

Streaming
All the things!
with



Akka streams

Johan Andrén

Voxxeddays, Zürich, 2017-02-23



Johan Andrén

Akka Team

Stockholm Scala User Group



@apnylle

johan.andren@lightbend.com



Akka

Make building powerful concurrent & distributed applications **simple**.

Akka is a toolkit and runtime for building highly concurrent, distributed, and resilient **message-driven** applications on the JVM

What's in the toolkit?

Actors – simple & high performance concurrency

Cluster / Remoting – location transparency, resilience

Cluster tools – and more prepackaged patterns

Streams – back-pressured stream processing

Persistence – Event Sourcing

HTTP – complete, fully async and reactive HTTP Server

Official **Kafka, Cassandra, DynamoDB integrations**, tons more in the community

Complete **Java & Scala APIs** for all features

Reactive Streams

Reactive Streams timeline

Oct 2013

RxJava, Akka and Twitter-
people meeting

Apr 2015

Reactive Streams Spec 1.0
TCK
5+ impls

Akka Streams, RxJava
Vert.x, MongoDB, ...

“Soon thereafter” 2013

Reactive Streams
Expert group formed

??? 2015

JEP-266
inclusion in JDK9



Reactive Streams

“Reactive Streams is an initiative to provide a standard for asynchronous stream processing with non-blocking back pressure. This encompasses efforts aimed at runtime environments (JVM and JavaScript) as well as network protocols

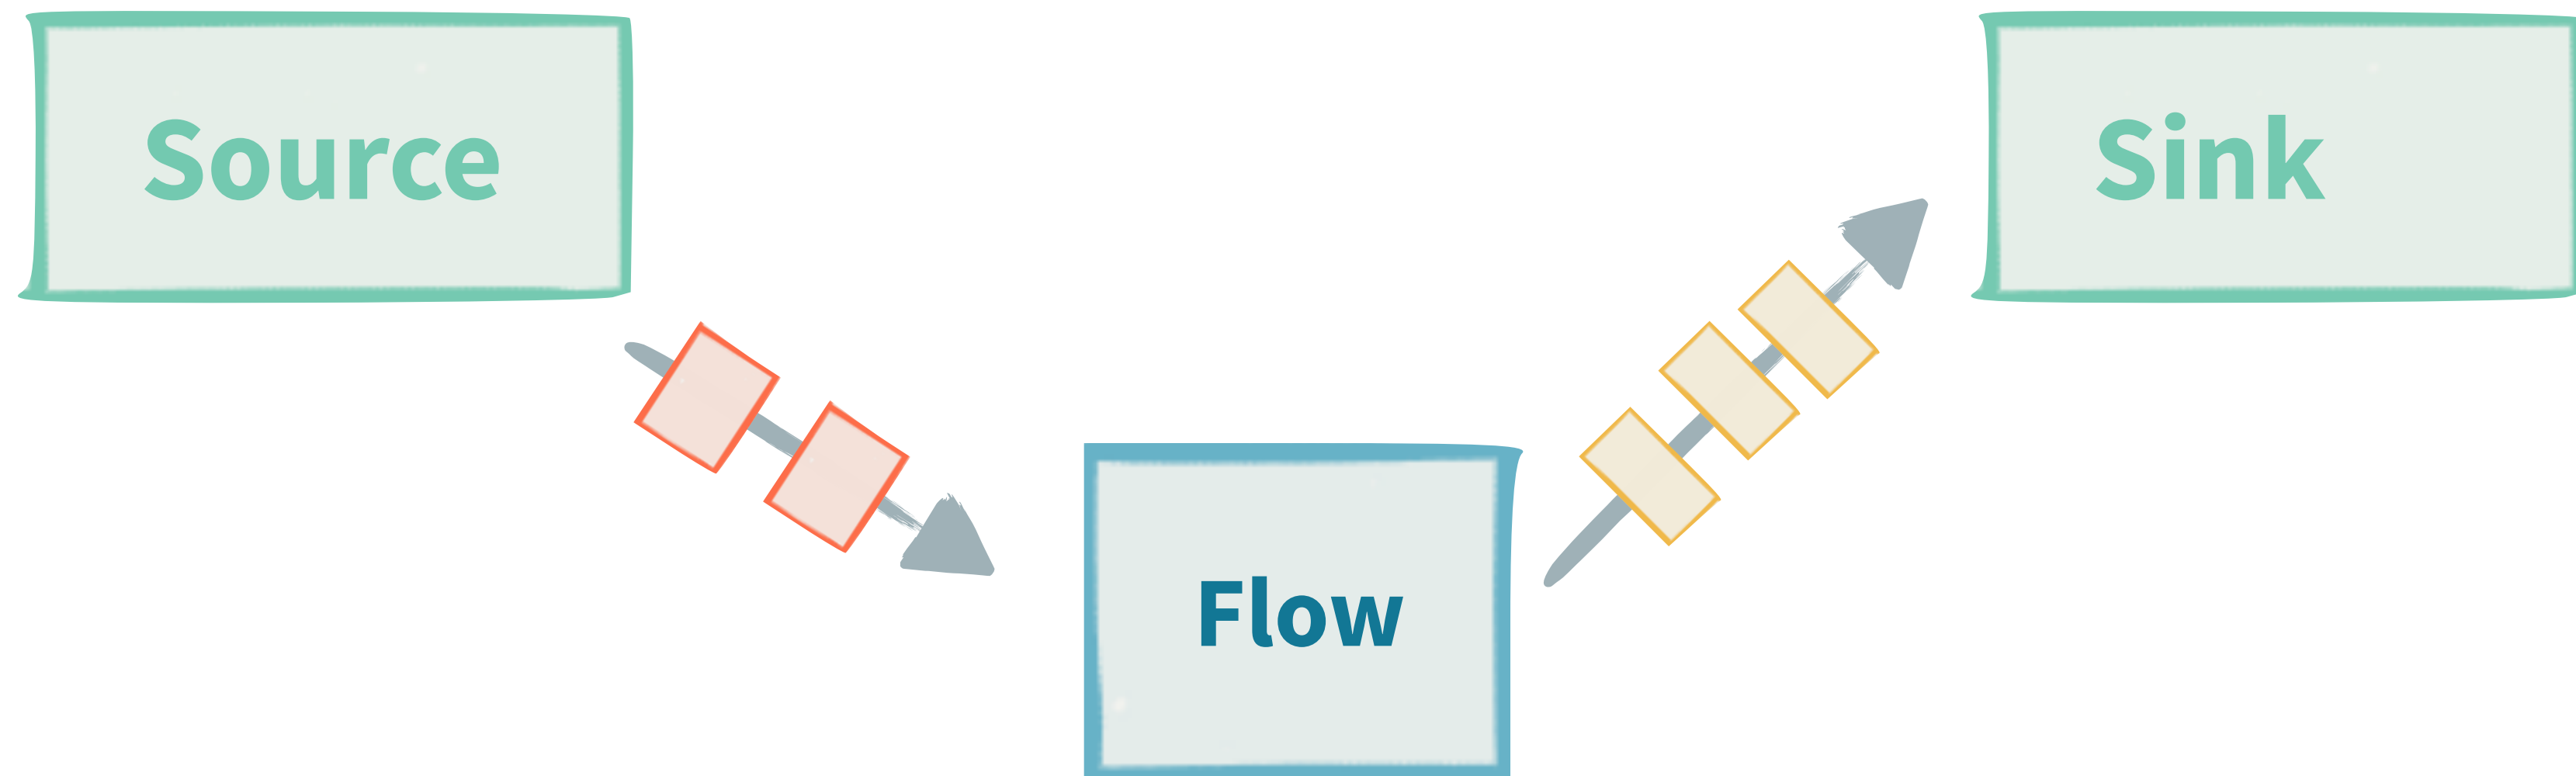
<http://www.reactive-streams.org>

Reactive Streams

“Reactive Streams is an initiative to provide a standard for asynchronous stream processing with non-blocking back pressure. This encompasses efforts aimed at runtime environments (JVM and JavaScript) as well as network protocols

