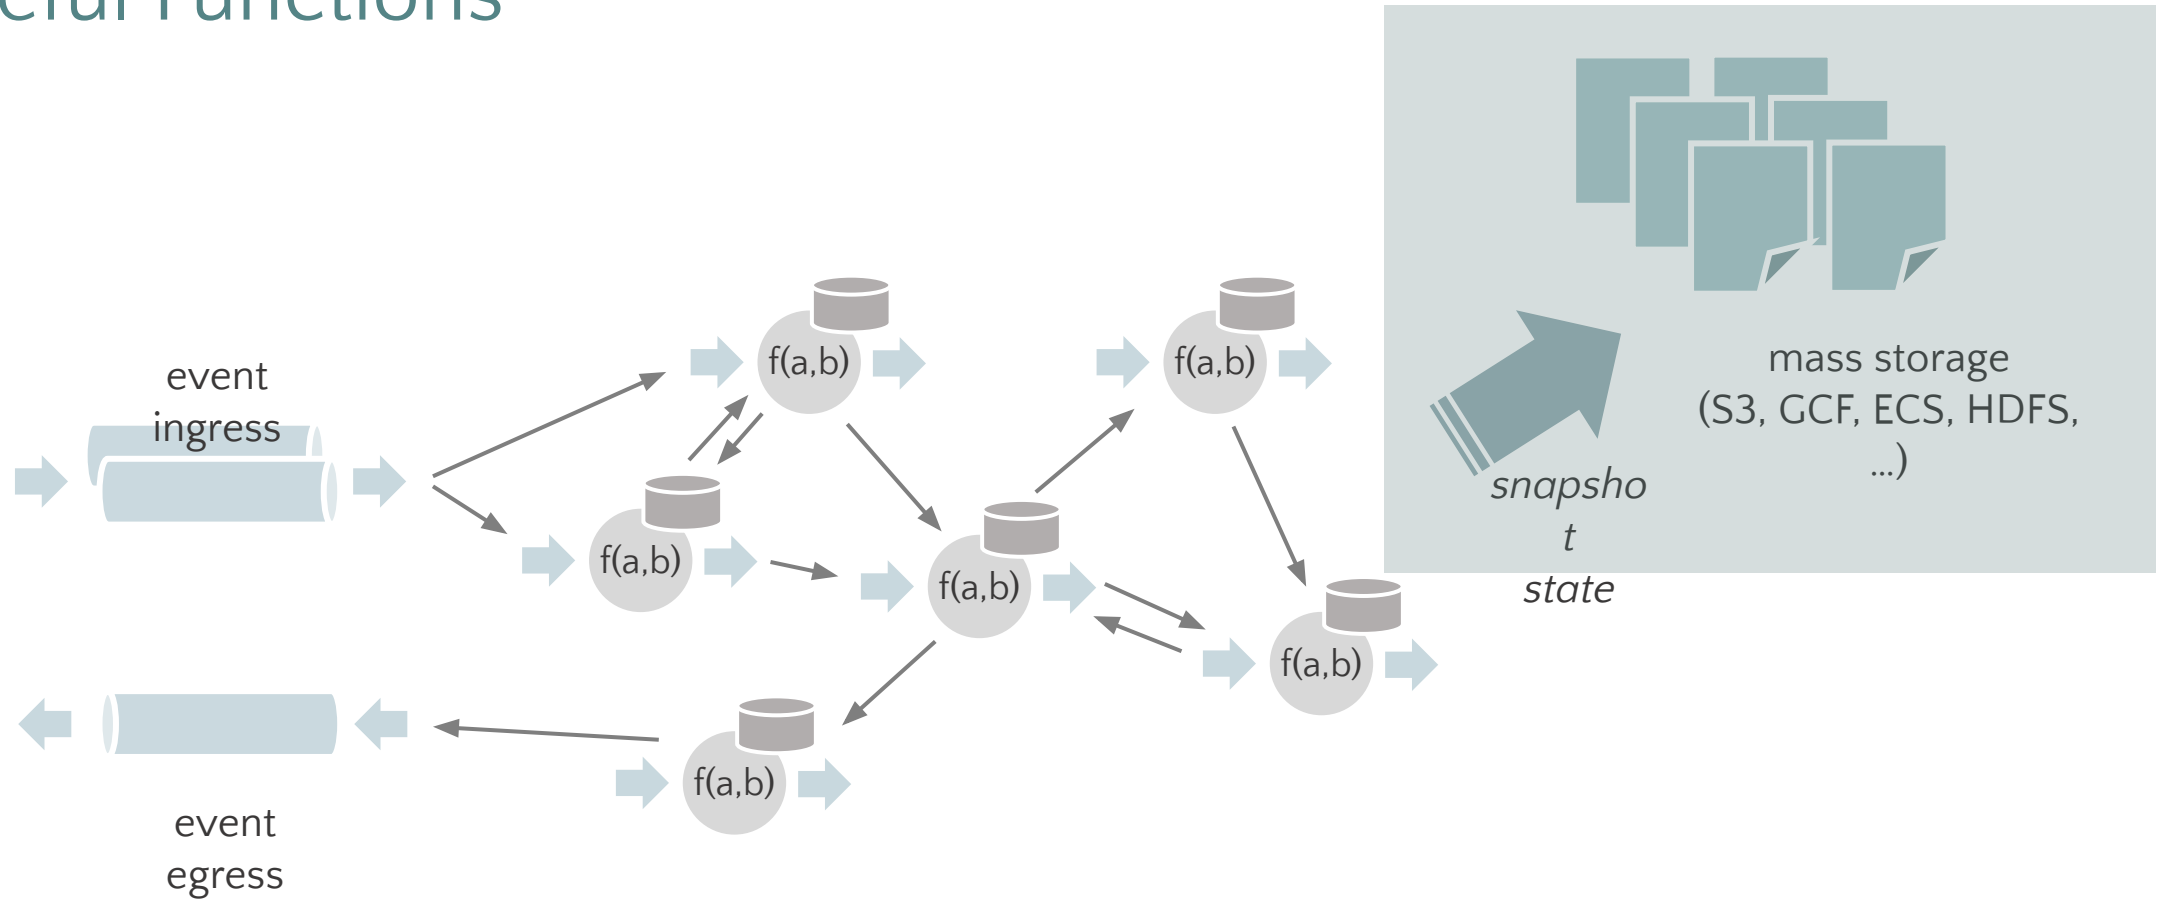
# Stateful Functions



State and messaging are consistent  
with exactly-once semantics



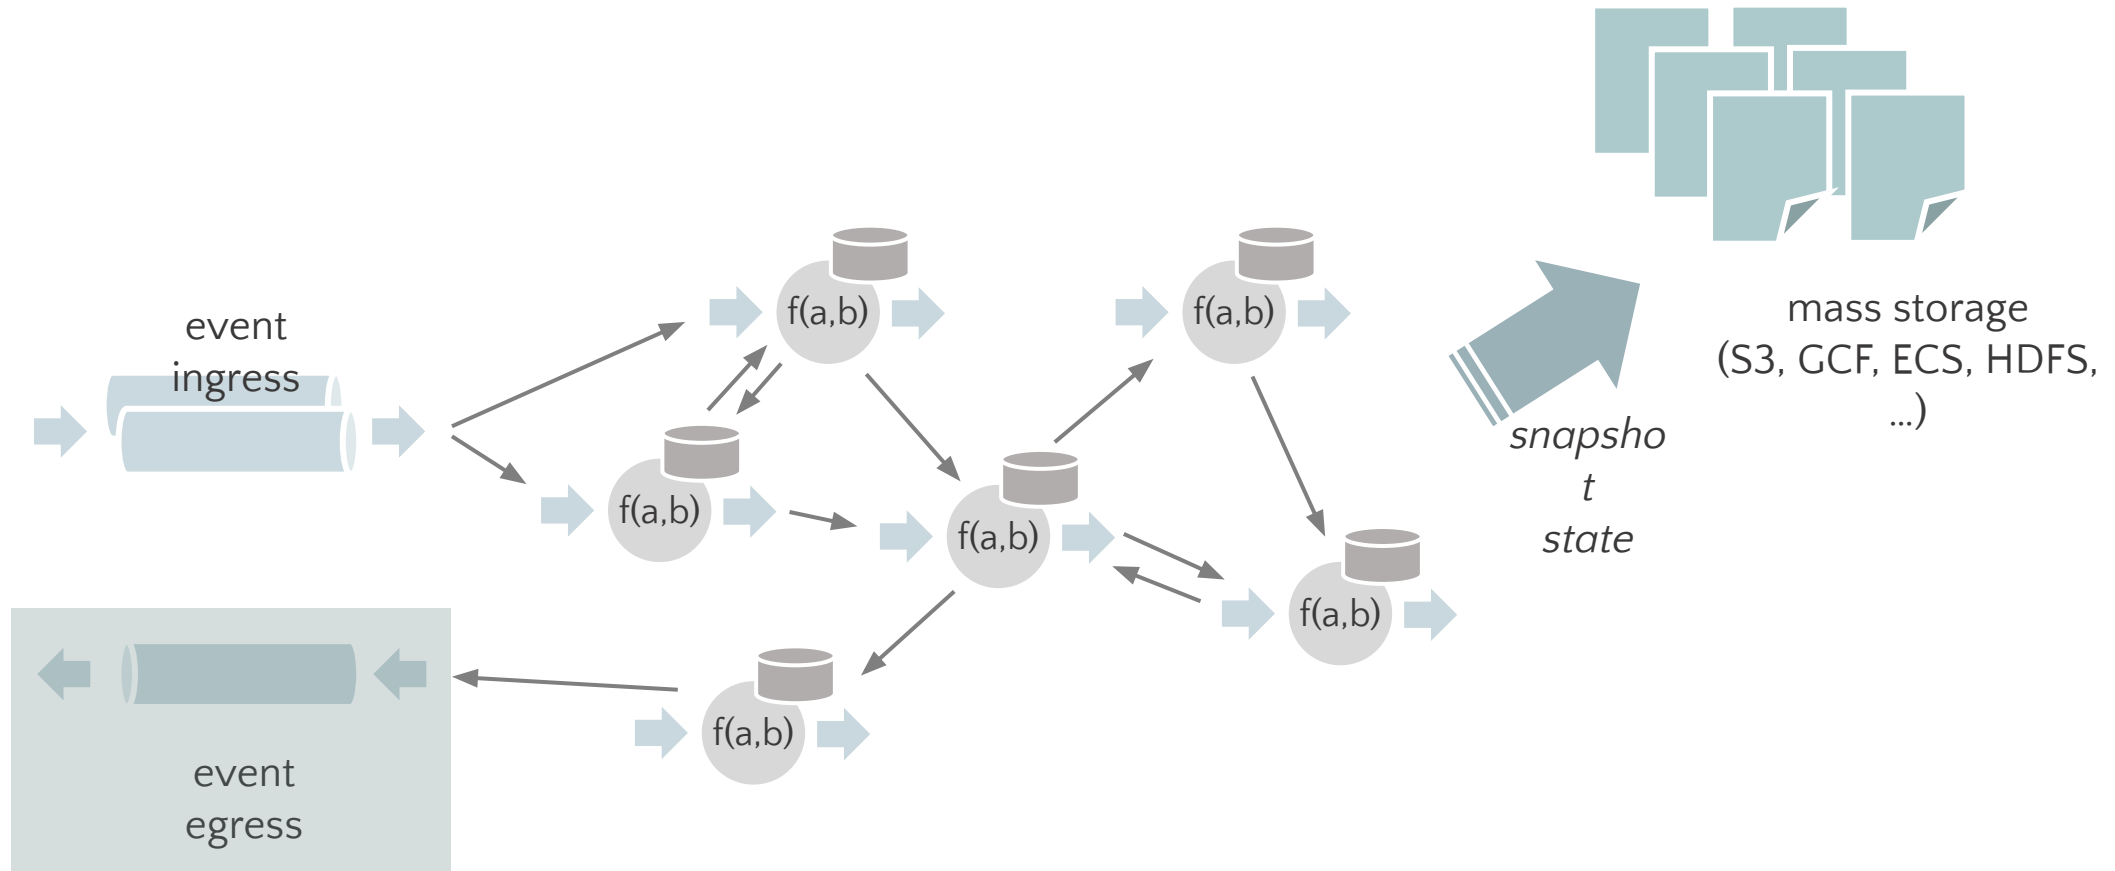
# Stateful Functions



No database required  
All persistence goes directly to blob storage



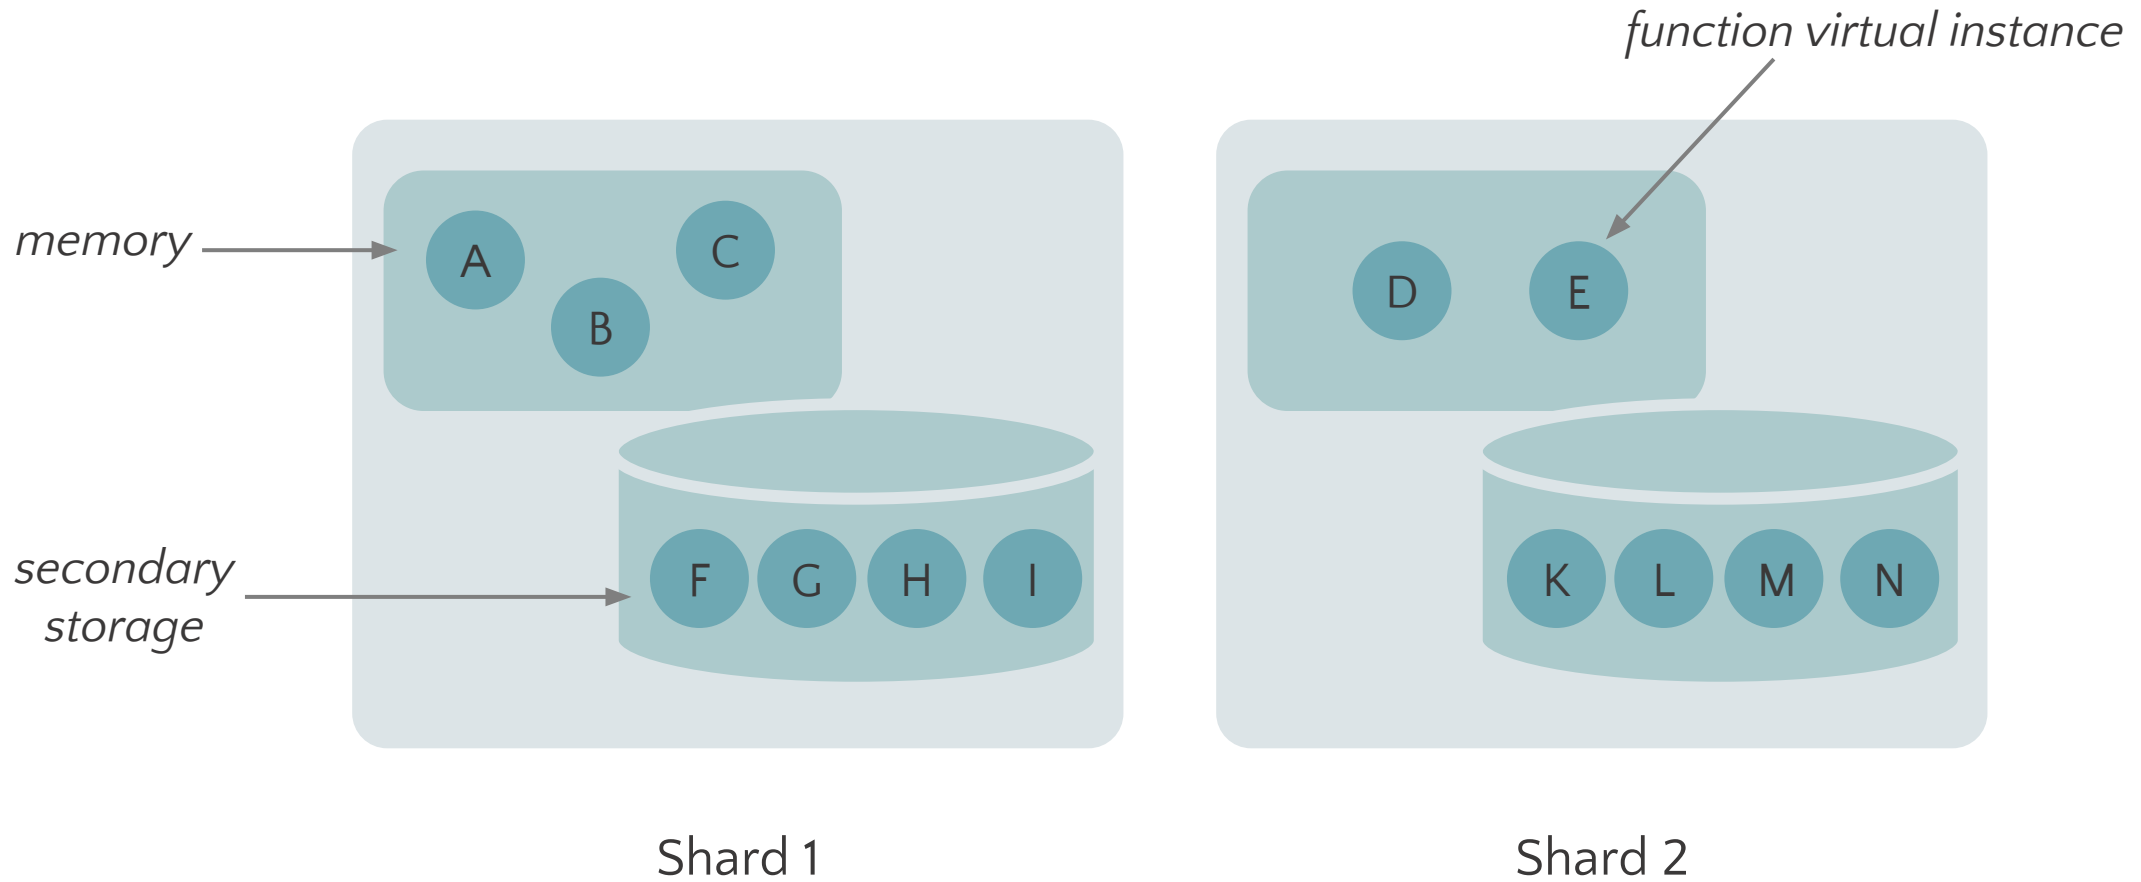
# Stateful Functions



Event egresses to respond via event streams

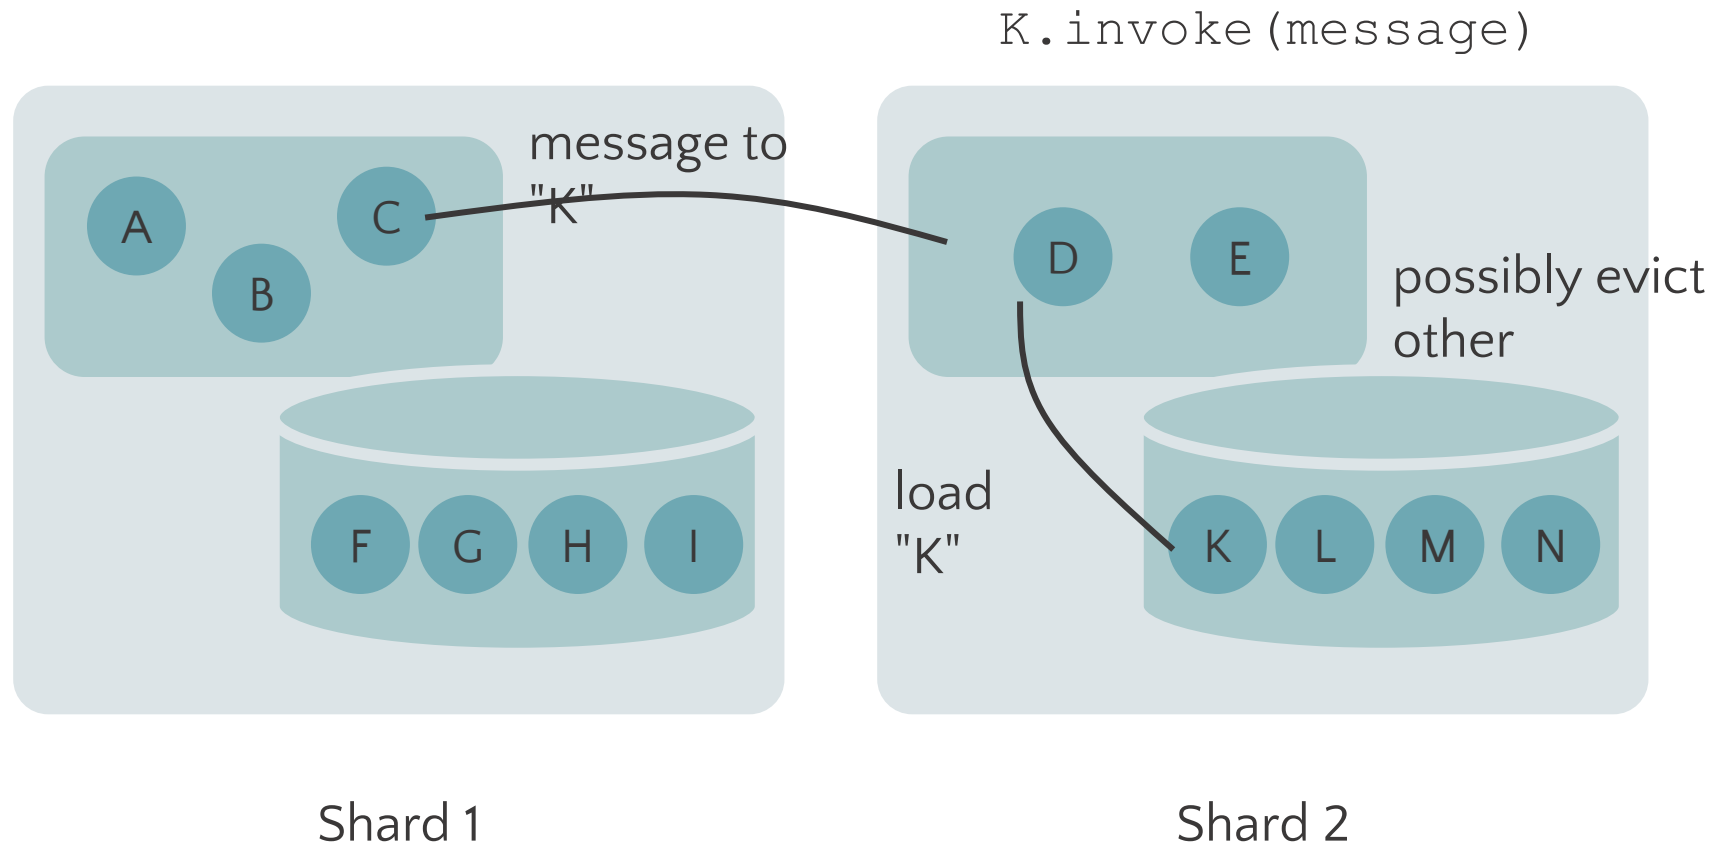


# Logical/Virtual Instances






# Logical/Virtual Instances



# SDK Concepts

- Each function is associated with a **FunctionType** and **id**
  - **FunctionType** + **id**  $\Rightarrow$  **Address**
- An **Address** is logical not physical
  - No service discovery required

**FunctionType**  $\Rightarrow$  Greeter  
**id**  $\Rightarrow$  User Id




# SDK Concepts

```
public class GreetRouter implements Router<GreetRequest> {  
  
    @Override  
