

Self Healing in Streaming Systems

#UW Database Day

Dec 2nd, 2016

Karthik Ramasamy

Twitter

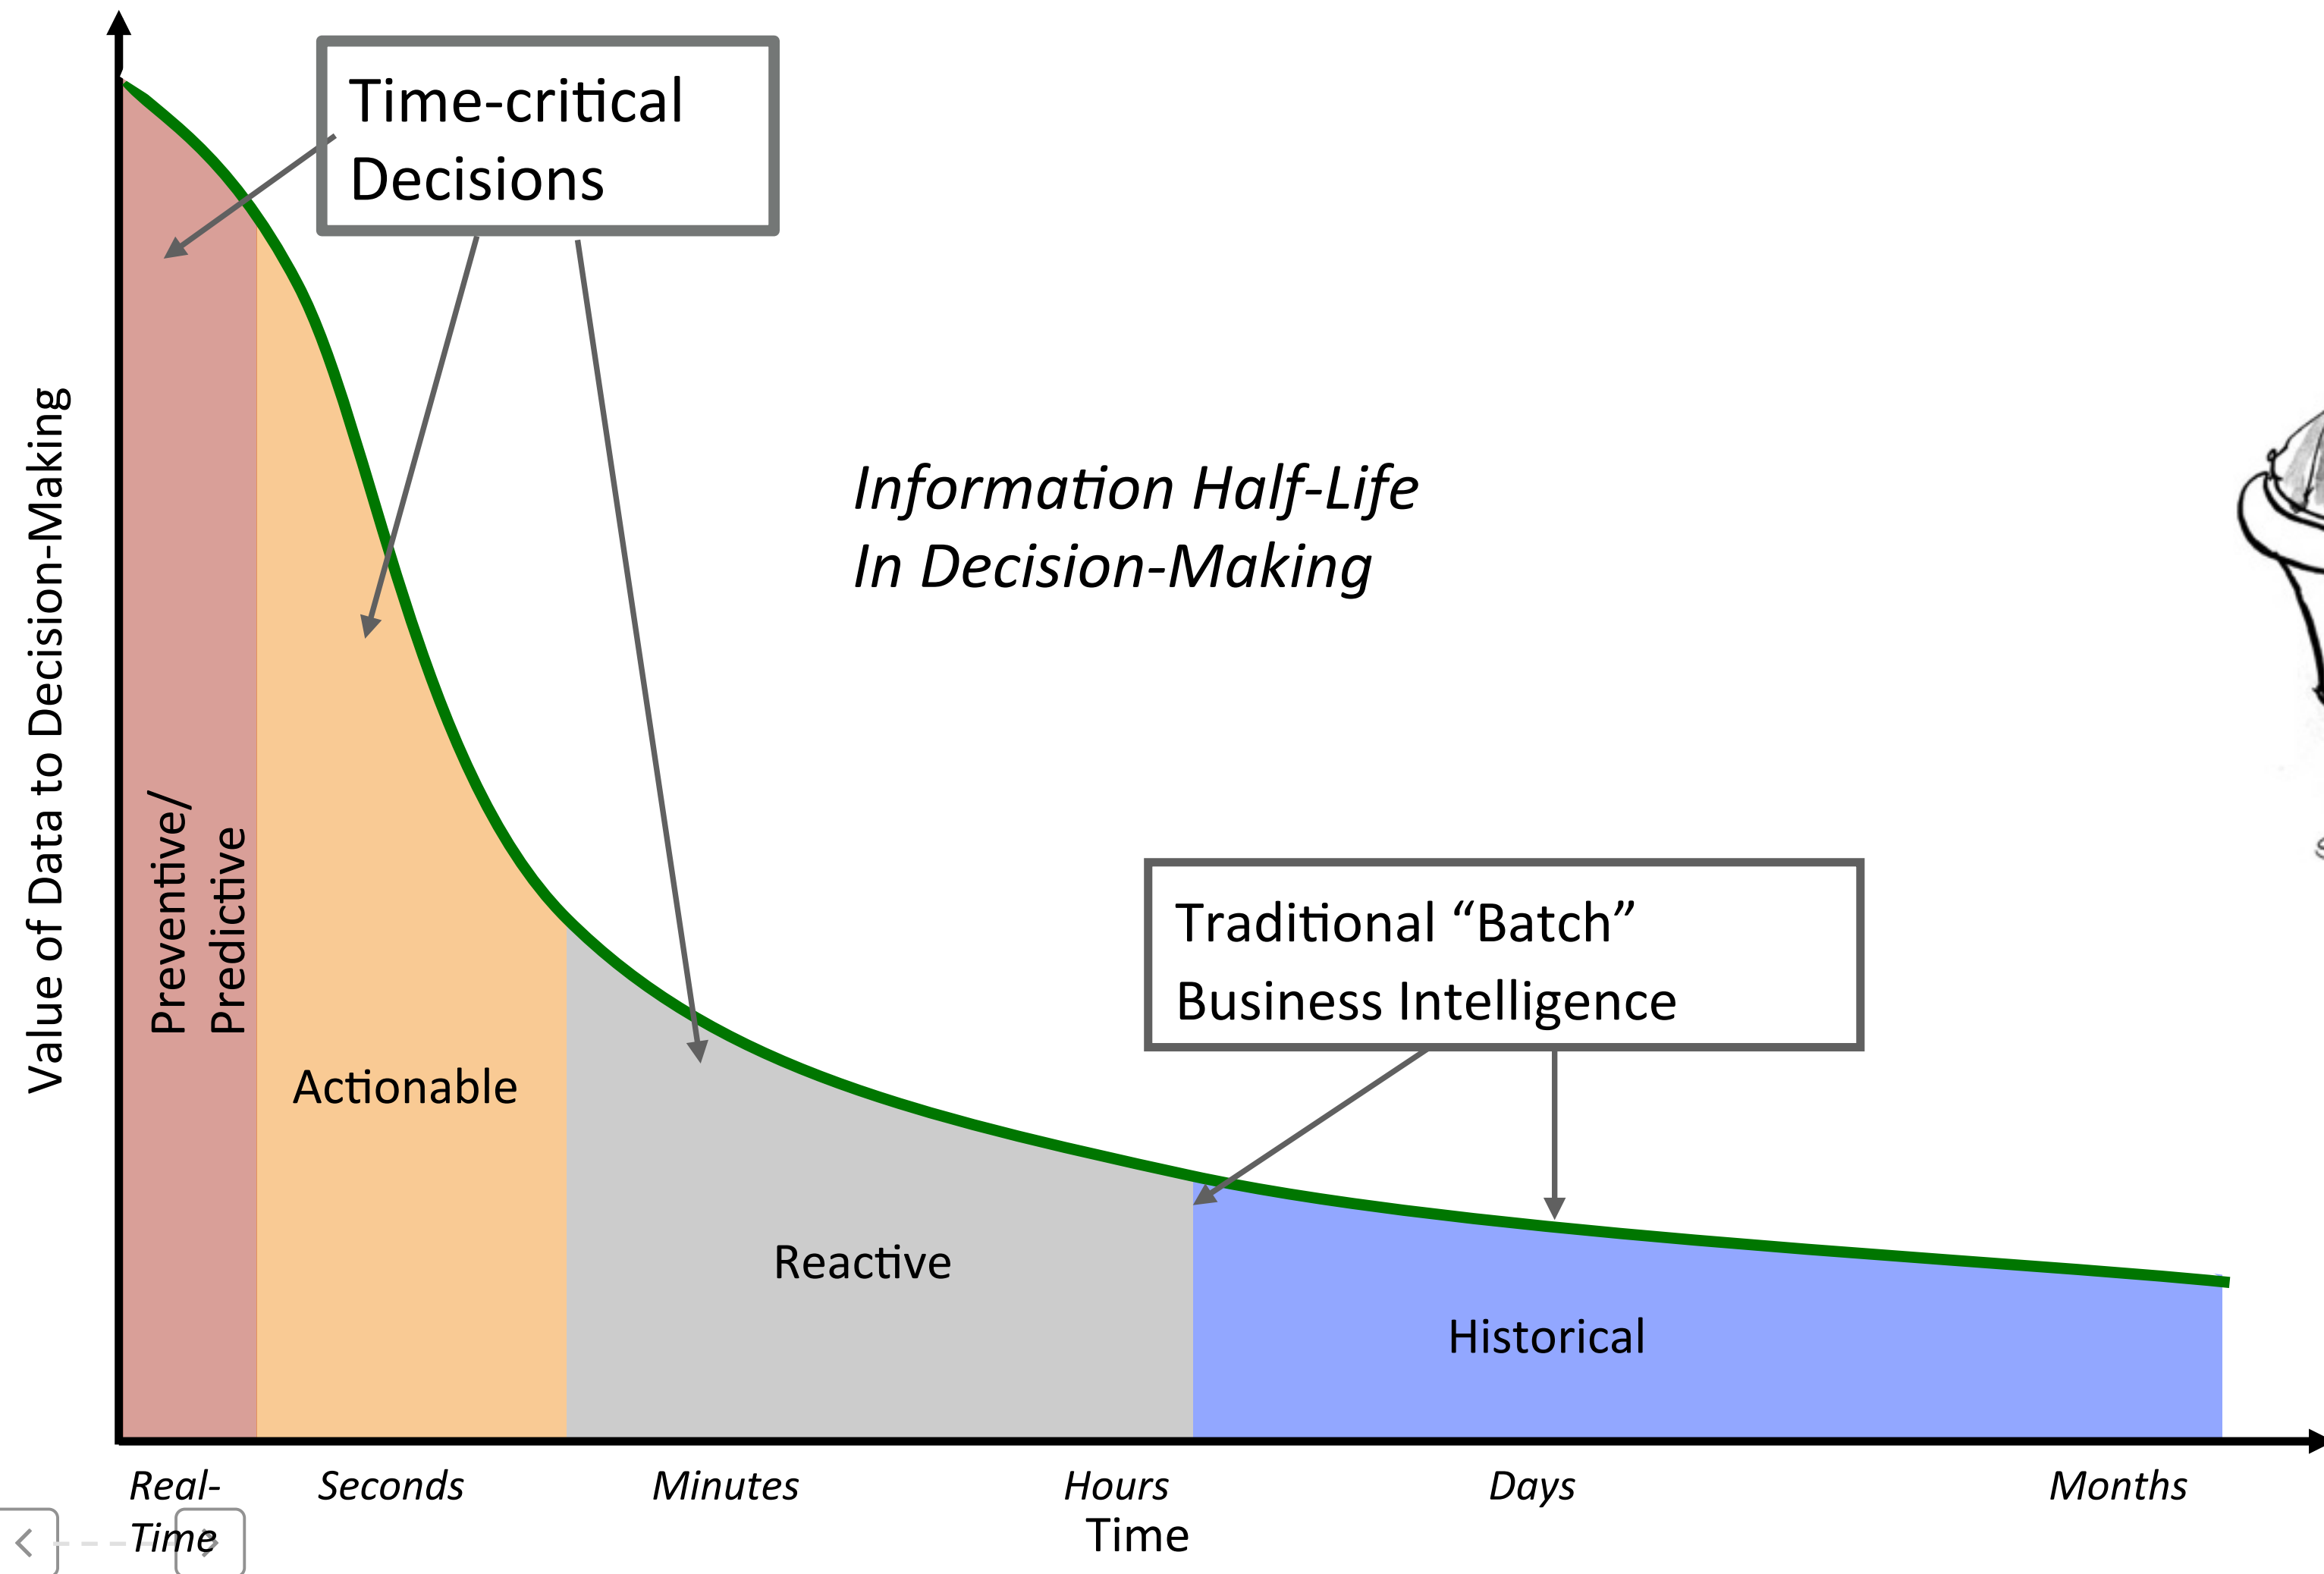
@karthikz



What is self healing?