<http://www.reactive-streams.org>

Stream processing

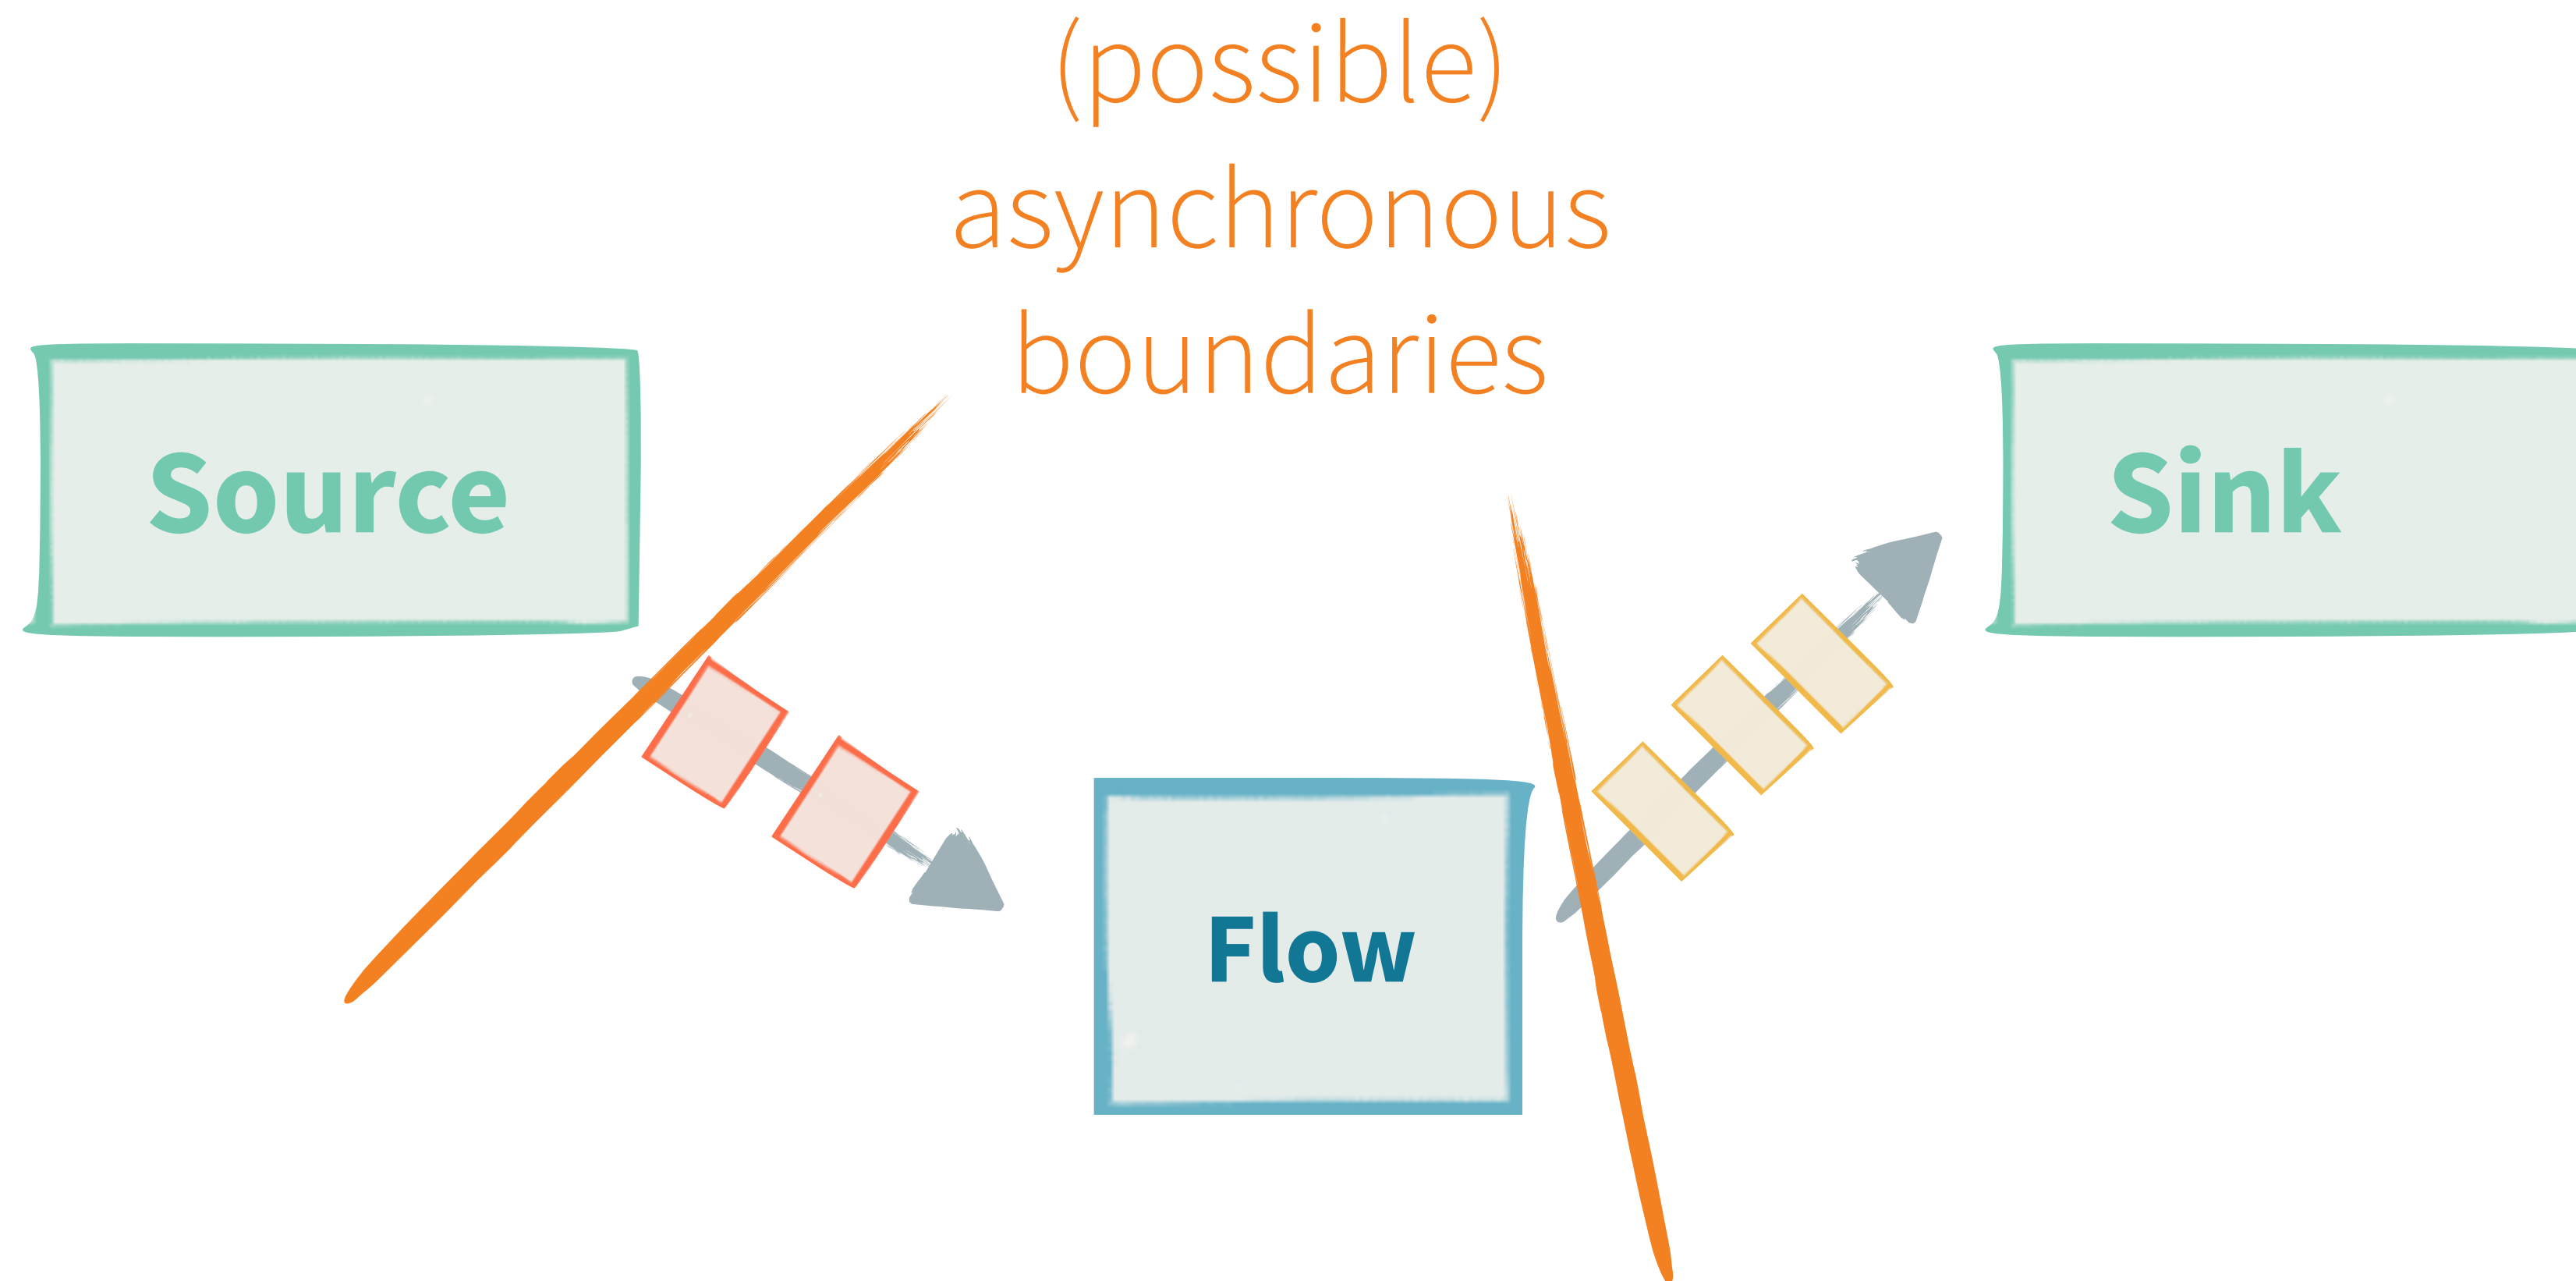


Reactive Streams

“Reactive Streams is an initiative to provide a standard for asynchronous stream processing with non-blocking back pressure. This encompasses efforts aimed at runtime environments (JVM and JavaScript) as well as network protocols