    public void route(GreetRequest message, Downstream<GreetRequest> downstream) {  
        Address address = new Address(GreetFunction.TYPE, message.getUserId());  
        downstream.forward(address, message);  
    }  
}
```



# SDK Concepts

- Applications are bundles of **StatefulFunction**'s
- A stateful function reacts to incoming events and can:
  - Perform a local computation
  - Access & modify local state
  - Send a message to any other stateful function
  - Send a message to external systems
  - Send a message with delay
  - Complete an asynchronous request



# SDK Concepts

```
public class GreeterFunction implements StatefulFunction {  
  
    public static FunctionType TYPE = new FunctionType("ververica", "greeter");  
  
    @Persisted  
    PersistedValue<Integer> seenCount = PersistedValue.of("seen", Integer.class);  
  
    @Override  
    public void invoke(Context ctx, Object message) {  
        String userId = ctx.self().id();  
        int seen = seenCount.getDefault(0) + 1;  
        seenCount.set(seen);  
  
        String greeting = createGreeting(userId, seen);  
        ctx.send(Identifier.greetings, greeting)  
    }  
}
```



# SDK Concepts

```
public class GreeterFunction implements StatefulFunction {  
    public static FunctionType TYPE = new FunctionType("ververica", "greeter");  
  
    @Persisted  
    PersistedValue<Integer> seenCount = PersistedValue.of("seen", Integer.class);  
  
    @Override  
    public void invoke(Context ctx, Object message) {  
        String userId = ctx.self().id();  
        int seen = seenCount.getDefault(0) + 1;  
        seenCount.set(seen);  
  
        String greeting = createGreeting(userId, seen);  
        ctx.send(Identifier.greetings, greeting)  
    }  
}
```



# SDK Concepts

```
public class GreeterFunction implements StatefulFunction {  
  
    public static FunctionType TYPE = new FunctionType("ververica", "greeter");  
  
    @Persisted  
    PersistedValue<Integer> seenCount = PersistedValue.of("seen", Integer.class);  
  
    @Override  
    public void invoke(Context ctx, Object message) {  
        String userId = ctx.self().id();  
        int seen = seenCount.getDefault(0) + 1;  
        seenCount.set(seen);  
  
        String greeting = createGreeting(userId, seen);  
        ctx.send(Identifier.greetings, greeting)  
    }  
}
```



# SDK Concepts

```
public class GreeterFunction implements StatefulFunction {  
  
    public static FunctionType TYPE = new FunctionType("ververica", "greeter");  
  
    @Persisted  
    PersistedValue<Integer> seenCount = PersistedValue.of("seen", Integer.class);  
  
    @Override  
    public void invoke(Context ctx, Object message) {  
        String userId = ctx.self().id();  
        int seen = seenCount.getDefault(0) + 1;  
        seenCount.set(seen);  
  
        String greeting = createGreeting(userId, seen);  
        ctx.send(Identifier.greetings, greeting)  
    }  
}
```






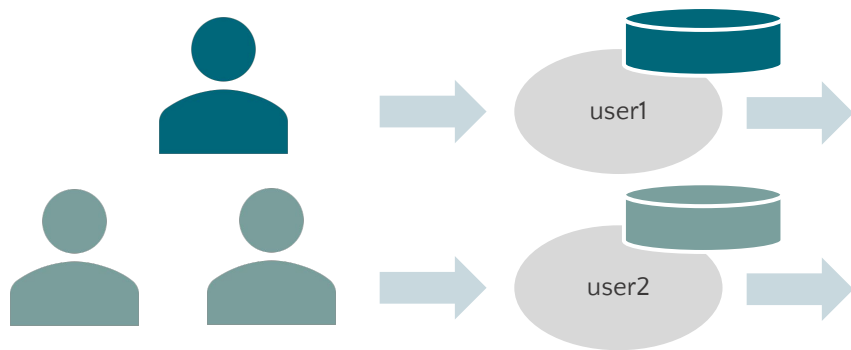
# SDK Concepts

```
public class GreeterFunction implements StatefulFunction {  
  
    public static FunctionType TYPE = new FunctionType("ververica", "greeter");  
  
    @Persisted  
    PersistedValue<Integer> seenCount = PersistedValue.of("seen", Integer.class);  
  
    @Override  
    public void invoke(Context ctx, Object message) {  
        String userId = ctx.self().id();  
        int seen = seenCount.getDefault(0) + 1;  
        seenCount.set(seen);  
  
        String greeting = createGreeting(userId, seen);  
        ctx.send(Identifier.greetings, greeting)  
    }  
}
```