<http://www.reactive-streams.org>

Asynchronous stream processing

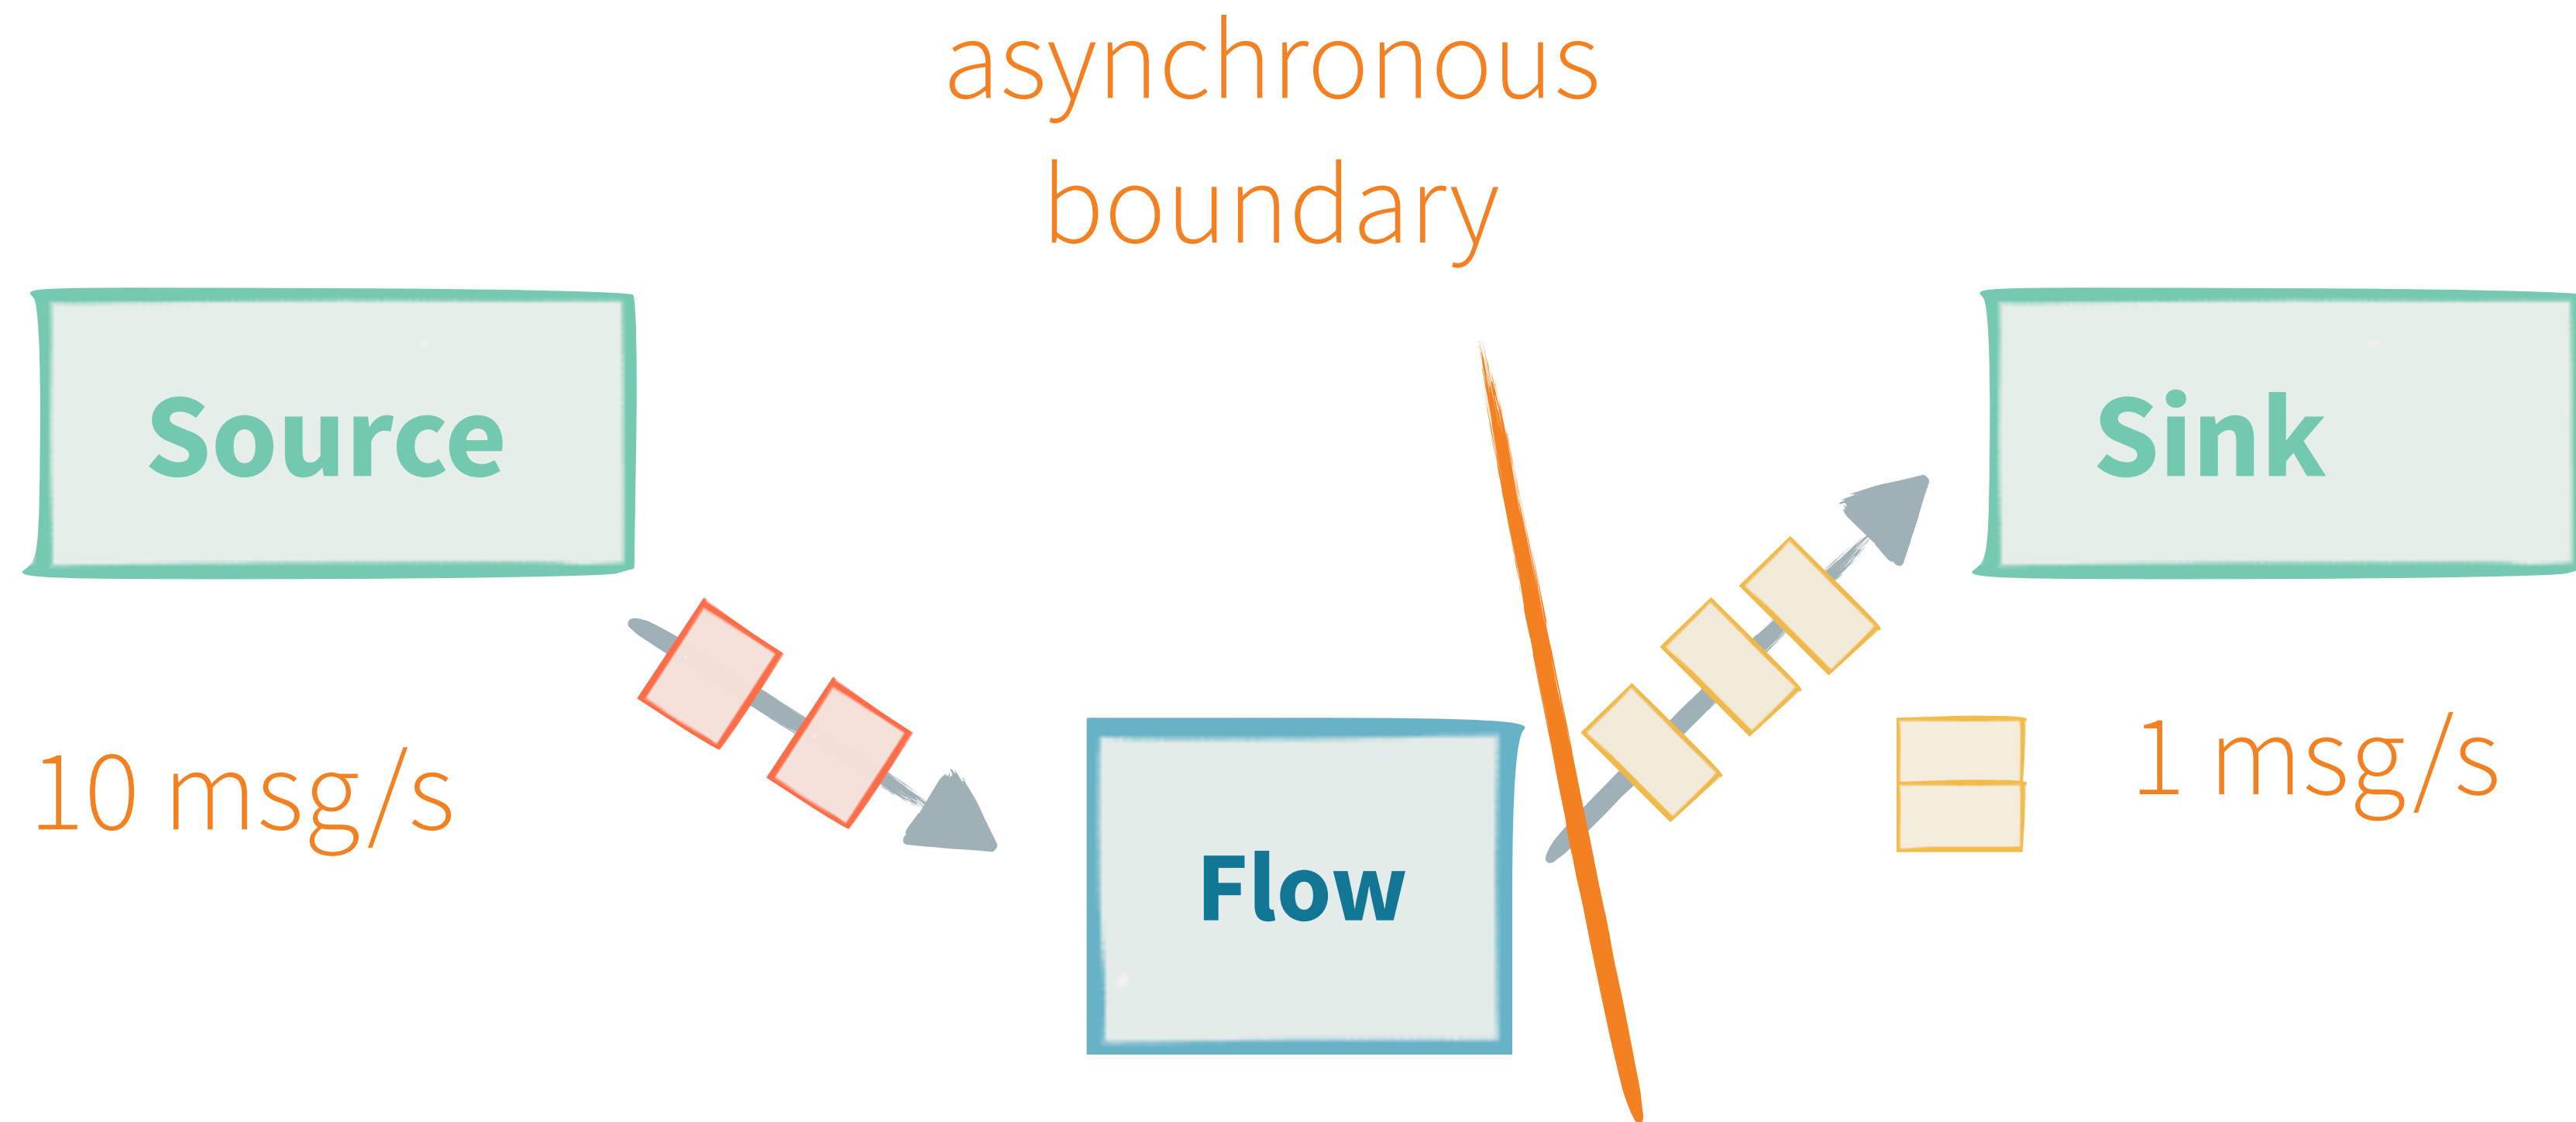


Reactive Streams

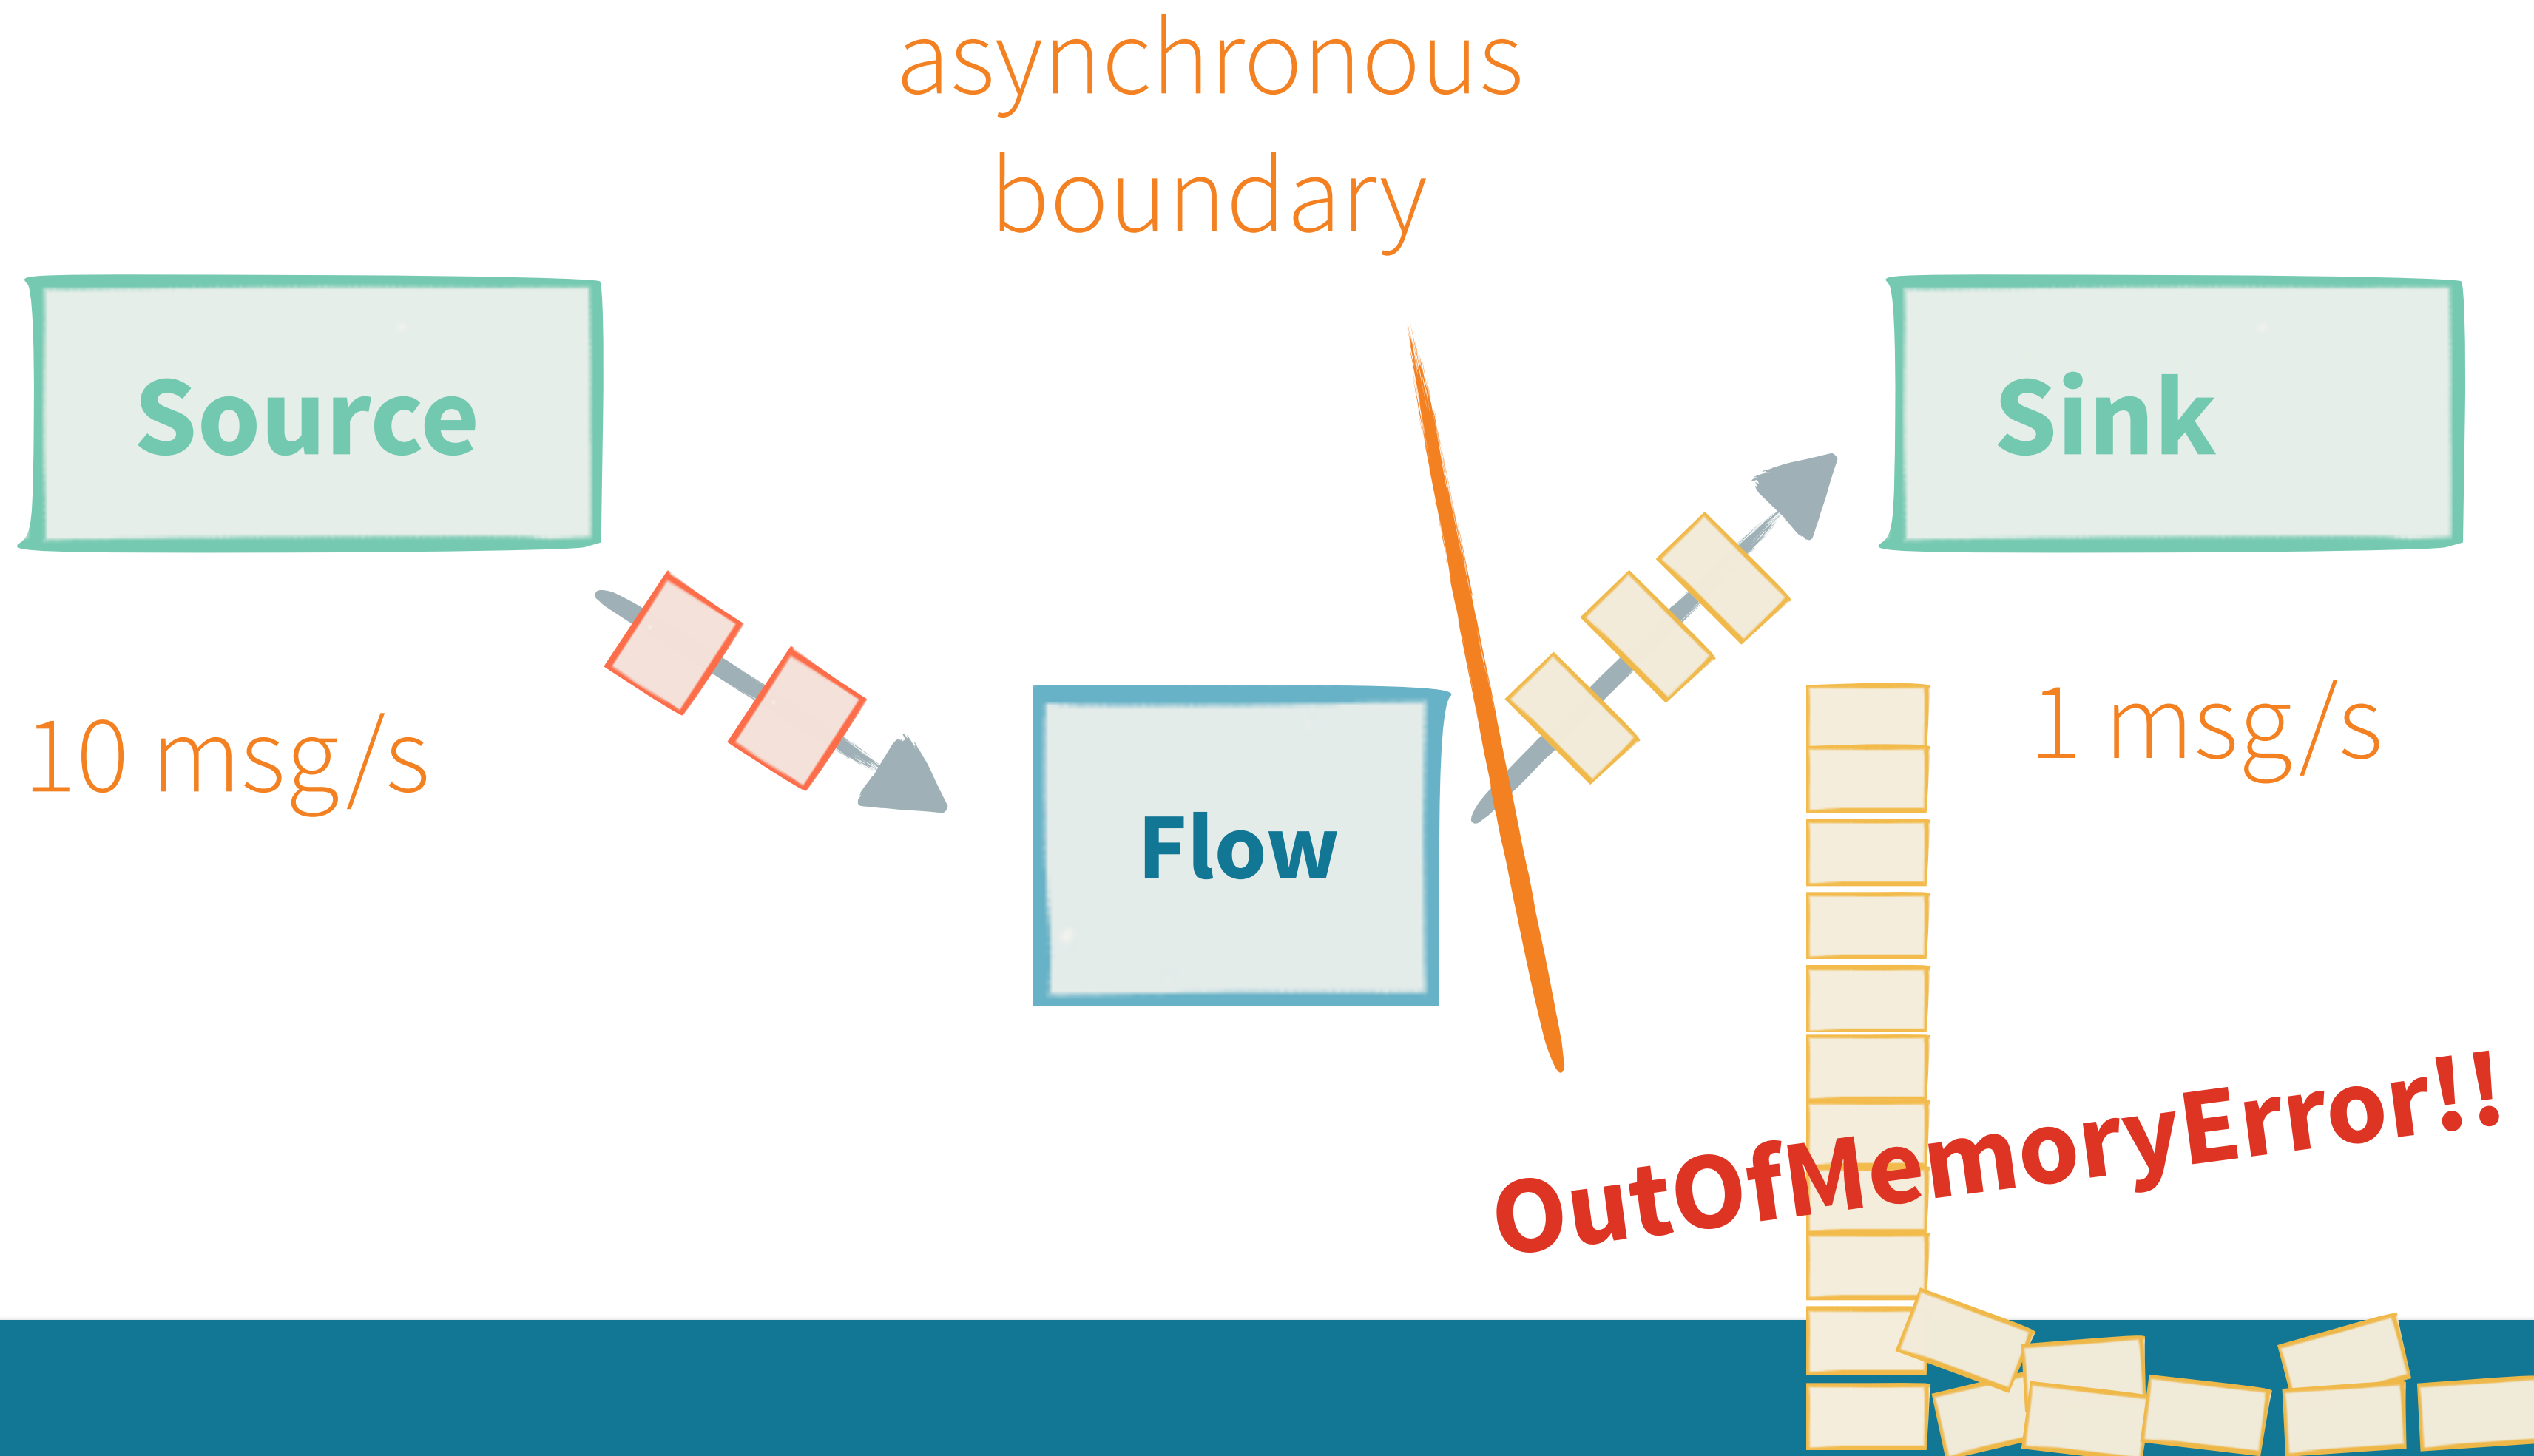
“Reactive Streams is an initiative to provide a standard for asynchronous stream processing with non-blocking back pressure. This encompasses efforts aimed at runtime environments (JVM and JavaScript) as well as network protocols