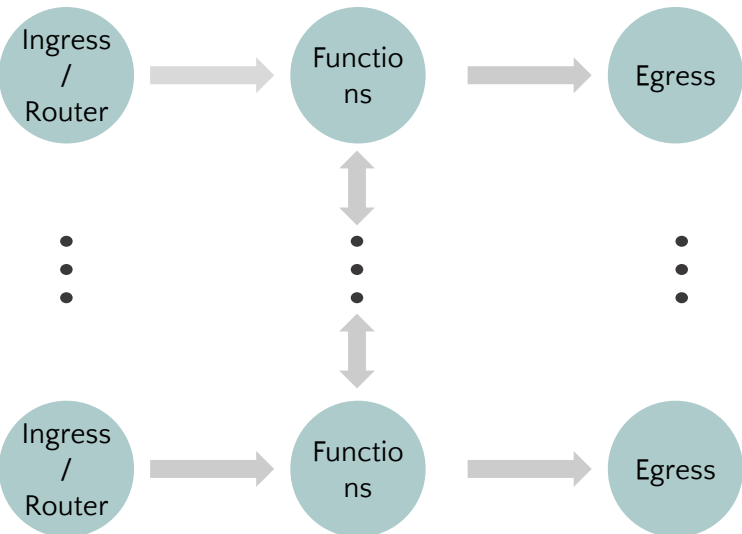


# Execution Model

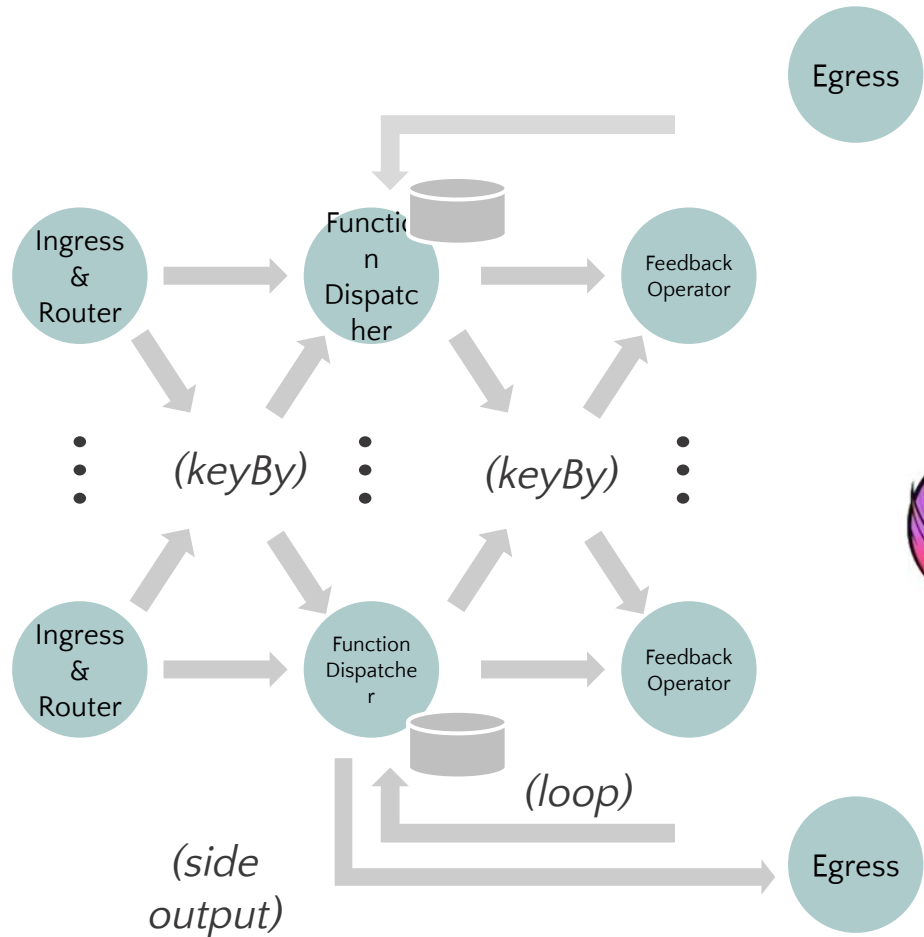
- There is a single (logical) instance of a function per **Address**
- Each **StatefulFunction** is created on demand, transparently by the runtime
- Messages sent to an **Address** are processed by a single thread
- Messages sent from function A to function B are always received in FIFO order



# Apache Flink is the State and Event Streaming Fabric



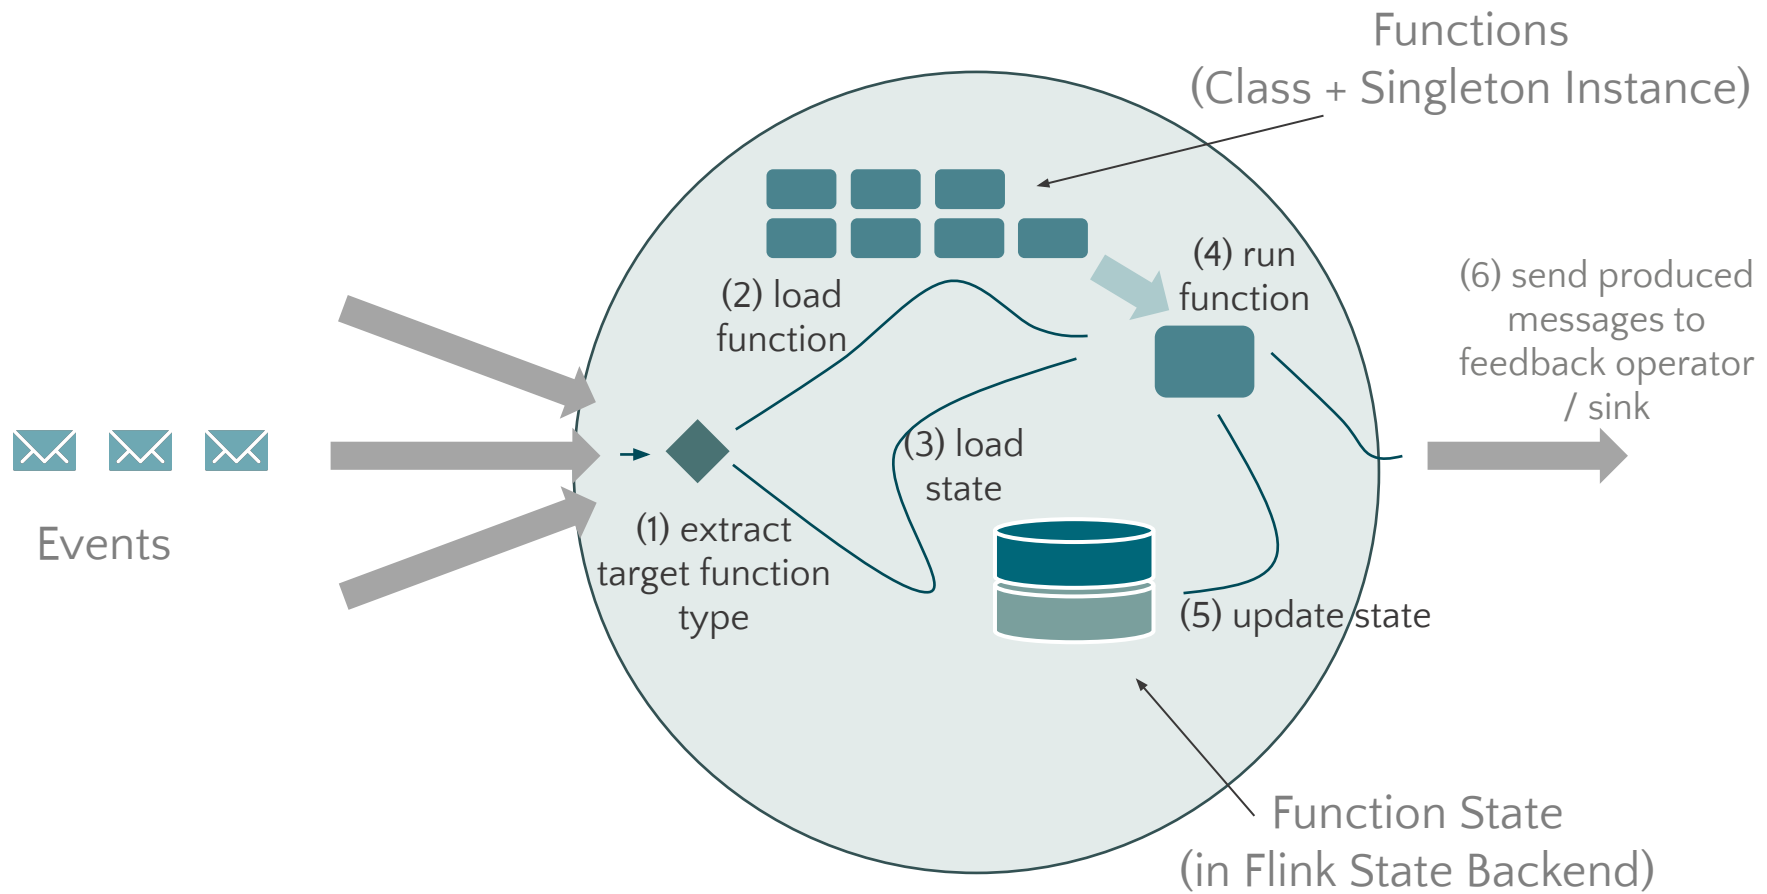
Conceptual Dataflow



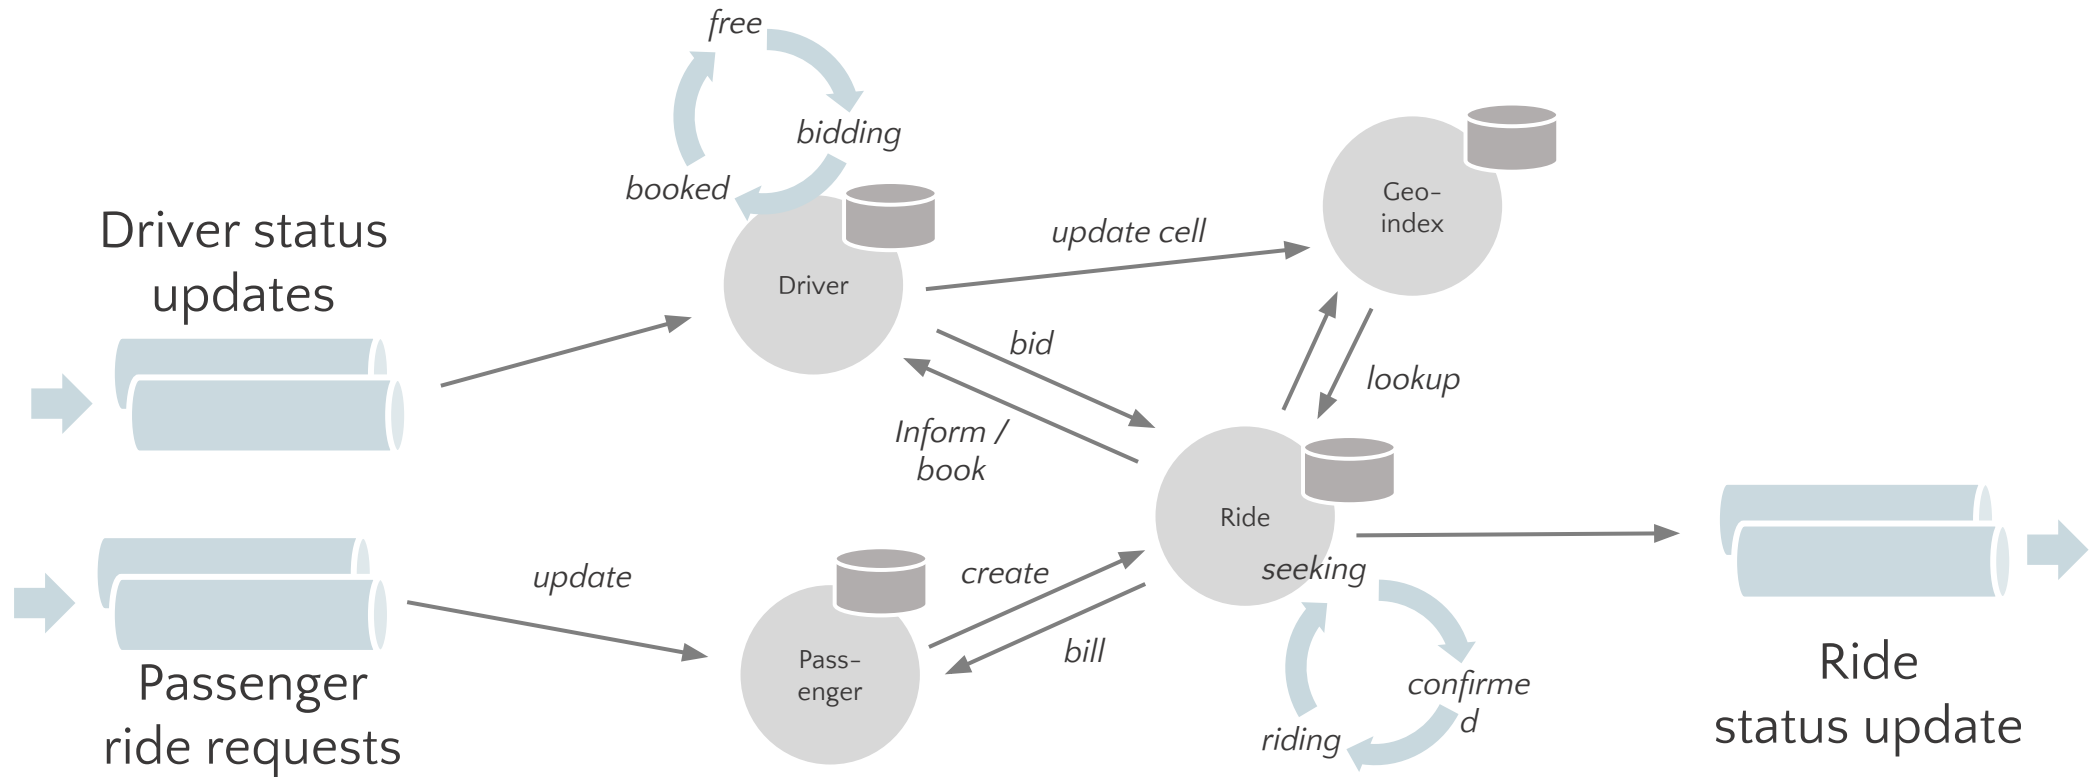
Apache Flink Dataflow Graph



# A Peek Under The Hood



# Example: Ride Sharing App



# Stream Processing Streaming SQL

*event/stream-centric*  
c

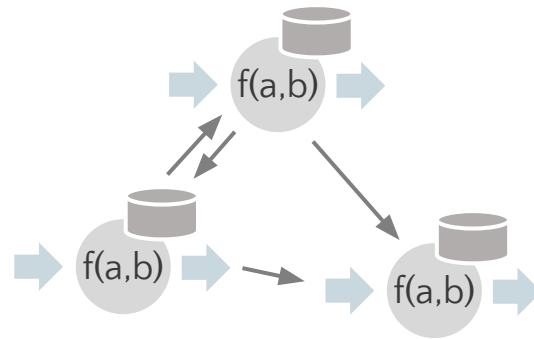


data preparation  
combining knowledge/information

filtering, enriching,  
aggregating, joining events

# Stateful Functions

*state-centric*  
c

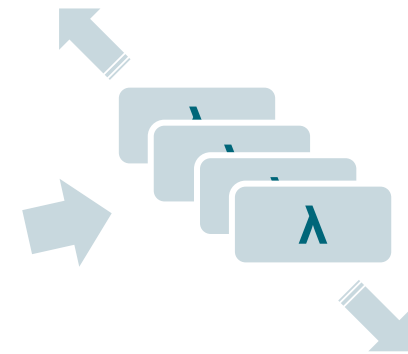


coordination,  
(interacting) state machines

complex event/state  
interactions

# F-a-a-S

*stateless /  
compute-centric*




“occasional” actions or  
spiky loads

compute-intensive  
or blocking




# Putting it all together: Ridesharing again

## Stream Processing




## Stateful Functions



ride life-cycle  
driver-to-ride matching

## FaaS

render map/route image  
create a receipt PDF  
send email





THANK YOU!

[www.statefun.io](http://www.statefun.io)

[@statefun\\_io](https://twitter.com/statefun_io)

[github.com/ververica/stateful-functions](https://github.com/ververica/stateful-functions)