<http://www.reactive-streams.org>

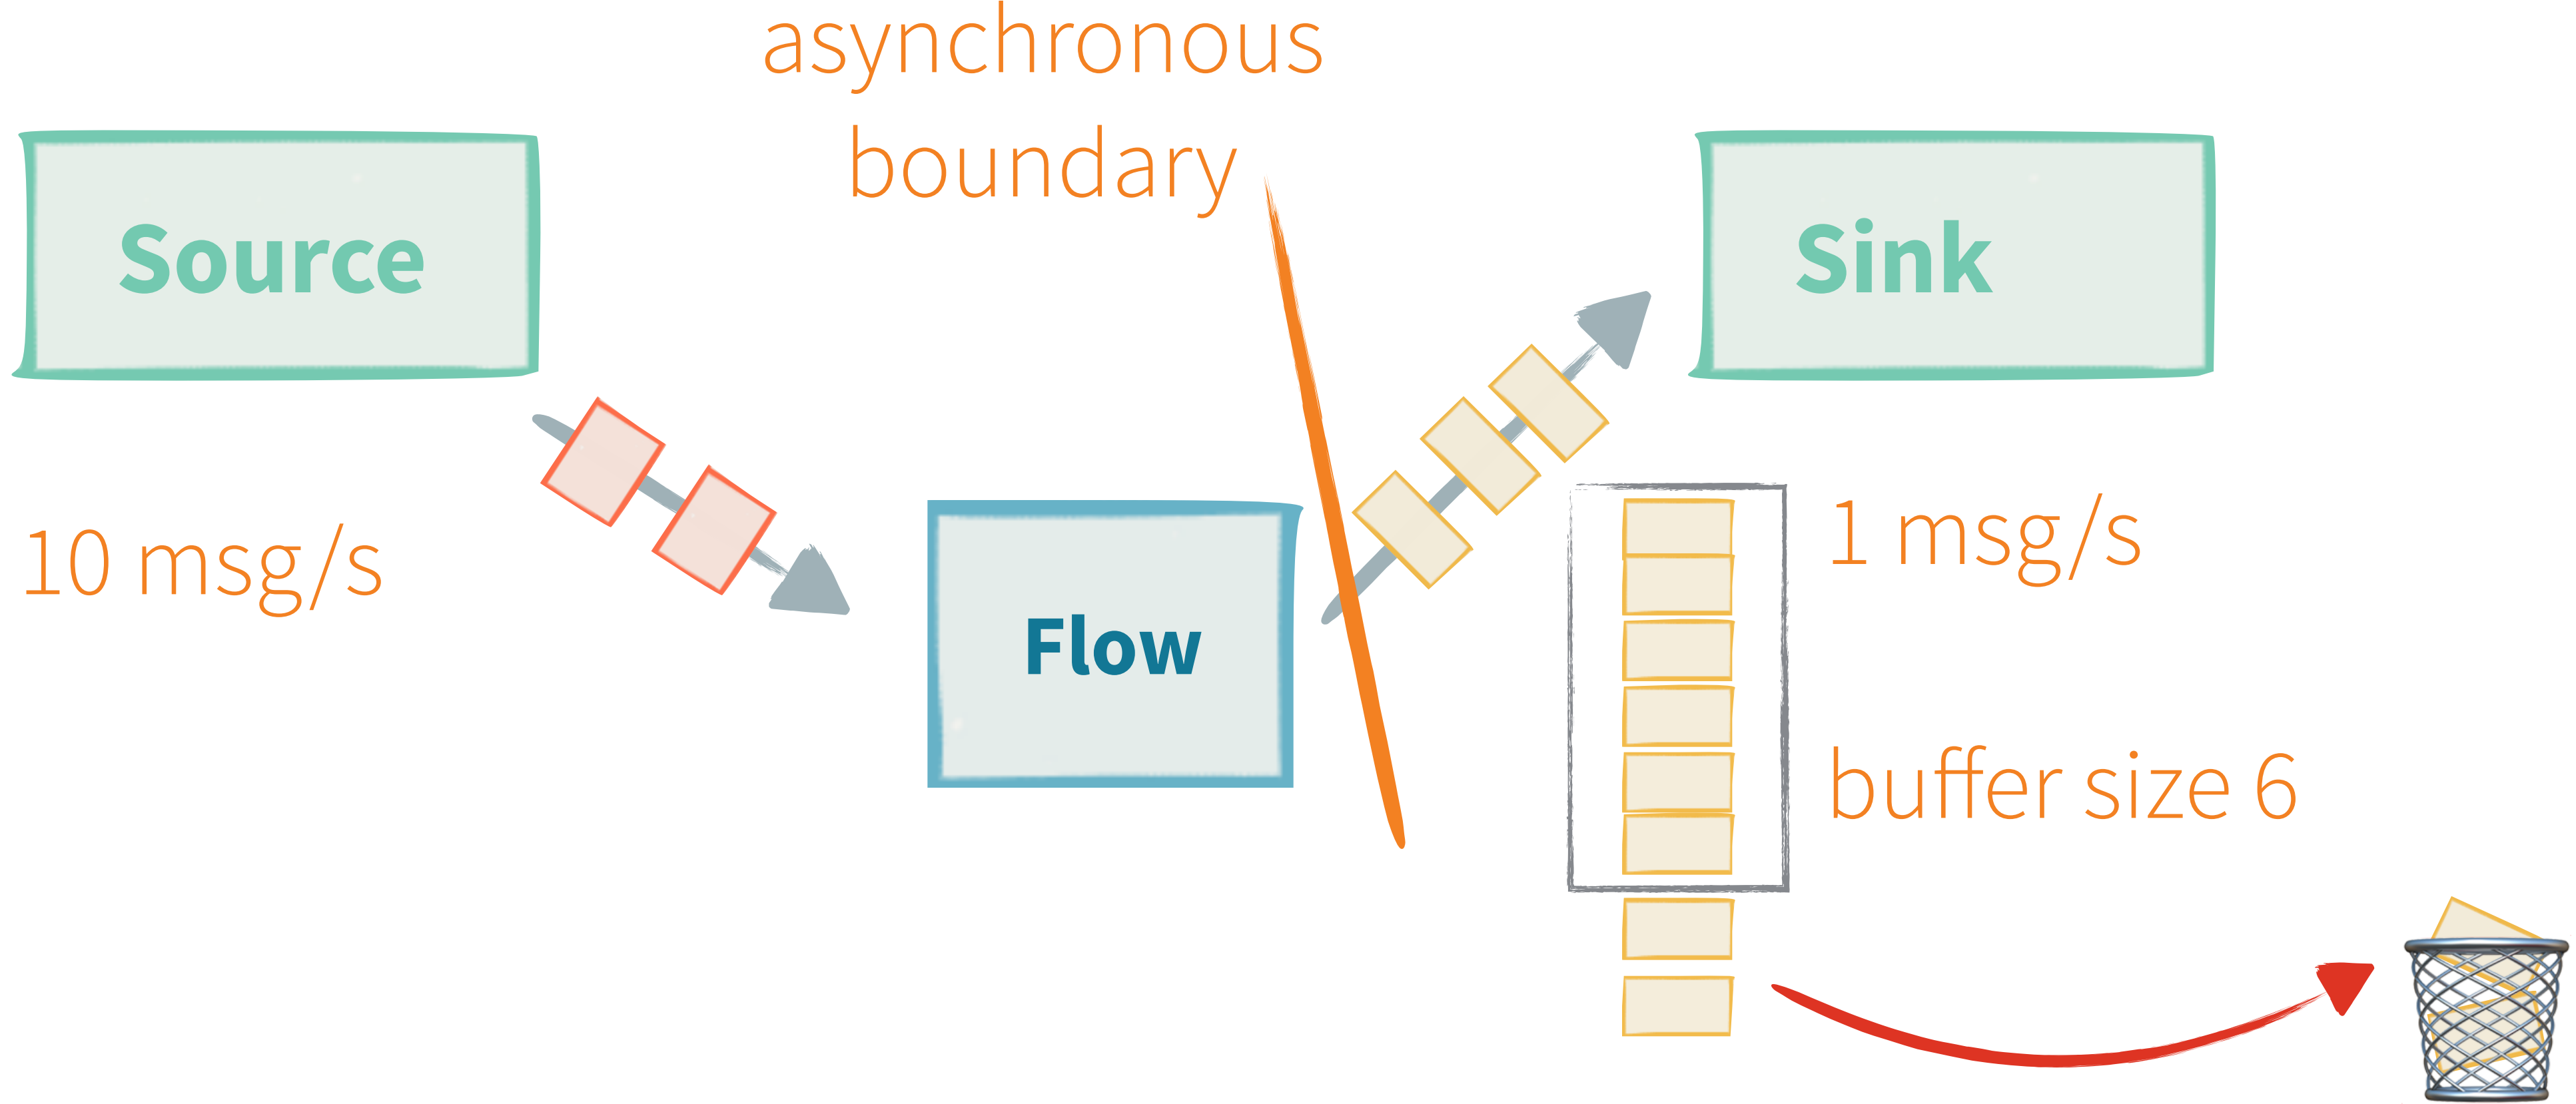
No back pressure



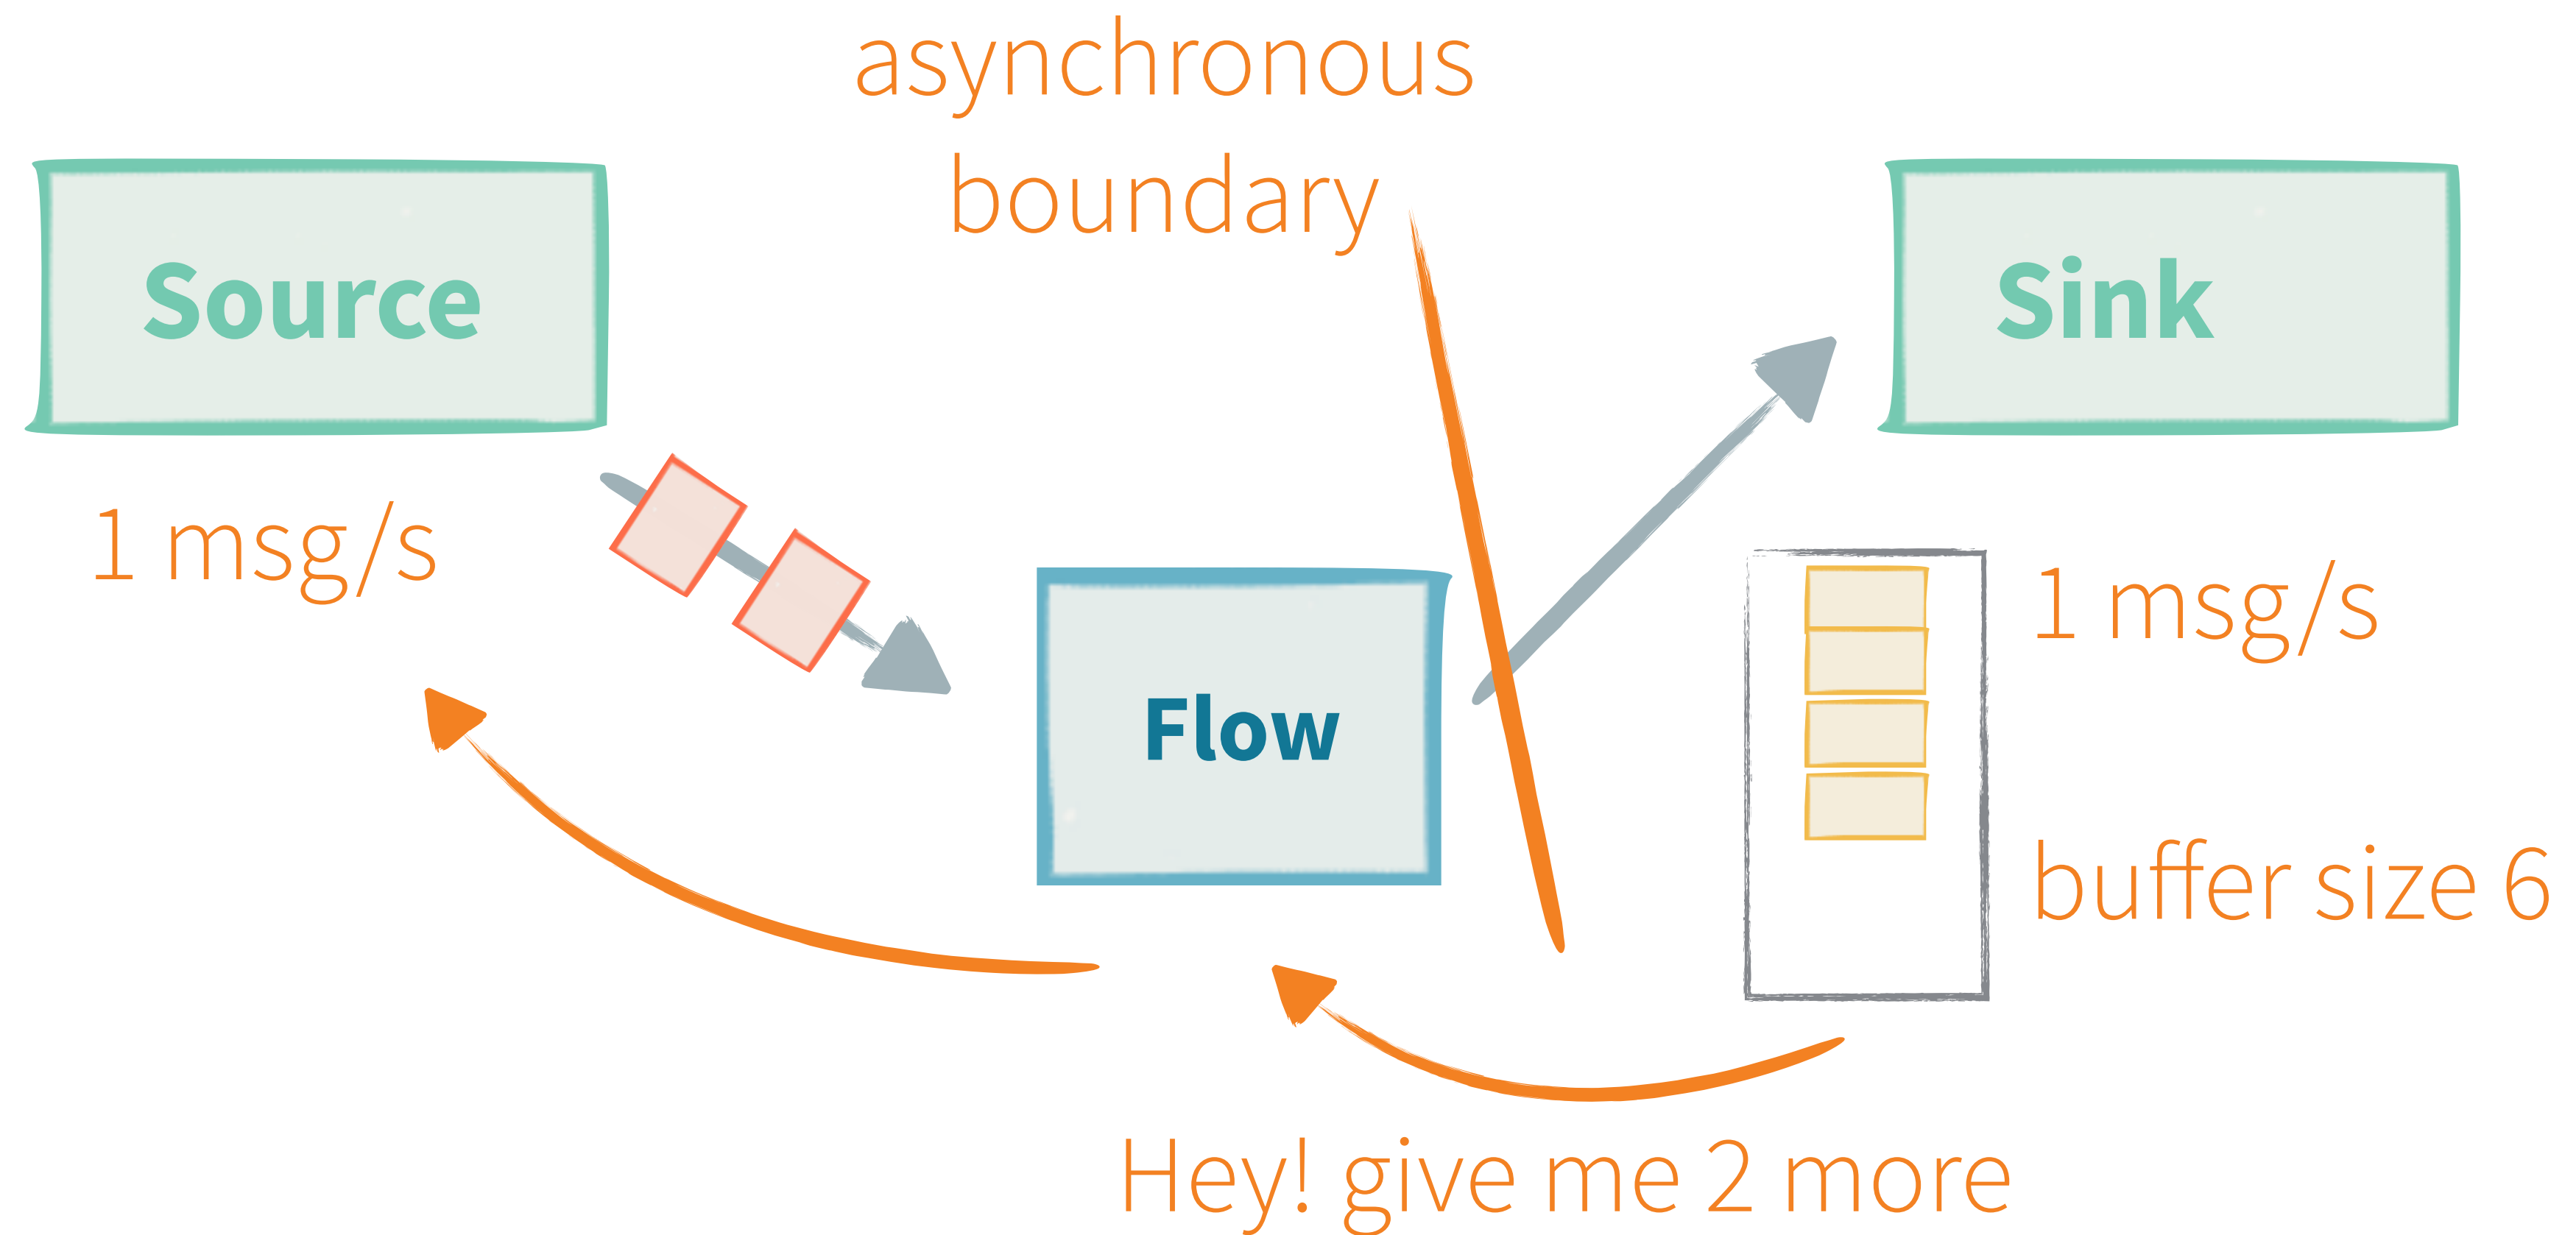
No back pressure



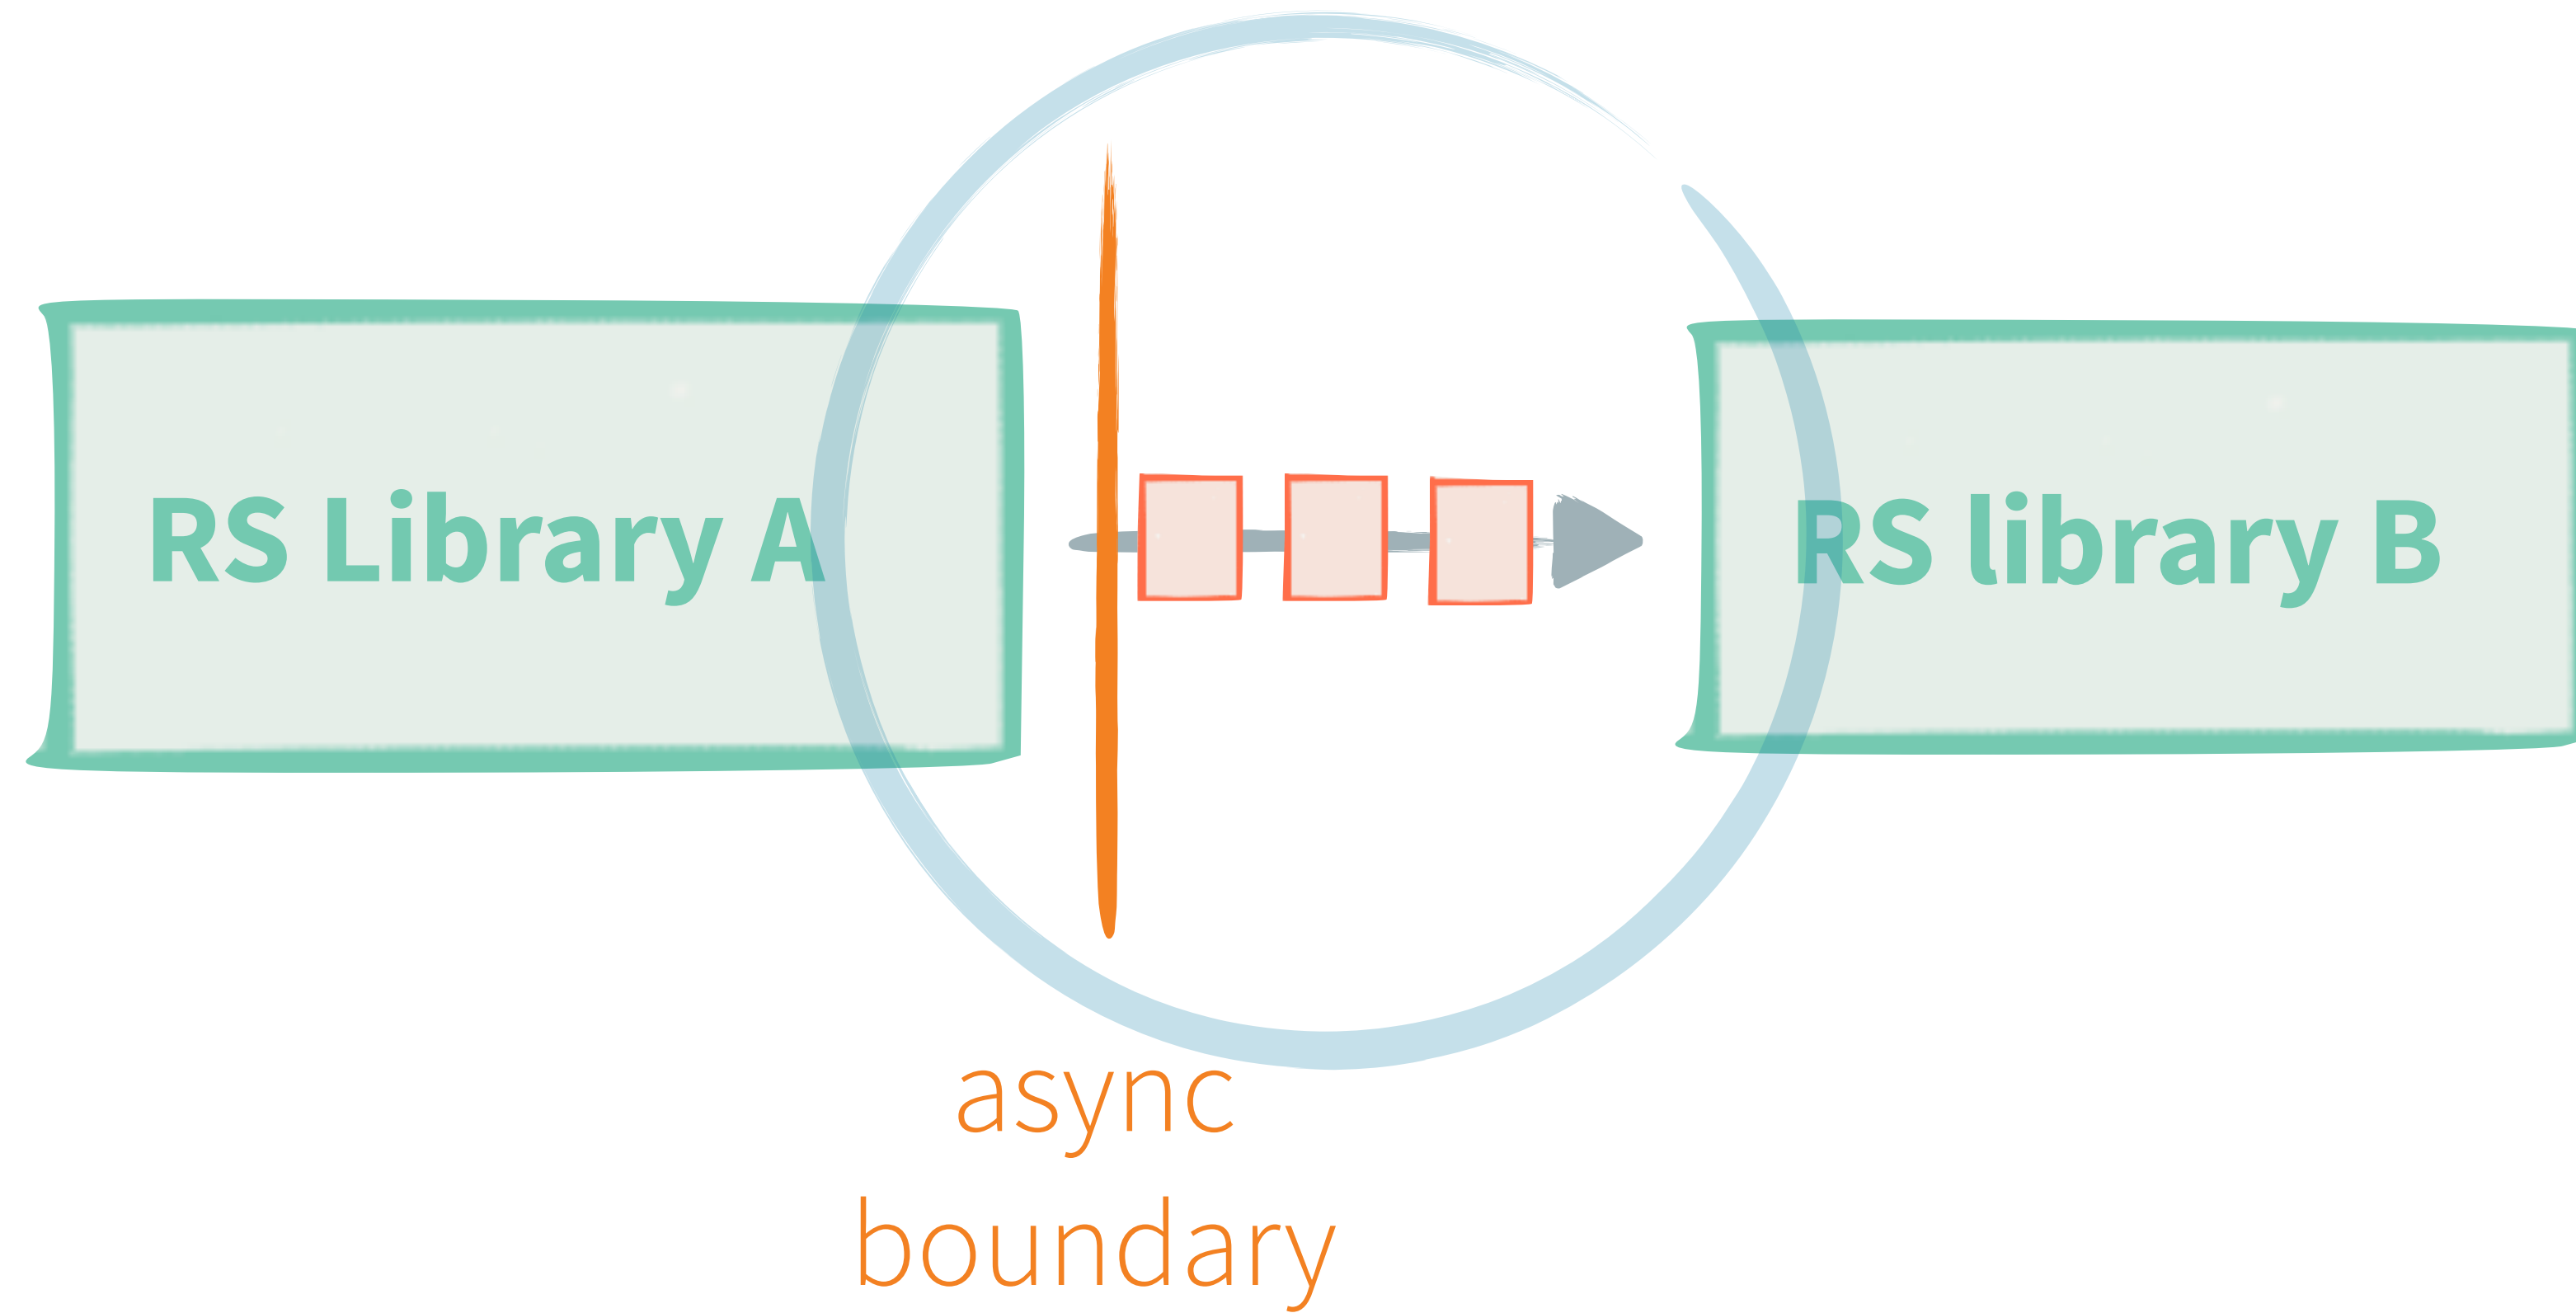
No back pressure - bounded buffer



Async non blocking back pressure



Reactive Streams



Reactive Streams

2. Subscriber (Code)

```
public interface Subscriber<T> {  
    public void onSubscribe(Subscription s);  
    public void onNext(T t);  
    public void onError(Throwable t);  
    public void onComplete();  
}
```

ID	Rule
1	A Subscriber MUST signal demand via <code>Subscription.request(long n)</code> to receive <code>onNext</code> signals.
2	If a Subscriber suspects that its processing of signals will negatively impact its Publisher's responsivity, it is RECOMMENDED that it asynchronously dispatches its signals.
3	<code>Subscriber.onComplete()</code> and <code>Subscriber.onError(Throwable t)</code> MUST NOT call any methods on the <code>Subscription</code> or the <code>Publisher</code> .
4	<code>Subscriber.onComplete()</code> and <code>Subscriber.onError(Throwable t)</code> MUST consider the <code>Subscription</code> cancelled after having received the signal.
5	A Subscriber MUST call <code>Subscription.cancel()</code> on the given <code>Subscription</code> after an <code>onSubscribe</code> signal if it already has an active <code>Subscription</code> .
6	A Subscriber MUST call <code>Subscription.cancel()</code> if it is no longer valid to the <code>Publisher</code> without the <code>Publisher</code> having signaled <code>onError</code> or <code>onComplete</code> .
7	A Subscriber MUST ensure that all calls on its <code>Subscription</code> take place from the same thread or provide for respective external synchronization.
8	A Subscriber MUST be prepared to receive one or more <code>onNext</code> signals after having called <code>Subscription.cancel()</code> if there are still requested elements pending [see 3.12]. <code>Subscription.cancel()</code> does not guarantee to perform the underlying cleaning operations immediately.
9	A Subscriber MUST be prepared to receive an <code>onComplete</code> signal with or without a preceding <code>Subscription.request(long n)</code> call.
10	A Subscriber MUST be prepared to receive an <code>onError</code> signal with or without a preceding <code>Subscription.request(long n)</code> call.
11	A Subscriber MUST make sure that all calls on its <code>onXXX</code> methods happen-before [1] the processing of the respective signals. I.e. the Subscriber must take care of properly publishing the signal to its processing logic.
12	<code>Subscriber.unsubscribe</code> MUST be called at most once for a given <code>Subscriber</code> (based on object equality).
13	Calling <code>onSubscribe</code> , <code>onNext</code> , <code>onError</code> or <code>onComplete</code> MUST return normally except when any provided parameter is <code>null</code> in which case it MUST throw a <code>java.lang.NullPointerException</code> to the caller, for all other situations the only legal way for a <code>Subscriber</code> to signal failure is by cancelling its <code>Subscription</code> . In the case that this rule is violated, any associated <code>Subscription</code> to the <code>Subscriber</code> MUST be considered as cancelled, and the caller MUST raise this error condition in a fashion that is adequate for the runtime environment.

[1] : See JMM definition of Happen-Before in section 17.4.5. on <http://docs.oracle.com/javase/specs/jls/se7/html/jls-17.html>

3. Publisher (Code)

```
public interface Publisher<T> {  
    void subscribe(Subscriber<? super T> s);  
}
```

ID	Rule
1	A number of <code>onNext</code> signals sent by a <code>Publisher</code> to a <code>Subscriber</code> MUST be greater than or equal to the number of elements requested by that <code>Subscriber</code> 's <code>Subscription</code> at all times.
2	A <code>Publisher</code> MAY signal less <code>onNext</code> than requested and terminate the <code>Subscription</code> with <code>onComplete</code> or <code>onError</code> .
3	When a <code>Subscriber</code> subscribes, <code>onNext</code> , <code>onError</code> and <code>onComplete</code> signaled to a <code>Subscriber</code> MUST be processed in the order of current notifications.
4	If a <code>Publisher</code> fails it MUST signal an <code>onError</code> .
5	When a <code>Publisher</code> terminates successfully (finite stream) it MUST signal an <code>onComplete</code> .
6	When a <code>Publisher</code> signals either <code>onError</code> or <code>onComplete</code> on a <code>Subscriber</code> , that <code>Subscriber</code> is considered cancelled.
7	When a <code>Subscriber</code> reaches its terminal state has been signaled (<code>onError</code> , <code>onComplete</code>) it is REQUIRED that the <code>Subscription</code> is cancelled its <code>Subscriber</code> MUST eventually stop being signaled.
8	A <code>Publisher.subscribe</code> MUST call <code>onSubscribe</code> on the provided <code>Subscriber</code> prior to signaling any <code>onNext</code> and MUST return normally, except when the provided <code>Subscriber</code> is <code>null</code> in which case it MUST throw a <code>java.lang.NullPointerException</code> to the caller, for all other situations the only legal way for a <code>Publisher</code> to signal failure (or reject the <code>Subscriber</code>) is by calling <code>onError</code> (after calling <code>onSubscribe</code>).
9	A <code>Publisher.subscribe</code> MAY be called as many times as needed.
10	A <code>Publisher</code> MAY support multiple <code>Subscription</code> s.

3. Subscription (Code)

```
public interface Subscription {  
    public void request(long n);  
    public void cancel();  
}
```

ID	Rule
1	<code>Subscription.request</code> and <code>Subscription.cancel</code> MUST only be called inside of its <code>Subscriber</code> context. A <code>Subscription</code> represents the unique relationship between a <code>Subscriber</code> and a <code>Publisher</code> [see 2.12].
2	The <code>Subscription</code> MUST allow the <code>Subscriber</code> to call <code>Subscription.request</code> synchronously from within <code>onNext</code> or <code>onSubscribe</code> .
3	<code>Subscription.request</code> MUST place an upper bound on possible synchronous recursion between <code>Publisher</code> and <code>Subscriber</code> [1].
4	<code>Subscription.request</code> SHOULD respect the responsivity of its caller by returning in a timely manner[2].
5	<code>Subscription.cancel</code> MUST respect the responsivity of its caller by returning in a timely manner[2], MUST be idempotent and MUST be thread-safe.
6	After the <code>Subscription</code> is cancelled, additional <code>Subscription.request(long n)</code> MUST be NOPs.
7	After the <code>Subscription</code> is cancelled, additional <code>Subscription.cancel()</code> MUST be NOPs.
8	While the <code>Subscription</code> is not cancelled, <code>Subscription.request(long n)</code> MUST register the given number of additional elements to be produced to the respective subscriber.
9	While the <code>Subscription</code> is not cancelled, <code>Subscription.request(long n)</code> MUST signal <code>onError</code> with a <code>java.lang.IllegalArgumentException</code> if the argument is ≤ 0 . The cause message MUST include a reference to this rule and/or quote the full rule.
10	While the <code>Subscription</code> is not cancelled, <code>Subscription.request(long n)</code> MAY synchronously call <code>onNext</code> on this (or other) subscriber(s).
11	While the <code>Subscription</code> is not cancelled, <code>Subscription.request(long n)</code> MAY synchronously call <code>onComplete</code> or <code>onError</code> on this (or other) subscriber(s).
12	While the <code>Subscription</code> is not cancelled, the <code>Publisher</code> to the <code>Subscription</code> MUST not signal <code>onError</code> or <code>onComplete</code> until the <code>Subscription</code> is cancelled.
13	While the <code>Subscription</code> is not cancelled, the <code>Publisher</code> to the same <code>Subscription</code> MUST not signal <code>onError</code> or <code>onComplete</code> that is disallowed.
14	While the <code>Subscription</code> is not cancelled, the <code>Publisher</code> , if any other <code>Subscription</code> exists at this point [see 1.9], MUST not signal <code>onError</code> or <code>onComplete</code> .
15	Calling <code>Subscription.cancel</code> MUST return normally. The only legal way to signal failure to a <code>Subscriber</code> is via the <code>onError</code> method.
16	Calling <code>Subscription.request</code> MUST return normally. The only legal way to signal failure to a <code>Subscriber</code> is via the <code>onError</code> method.
17	A <code>Subscription</code> MUST support an unbounded number of calls to <code>request</code> and MUST support a demand


“Make building powerful concurrent & distributed applications **simple**.”

Akka Streams

Complete and awesome
Java and Scala APIs
(Just like everything in Akka)



Akka Streams in ~20 seconds:



```
final ActorSystem system = ActorSystem.create();
final Materializer materializer = ActorMaterializer.create(system);

final Source<Integer, NotUsed> source =
    Source.range(0, 200000000);

final Flow<Integer, String, NotUsed> flow =
    Flow.fromFunction((Integer n) -> n.toString());


final Sink<String, CompletionStage<Done>> sink =
    Sink.foreach(str -> System.out.println(str));

final RunnableGraph<NotUsed> runnable = source.via(flow).to(sink);

runnable.run(materializer);
```



Akka Streams in ~20 seconds:



```
final ActorSystem system = ActorSystem.create();
final Materializer materializer = ActorMaterializer.create(system);

final Source<Integer, NotUsed> source =
    Source.range(0, 20000000);

final Flow<Integer, String, NotUsed> flow =
    Flow.fromFunction((Integer n) -> n.toString());

final Sink<String, CompletionStage<Done>> sink =
    Sink.foreach(str -> System.out.println(str));

final RunnableGraph<NotUsed> runnable = source.via(flow).to(sink);

runnable.run(materializer);
```



Akka Streams in ~20 seconds:



```
implicit val system = ActorSystem()
implicit val mat = ActorMaterializer()

val source = Source(0 to 200000000)

val flow = Flow[Int].map(_.toString())


val sink = Sink.foreach[String](println(_))

val runnable = source.via(flow).to(sink)

runnable.run()
```



Akka Streams in ~20 seconds:



```
Source.range(0, 200000000)
  .map(Object::toString)
  .runForeach(str -> System.out.println(str), materializer);
```




Akka Streams in ~20 seconds:



```
Source(0 to 200000000)  
  .map(_.toString)  
  .runForeach(println)
```




Numbers as a service



```
final Source<ByteString, NotUsed> numbers = Source.unfold(0L, n -> {  
    long next = n + 1;  
    return Optional.of(Pair.create(next, next));  
}).map(n -> ByteString.fromString(n.toString() + "\n"));
```

```
final Route route =  
    path("numbers", () ->  
        get(() ->  
            complete(HttpResponse.create()  
                .withStatus(StatusCodes.OK)  
                .withEntity(HttpEntities.create(  
                    ContentTypes.TEXT_PLAIN_UTF8,  
                    numbers  
                )))  
        )  
    );
```

```
final CompletionStage<ServerBinding> bindingCompletionStage =  
    http.bindAndHandle(route.flow(system, materializer), host, materializer);
```



Numbers as a service






```
val numbers =
  Source.unfold(0L) { (n) =>
    val next = n + 1
    Some((next, next))
  }.map(n => ByteString(n + "\n"))

val route =
  path("numbers") {
    get {
      complete(
        HttpResponse(entity = HttpEntity(`text/plain(UTF-8)`, numbers))
      )
    }
  }
val futureBinding = Http().bindAndHandle(route, "127.0.0.1", 8080)
```

Back pressure over TCP

Client

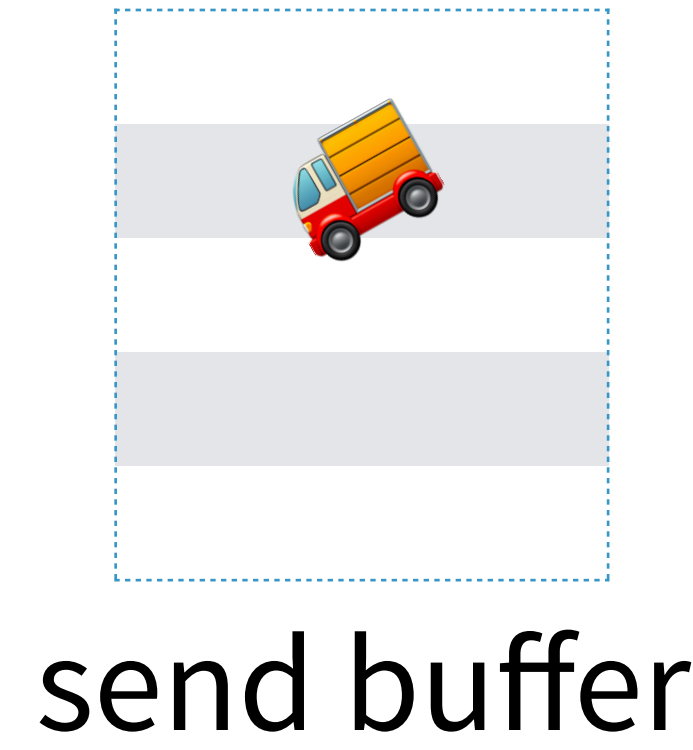



Server



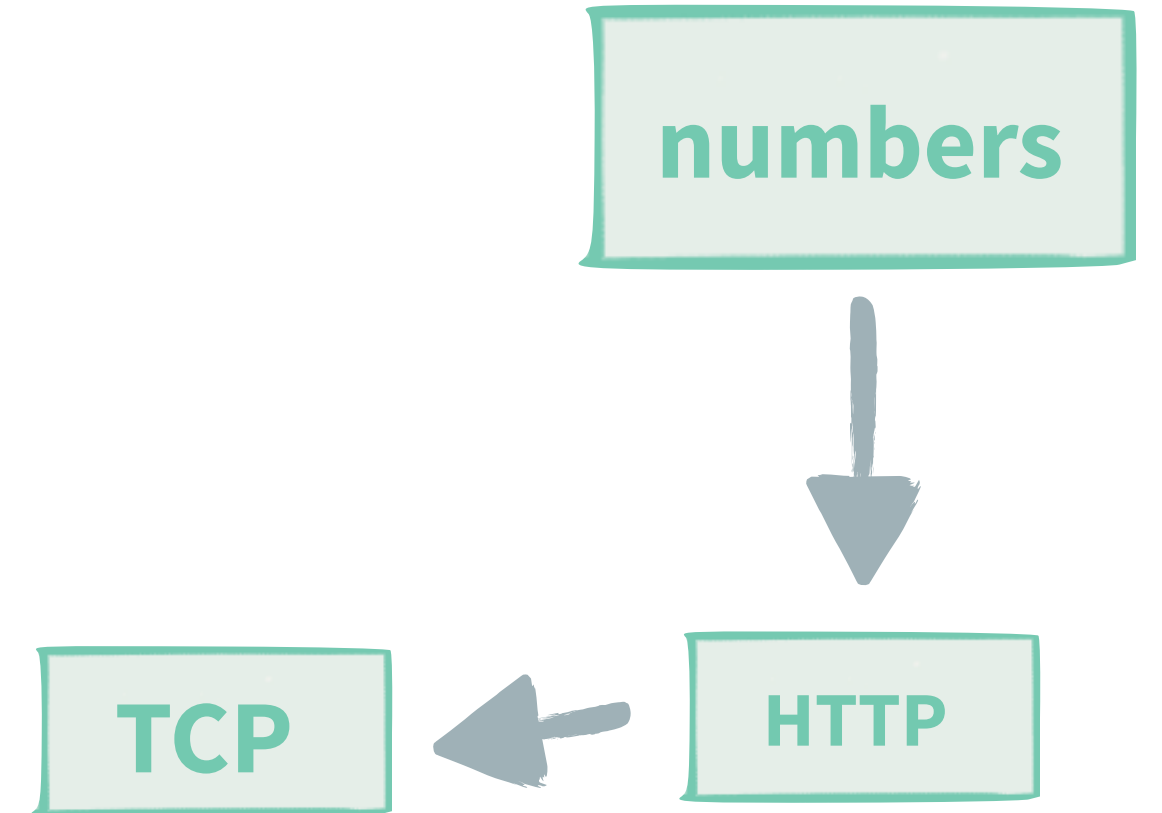
Back pressure over TCP

Client




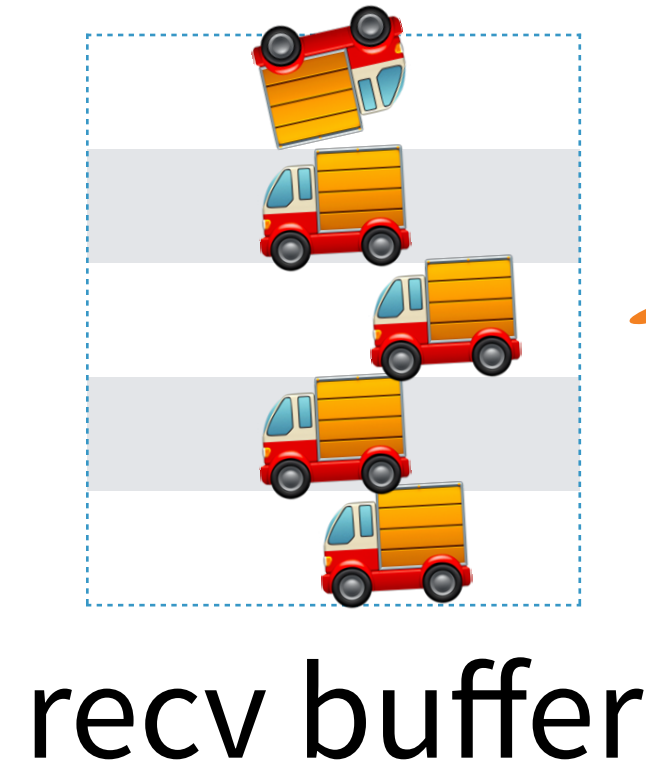
Backpressure

Server



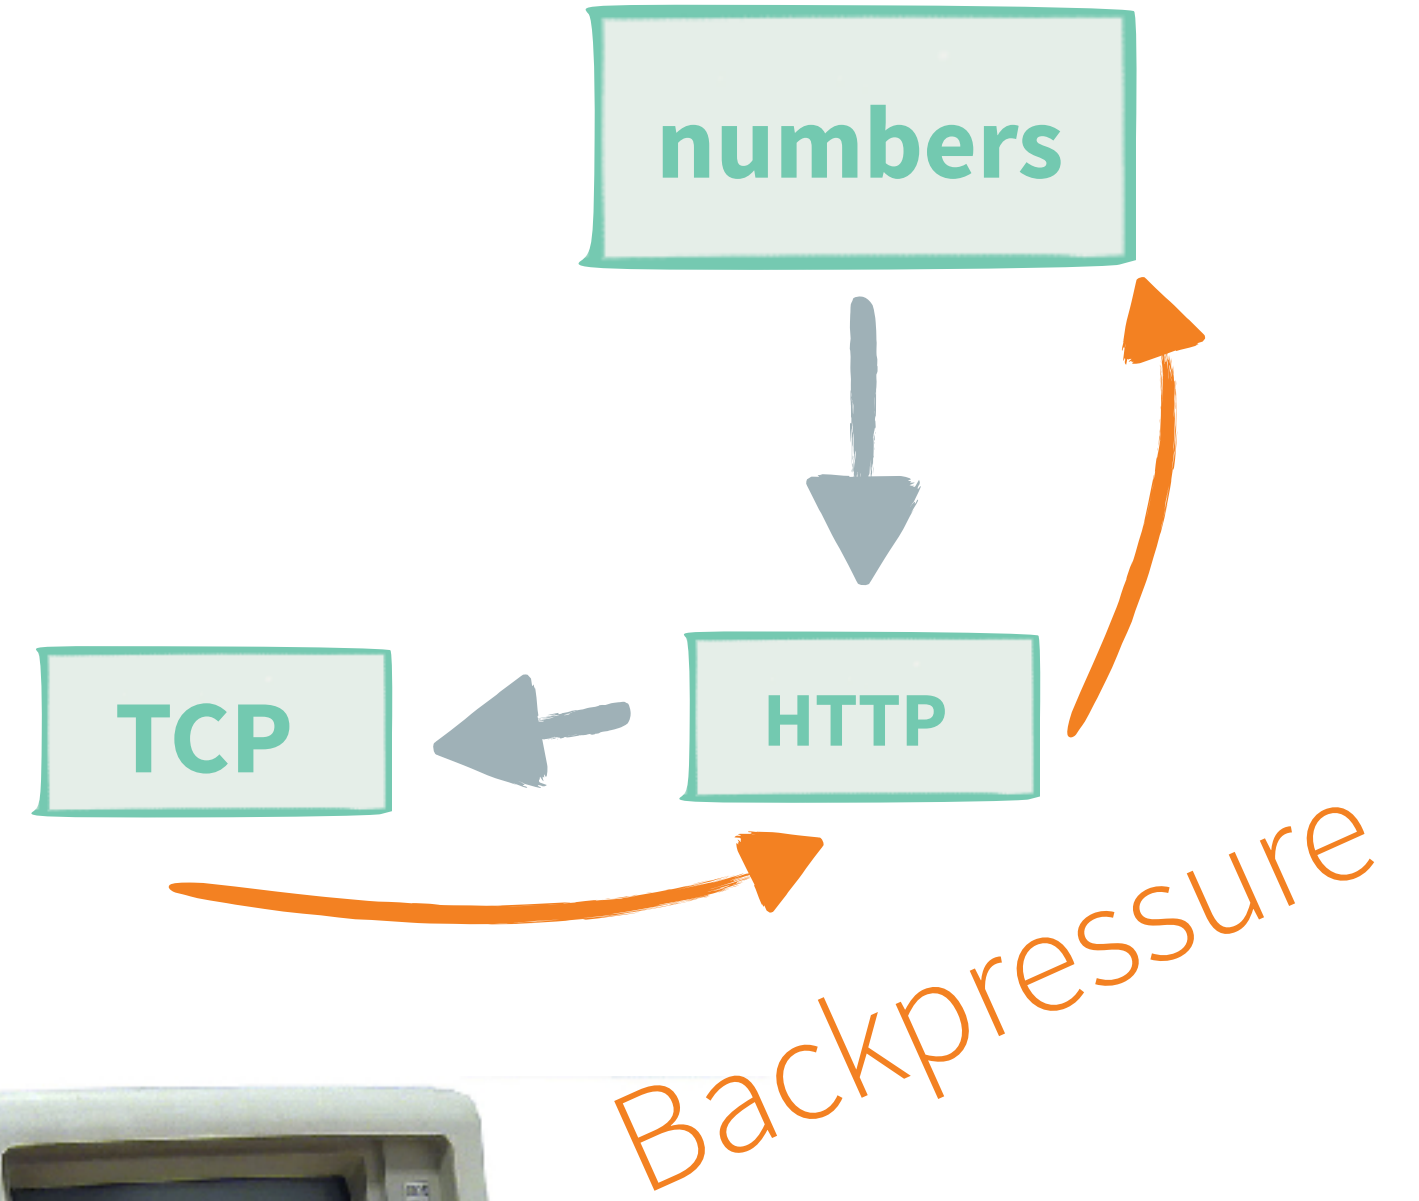
Back pressure over TCP

Client



Backpressure


Server





A more useful example

Credit to: Colin Breck



```
final Flow<Message, Message, NotUsed> measurementsFlow =  
    Flow.of(Message.class)  
        .flatMapConcat((Message message) ->  
            message.asTextMessage()  
                .getStreamedText()  
                .fold("", (acc, elem) -> acc + elem)  
        )  
        .groupedWithin(1000, FiniteDuration.create(1, SECONDS))  
        .mapAsync(5, database::asyncBulkInsert)  
        .map(written ->  
            TextMessage.create("wrote up to: " + written.get(written.size() - 1))  
        );  
  
final Route route = path("measurements", () ->  
    get(() ->  
        handleWebSocketMessages(measurementsFlow)  
    )  
);  
  
final CompletionStage<ServerBinding> bindingCompletionStage =  
    http.bindAndHandle(route.flow(system, materializer), host, materializer);
```



A more useful example





```
val measurementsFlow =
  Flow[Message].flatMapConcat(message =>
    message.asTextMessage.getStreamedText.fold("")( _ + _ )
  )
  .groupedWithin(1000, 1.second)
  .mapAsync(5)(Database.asyncBulkInsert)
  .map(written => TextMessage("wrote up to: " + written.last))

val route =
  path("measurements") {
    get {
      handleWebSocketMessages(measurementsFlow)
    }
  }

val futureBinding = Http().bindAndHandle(route, "127.0.0.1", 8080)
```

The tale of the two pancake chefs

Scoops of batter



Pancakes



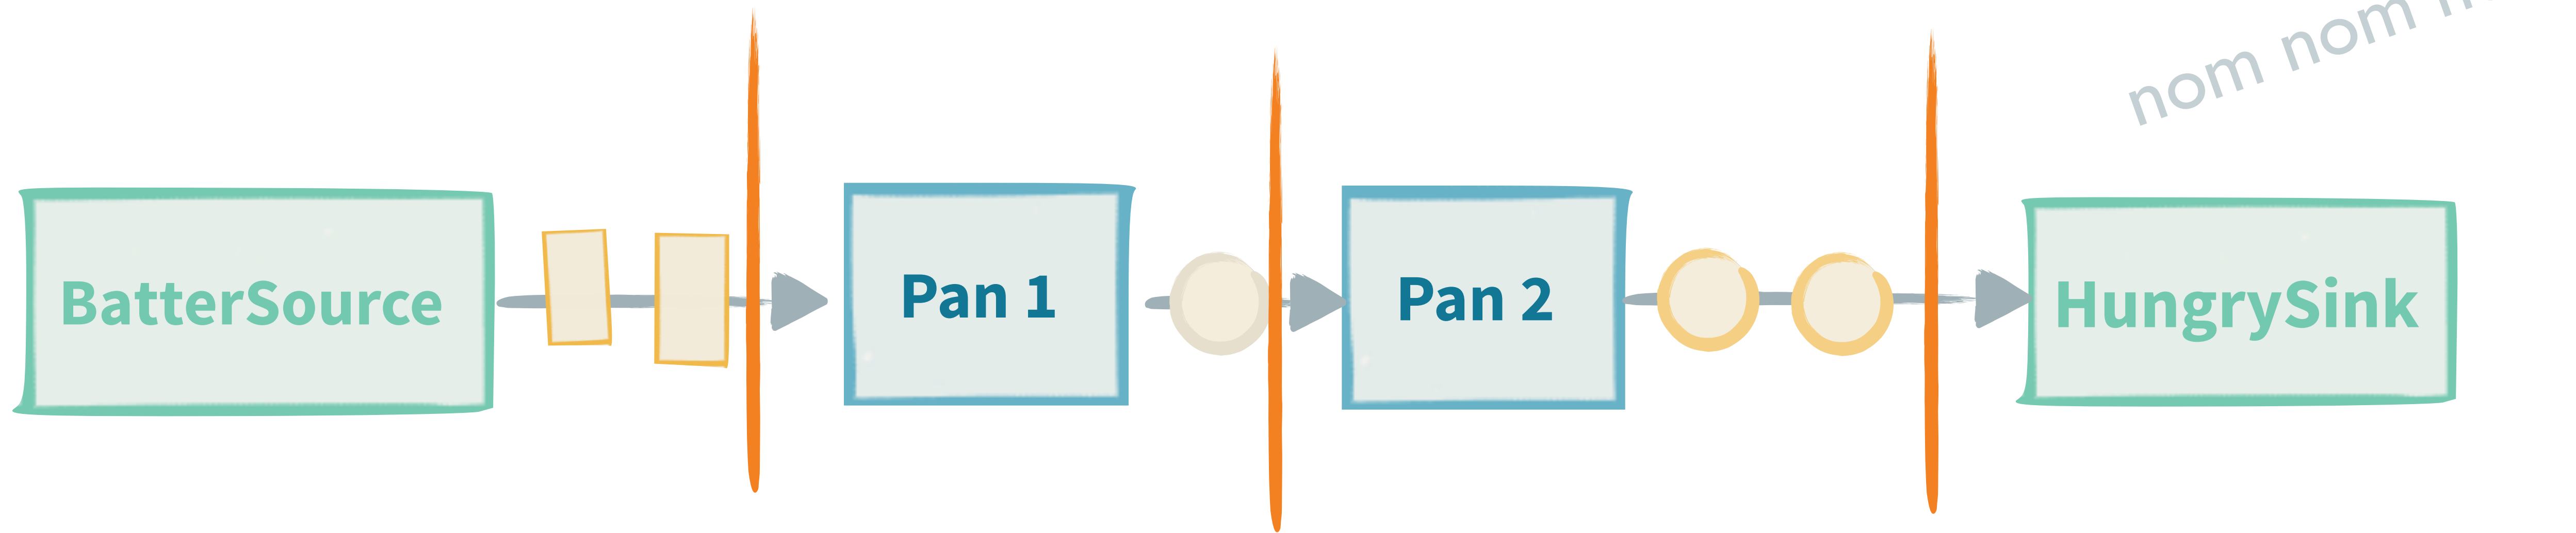
asynchronous
boundaries

nom nom nom






Rolands pipelined pancakes





Rolands pipelined pancakes



```
Flow<ScoopOfBatter, HalfCookedPancake, NotUsed> fryingPan1 =  
    Flow.of(ScoopOfBatter.class).map(batter -> new HalfCookedPancake());
```

```
Flow<HalfCookedPancake, Pancake, NotUsed> fryingPan2 =  
    Flow.of(HalfCookedPancake.class).map(halfCooked -> new Pancake());
```

```
Flow<ScoopOfBatter, Pancake, NotUsed> pancakeChef =  
    fryingPan1.async().via(fryingPan2.async());
```



Rolands pipelined pancakes




```
// Takes a scoop of batter and creates a pancake with one side cooked
val fryingPan1: Flow[ScoopOfBatter, HalfCookedPancake, NotUsed] =
    Flow[ScoopOfBatter].map { batter => HalfCookedPancake() }

// Finishes a half-cooked pancake
val fryingPan2: Flow[HalfCookedPancake, Pancake, NotUsed] =
    Flow[HalfCookedPancake].map { halfCooked => Pancake() }

// With the two frying pans we can fully cook pancakes
val pancakeChef: Flow[ScoopOfBatter, Pancake, NotUsed] =
    Flow[ScoopOfBatter].via(fryingPan1.async).via(fryingPan2.async)
```


Patriks parallel pancakes



nom nom nom



Patriks parallel pancakes



```
Flow<ScoopOfBatter, Pancake, NotUsed> fryingPan =  
    Flow.of(ScoopOfBatter.class).map(batter -> new Pancake());  
  
Flow<ScoopOfBatter, Pancake, NotUsed> pancakeChef =  
    Flow.fromGraph(GraphDSL.create(builder -> {  
        final UniformFanInShape<Pancake, Pancake> mergePancakes =  
            builder.add(Merge.create(2));  
        final UniformFanOutShape<ScoopOfBatter, ScoopOfBatter> dispatchBatter =  
            builder.add(Balance.create(2));  
  
        builder.from(dispatchBatter.out(0))  
            .via(builder.add(fryingPan.async()))  
            .toInlet(mergePancakes.in(0));  
  
        builder.from(dispatchBatter.out(1))  
            .via(builder.add(fryingPan.async()))  
            .toInlet(mergePancakes.in(1));  
  
        return FlowShape.of(dispatchBatter.in(), mergePancakes.out());  
    }));
```



Patriks parallel pancakes



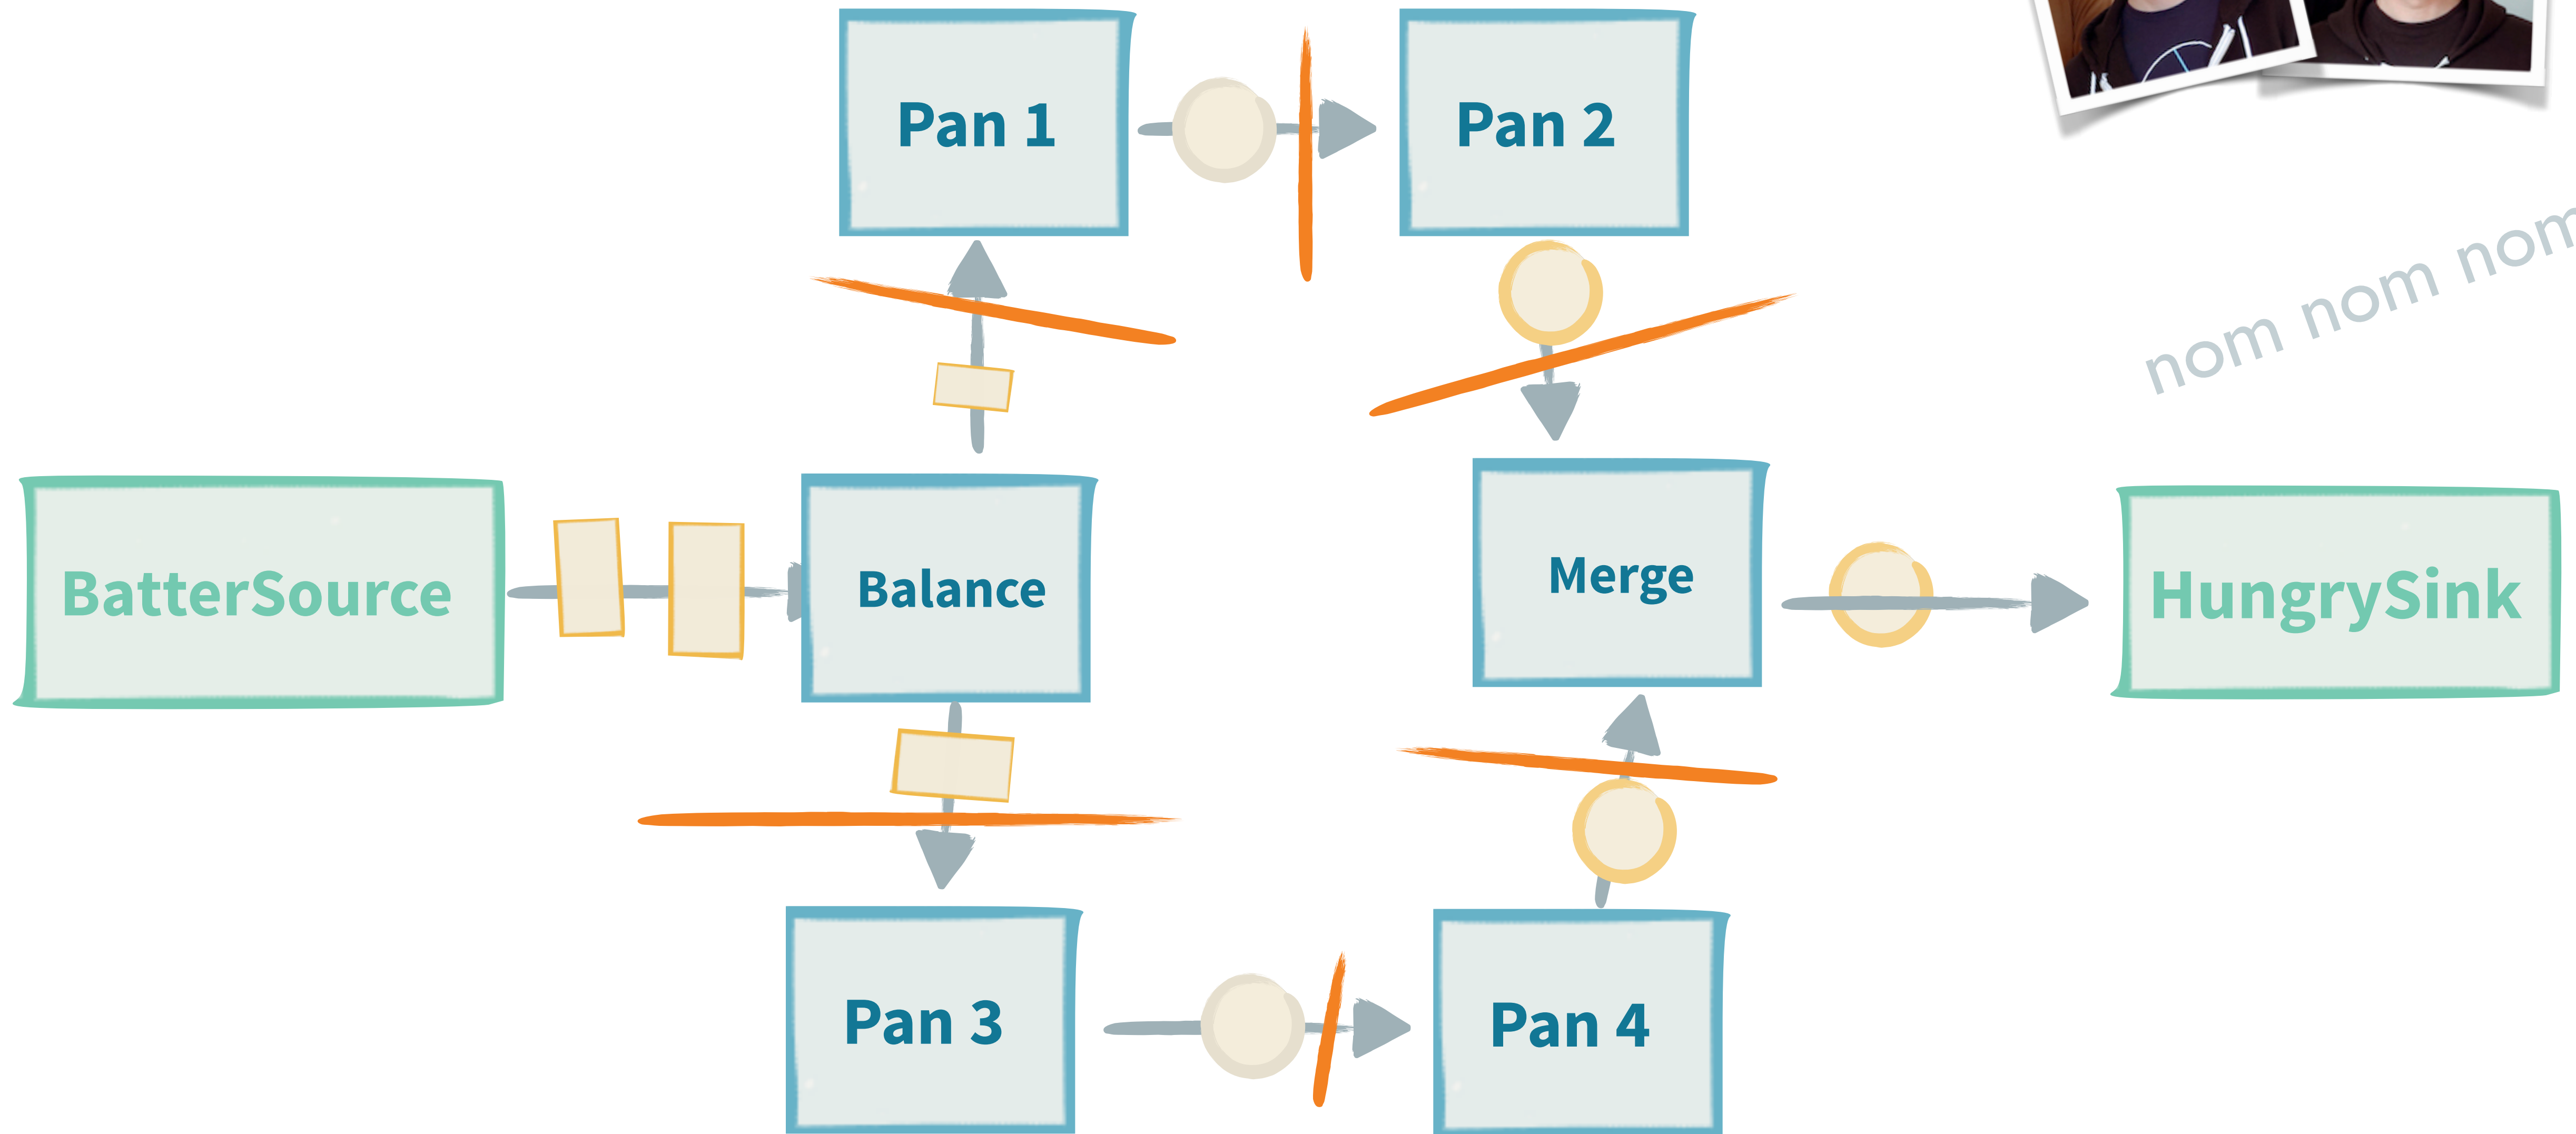
```
val pancakeChef: Flow[ScoopOfBatter, Pancake, NotUsed] =
  Flow.fromGraph(GraphDSL.create() { implicit builder =>
    import GraphDSL.Implicits._

    val dispatchBatter = builder.add(Balance[ScoopOfBatter](2))
    val mergePancakes = builder.add(Merge[Pancake](2))

    // Using two pipelines, having two frying pans each, in total using
    // four frying pans
    dispatchBatter.out(0) ~> fryingPan1.async ~> fryingPan2.async ~> mergePancakes.in(0)
    dispatchBatter.out(1) ~> fryingPan1.async ~> fryingPan2.async ~> mergePancakes.in(1)

    FlowShape(dispatchBatter.in, mergePancakes.out)
  })
```

Making pancakes together



Built in stages

Source stages

fromIterator, apply, single, repeat, cycle, tick, fromFuture, fromCompletionStage, unfold, unfoldAsync, empty, maybe, failed, lazily, actorPublisher, actorRef, combine, unfoldResource, unfoldResourceAsync, queue, asSubscriber, fromPublisher, zipN, zipWithN

Sink stages

head, headOption, last, lastOption, ignore, cancelled, seq, foreach, foreachParallel, onComplete, lazyInit, queue, fold, reduce, combine, actorRef, actorRefWithAck, actorSubscriber, asPublisher, fromSubscriber

Flow stages

map/fromFunction, mapConcat, statefulMapConcat, filter, filterNot, collect, grouped, sliding, scan, scanAsync, fold, foldAsync, reduce, drop, take, takeWhile, dropWhile, recover, recoverWith, recoverWithRetries, mapError, detach, throttle, intersperse, limit, limitWeighted, log, recoverWithRetries, mapAsync, mapAsyncUnordered, takeWithin, dropWithin, groupedWithin, initialDelay, delay, conflate, conflateWithSeed, batch, batchWeighted, expand, buffer, prefixAndTail, groupBy, splitWhen, splitAfter, flatMapConcat, flatMapMerge, initialTimeout, completionTimeout, idleTimeout, backpressureTimeout, keepAlive, initialDelay, merge, mergeSorted,

mergePreferred, zip, zipWith, zipWithIndex, concat, prepend, orElse, interleave, unzip, unzipWith, broadcast, balance, partition, watchTermination, monitor

File IO Sinks and Sources

fromPath, toPath

Additional Sink and Source converters

{from,as}OutputStream, {from,as}InputStream, {as,from}javaCollector, javaCollectorParallelUnordered

Even more

Framing, JSON framing, killswitch, BroadcastHub, MergeHub

**But I want to
connect other
things!**



@doggosdoingthings



Alpakka

A community for Akka Streams connectors

<http://github.com/akka/alpakka>

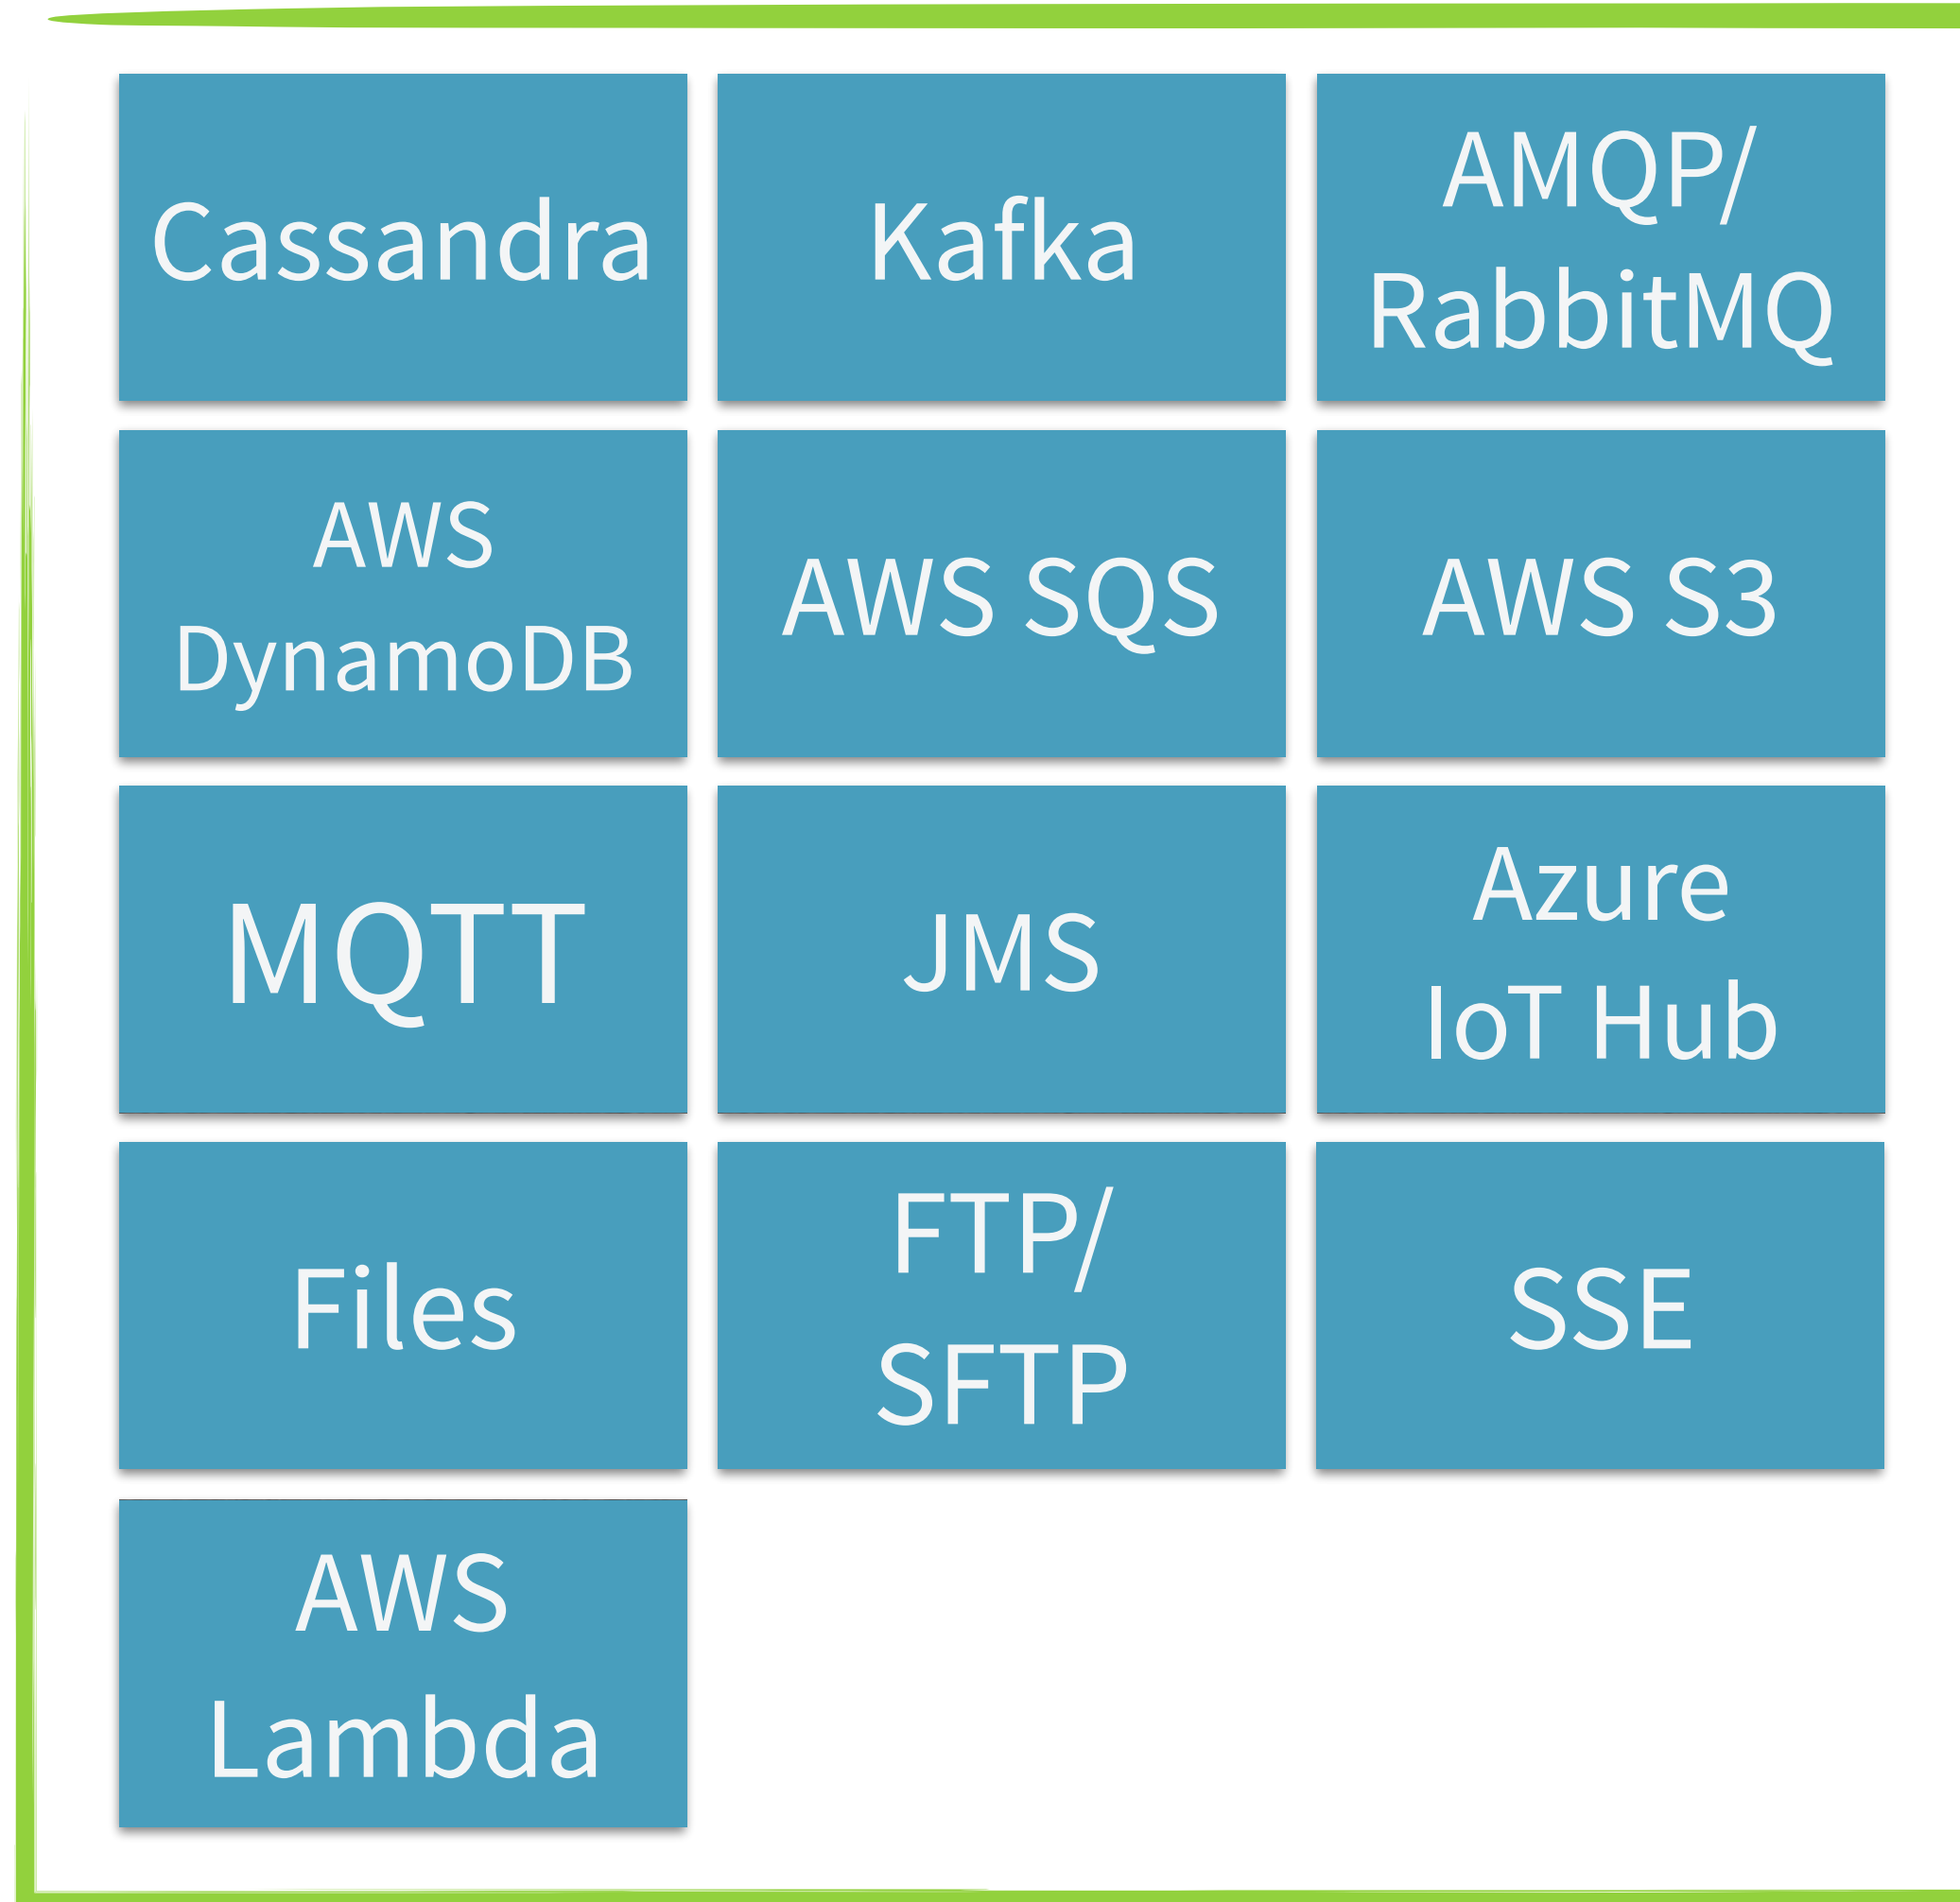


Alpakka – a community for Stream connectors

In Akka




Existing Alpakka



Alpakka PRs




**But my usecase is a
unique snowflake!**





GraphStage API



```
public class Map<A, B> extends GraphStage<FlowShape<A, B>> {
  private final Function<A, B> f;
  public final Inlet<A> in = Inlet.create("Map.in");
  public final Outlet<B> out = Outlet.create("Map.out");
  private final FlowShape<A, B> shape = FlowShape.of(in, out);
  public Map(Function<A, B> f) {
    this.f = f;
  }
  public FlowShape<A,B> shape() {
    return shape;
  }
  public GraphStageLogic createLogic(Attributes inheritedAttributes) {
    return new GraphStageLogic(shape) {
      {
        setHandler(in, new AbstractInHandler() {
          @Override
          public void onPush() throws Exception {
            push(out, f.apply(grab(in)));
          }
        });
        setHandler(out, new AbstractOutHandler() {
          @Override
          public void onPull() throws Exception {
            pull(in);
          }
        });
      }
    });
  }
}
```




GraphStage API



```
class Map[A, B](f: A => B) extends GraphStage[FlowShape[A, B]] {  
  
  val in = Inlet[A]("Map.in")  
  val out = Outlet[B]("Map.out")  
  override val shape = FlowShape.of(in, out)  
  
  override def createLogic(attr: Attributes): GraphStageLogic =  
    new GraphStageLogic(shape) {  
      setHandler(in, new InHandler {  
        override def onPush(): Unit = {  
          push(out, f(grab(in)))  
        }  
      })  
      setHandler(out, new OutHandler {  
        override def onPull(): Unit = {  
          pull(in)  
        }  
      })  
    }  
}
```

What about distributed/reactive systems?



The community

~200 active contributors!

Mailing list:

<https://groups.google.com/group/akka-user>

Public chat rooms:

<http://gitter.im/akka/dev> developing Akka

<http://gitter.im/akka/akka> using Akka

Easy to contribute tickets:

<https://github.com/akka/akka/issues?q=is%3Aissue+is%3Aopen+label%3Aeasy-to-contribute>

<https://github.com/akka/akka/issues?q=is%3Aissue+is%3Aopen+label%3A%22nice-to-have+%28low-prio%29%22>



Akka

<http://akka.io>

Thanks for listening!

Runnable sample sources (Java & Scala)

<https://github.com/johanandren/akka-stream-samples/tree/jfokus-2017>



@apnylle

johan.andren@lightbend.com