A self healing system adapts itself as their environmental conditions change and continue to produce results

Why?



"Hold on, shoppers, while I bring that request to management."

Why?

LOSS OF REVENUE



Impact of downtime
popular event such as
Super Bowl
Oscars, etc

SLA VIOLATIONS



Impact of not honoring
an SLA leading to
penalty payments

QUALITY OF LIFE



Engineers & SRE burnt
out attending to
incidents

INCREASED PRODUCTIVITY




With reduced incidents,
engineers can focus on
actual development

Real-Time Budget

> 1 HOUR

high throughput



ad hoc queries

monthly active users relevance for ads

10 MS - 1 SEC

approximate



ad impressions count
hash tag trends

< 500 MS

latency sensitive



deterministic workflows

fanout Tweets search for Tweets

< 1 MS

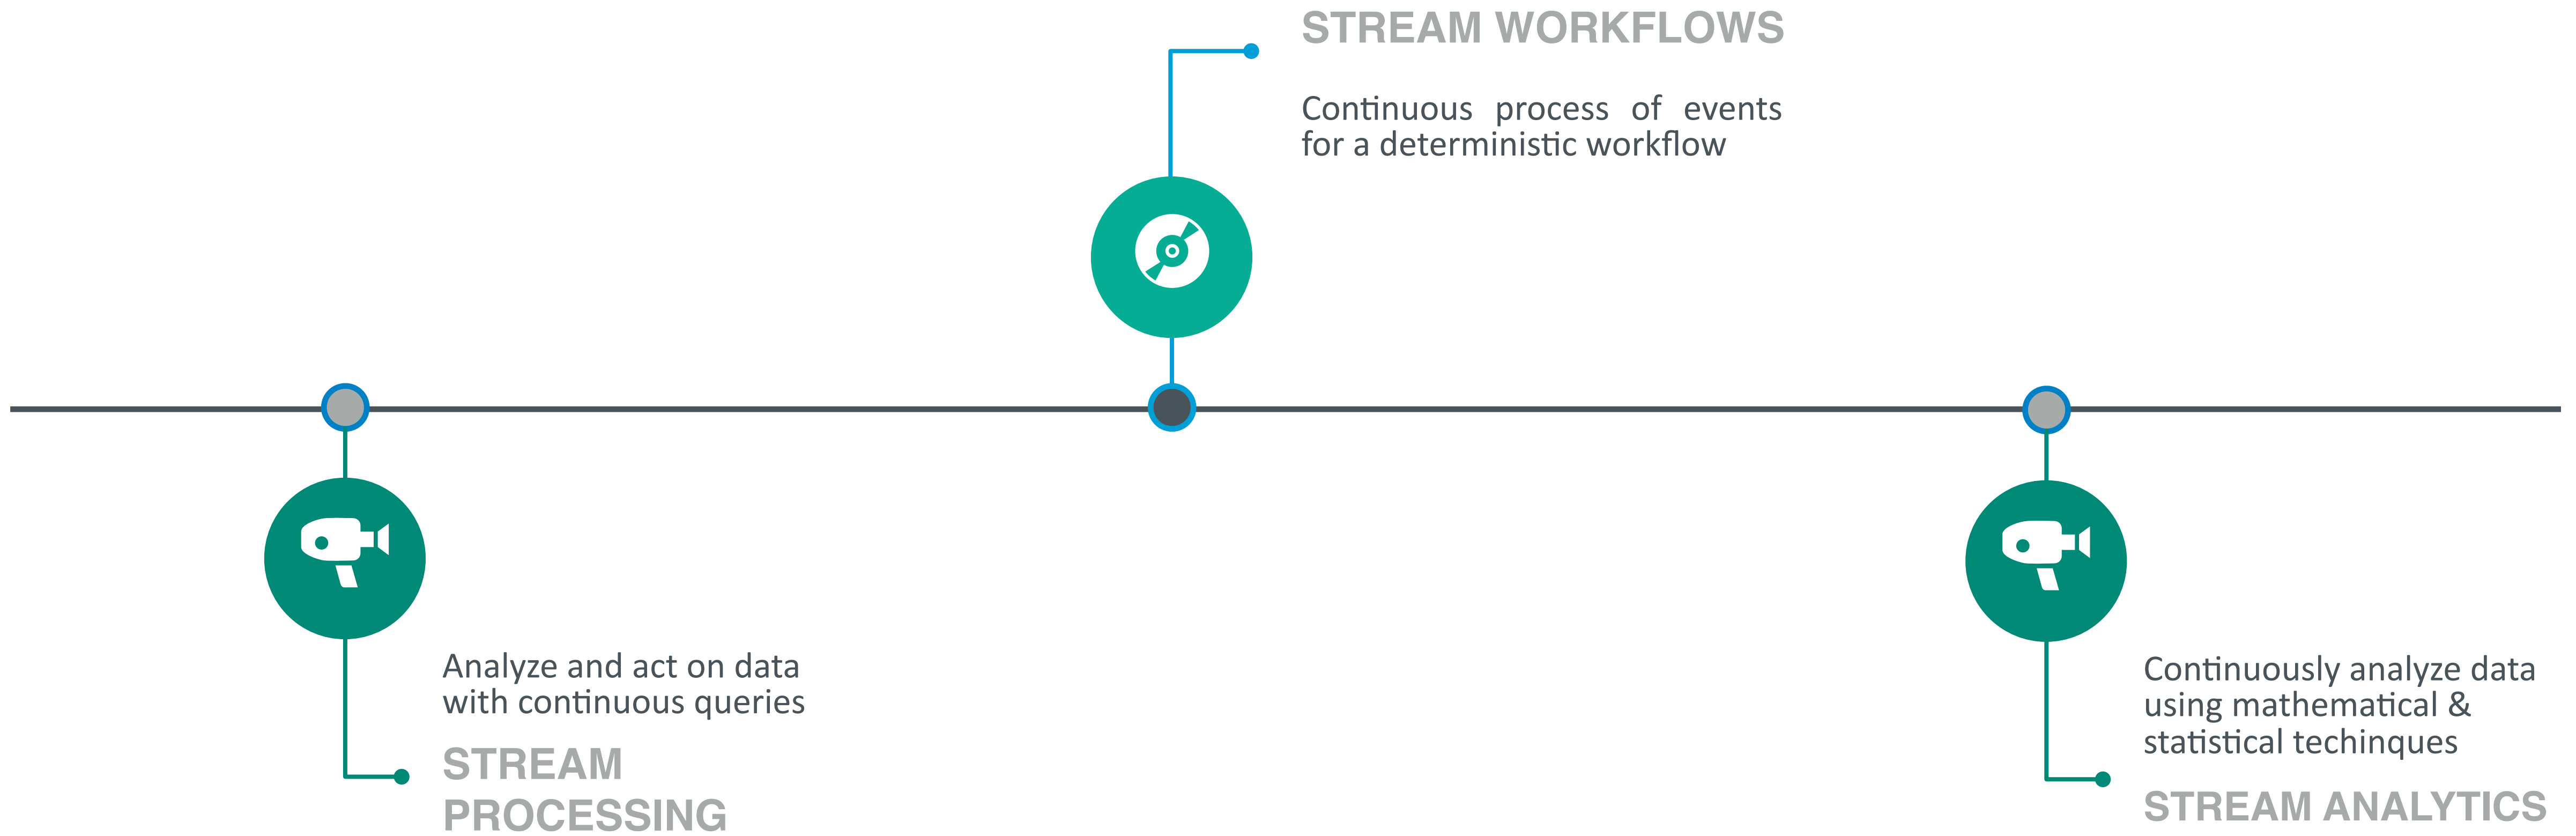
low latency



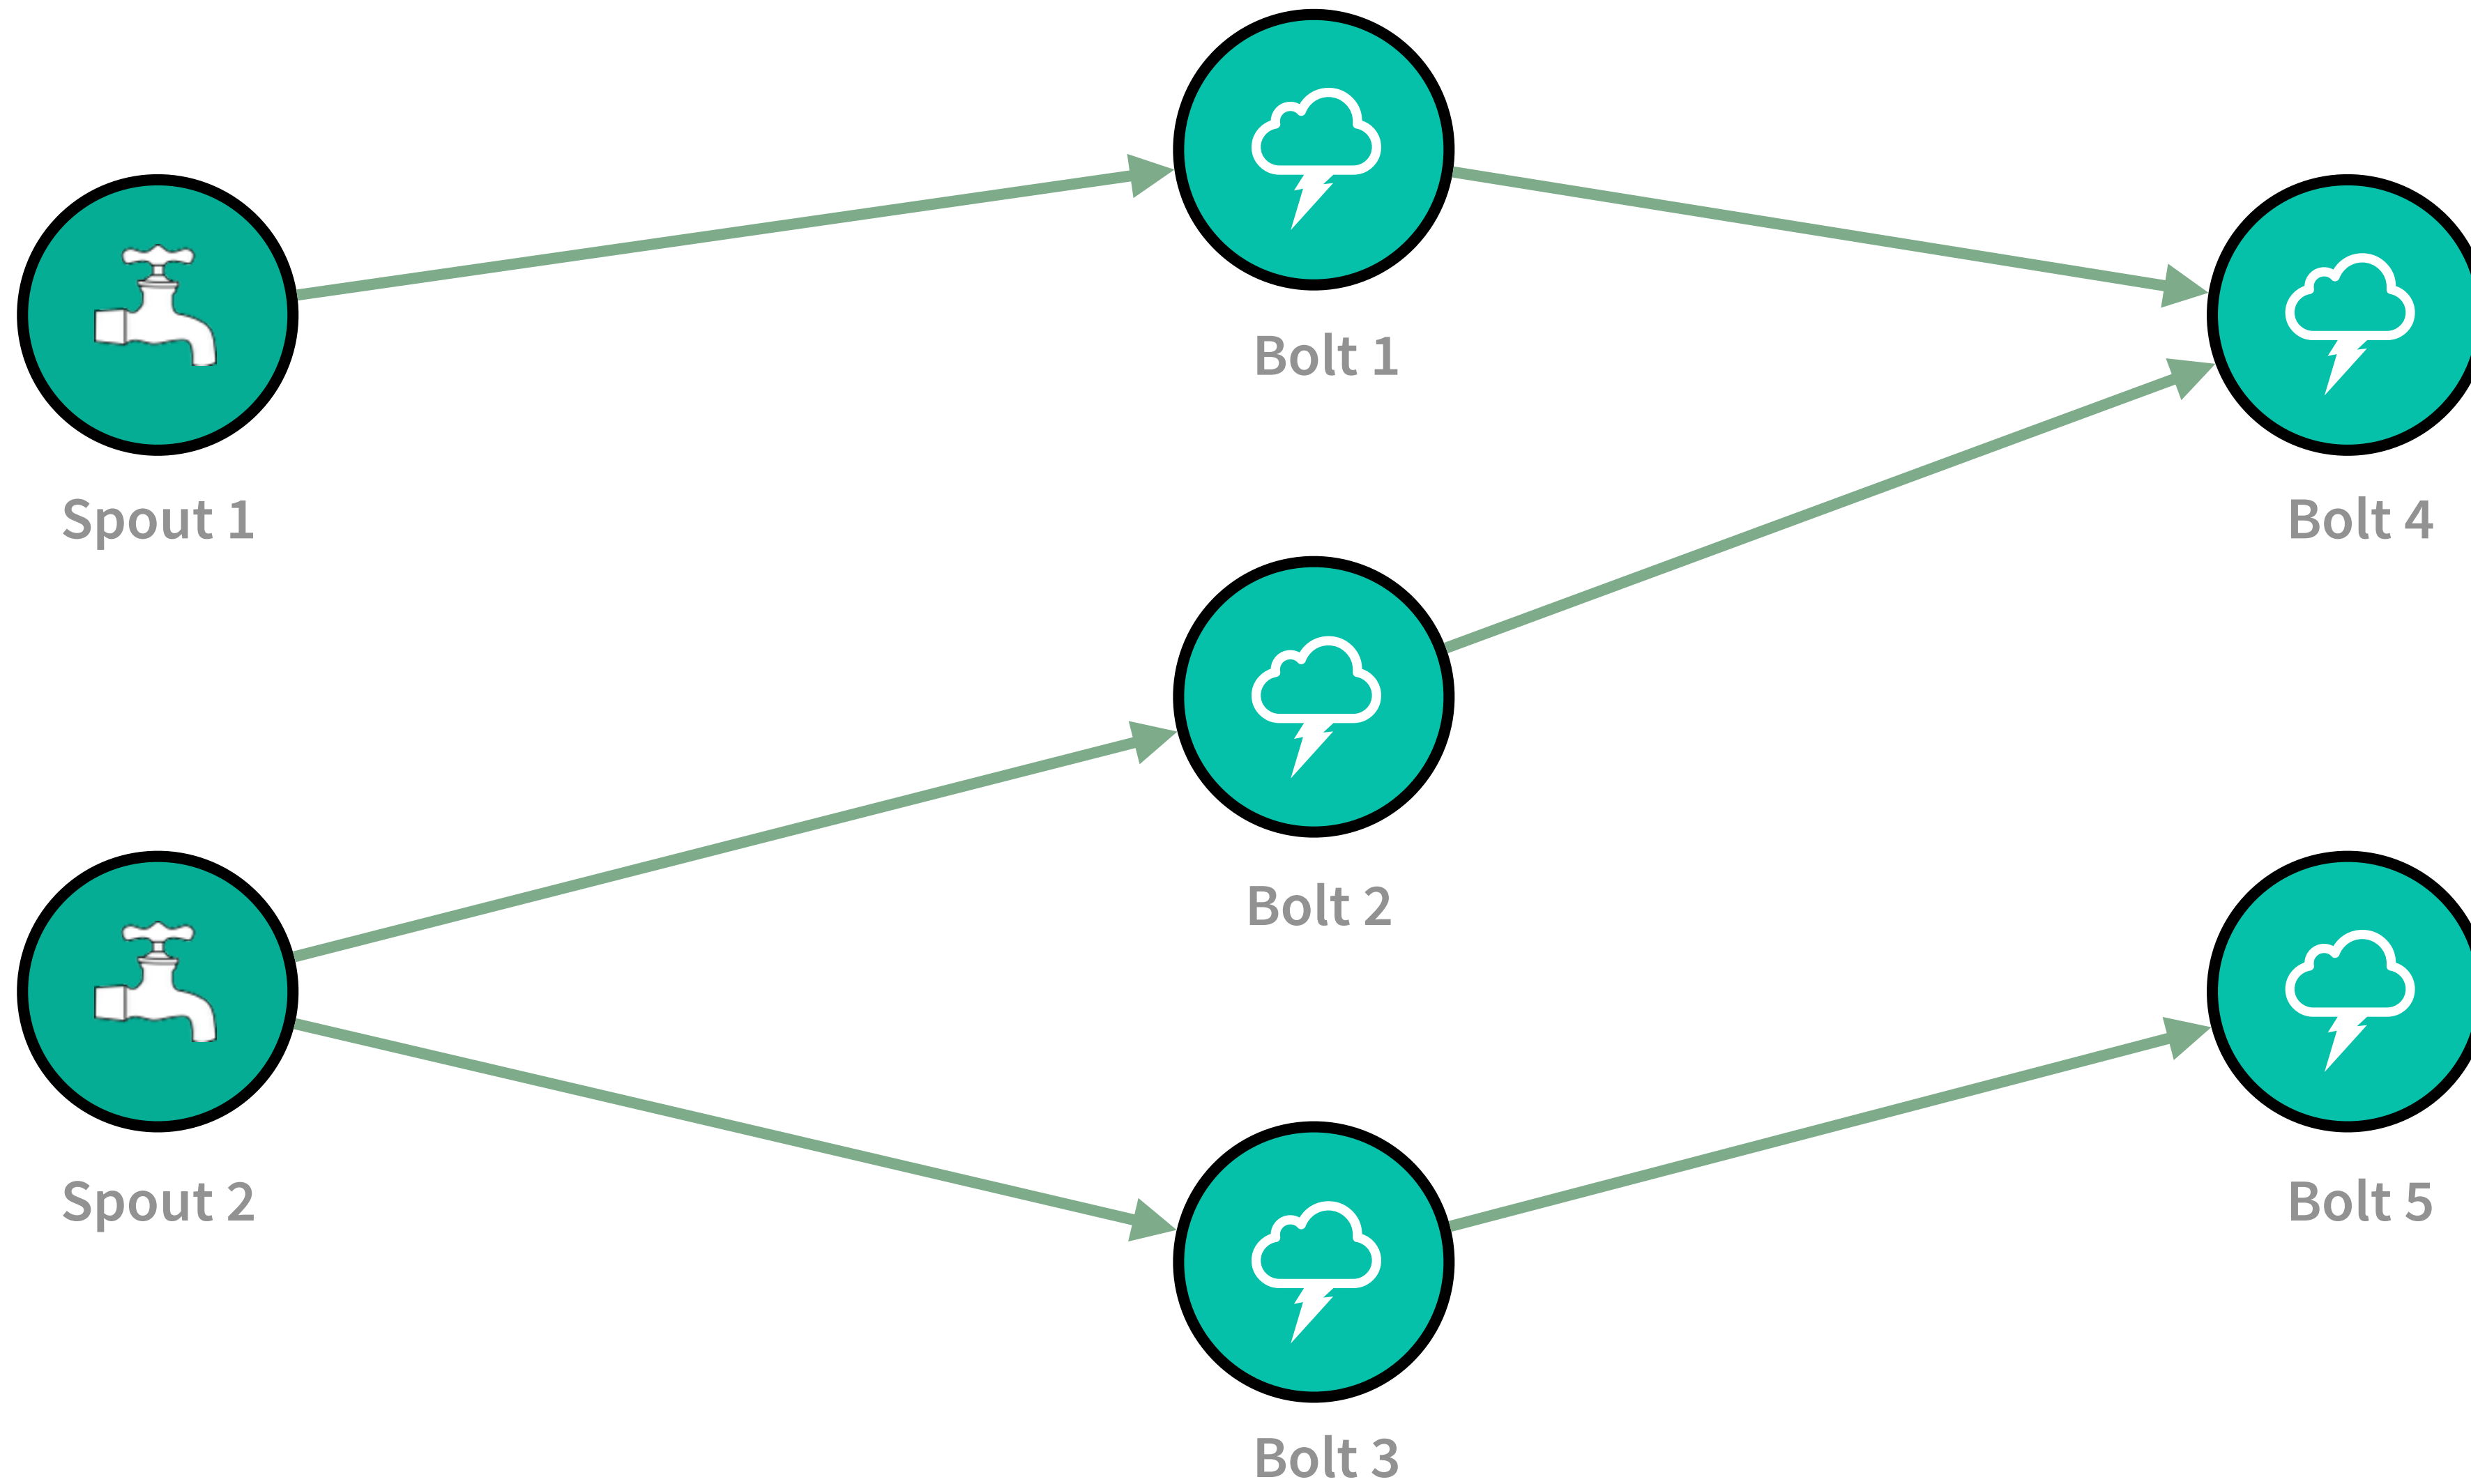
Financial Trading




Streaming Variants




Heron Topology



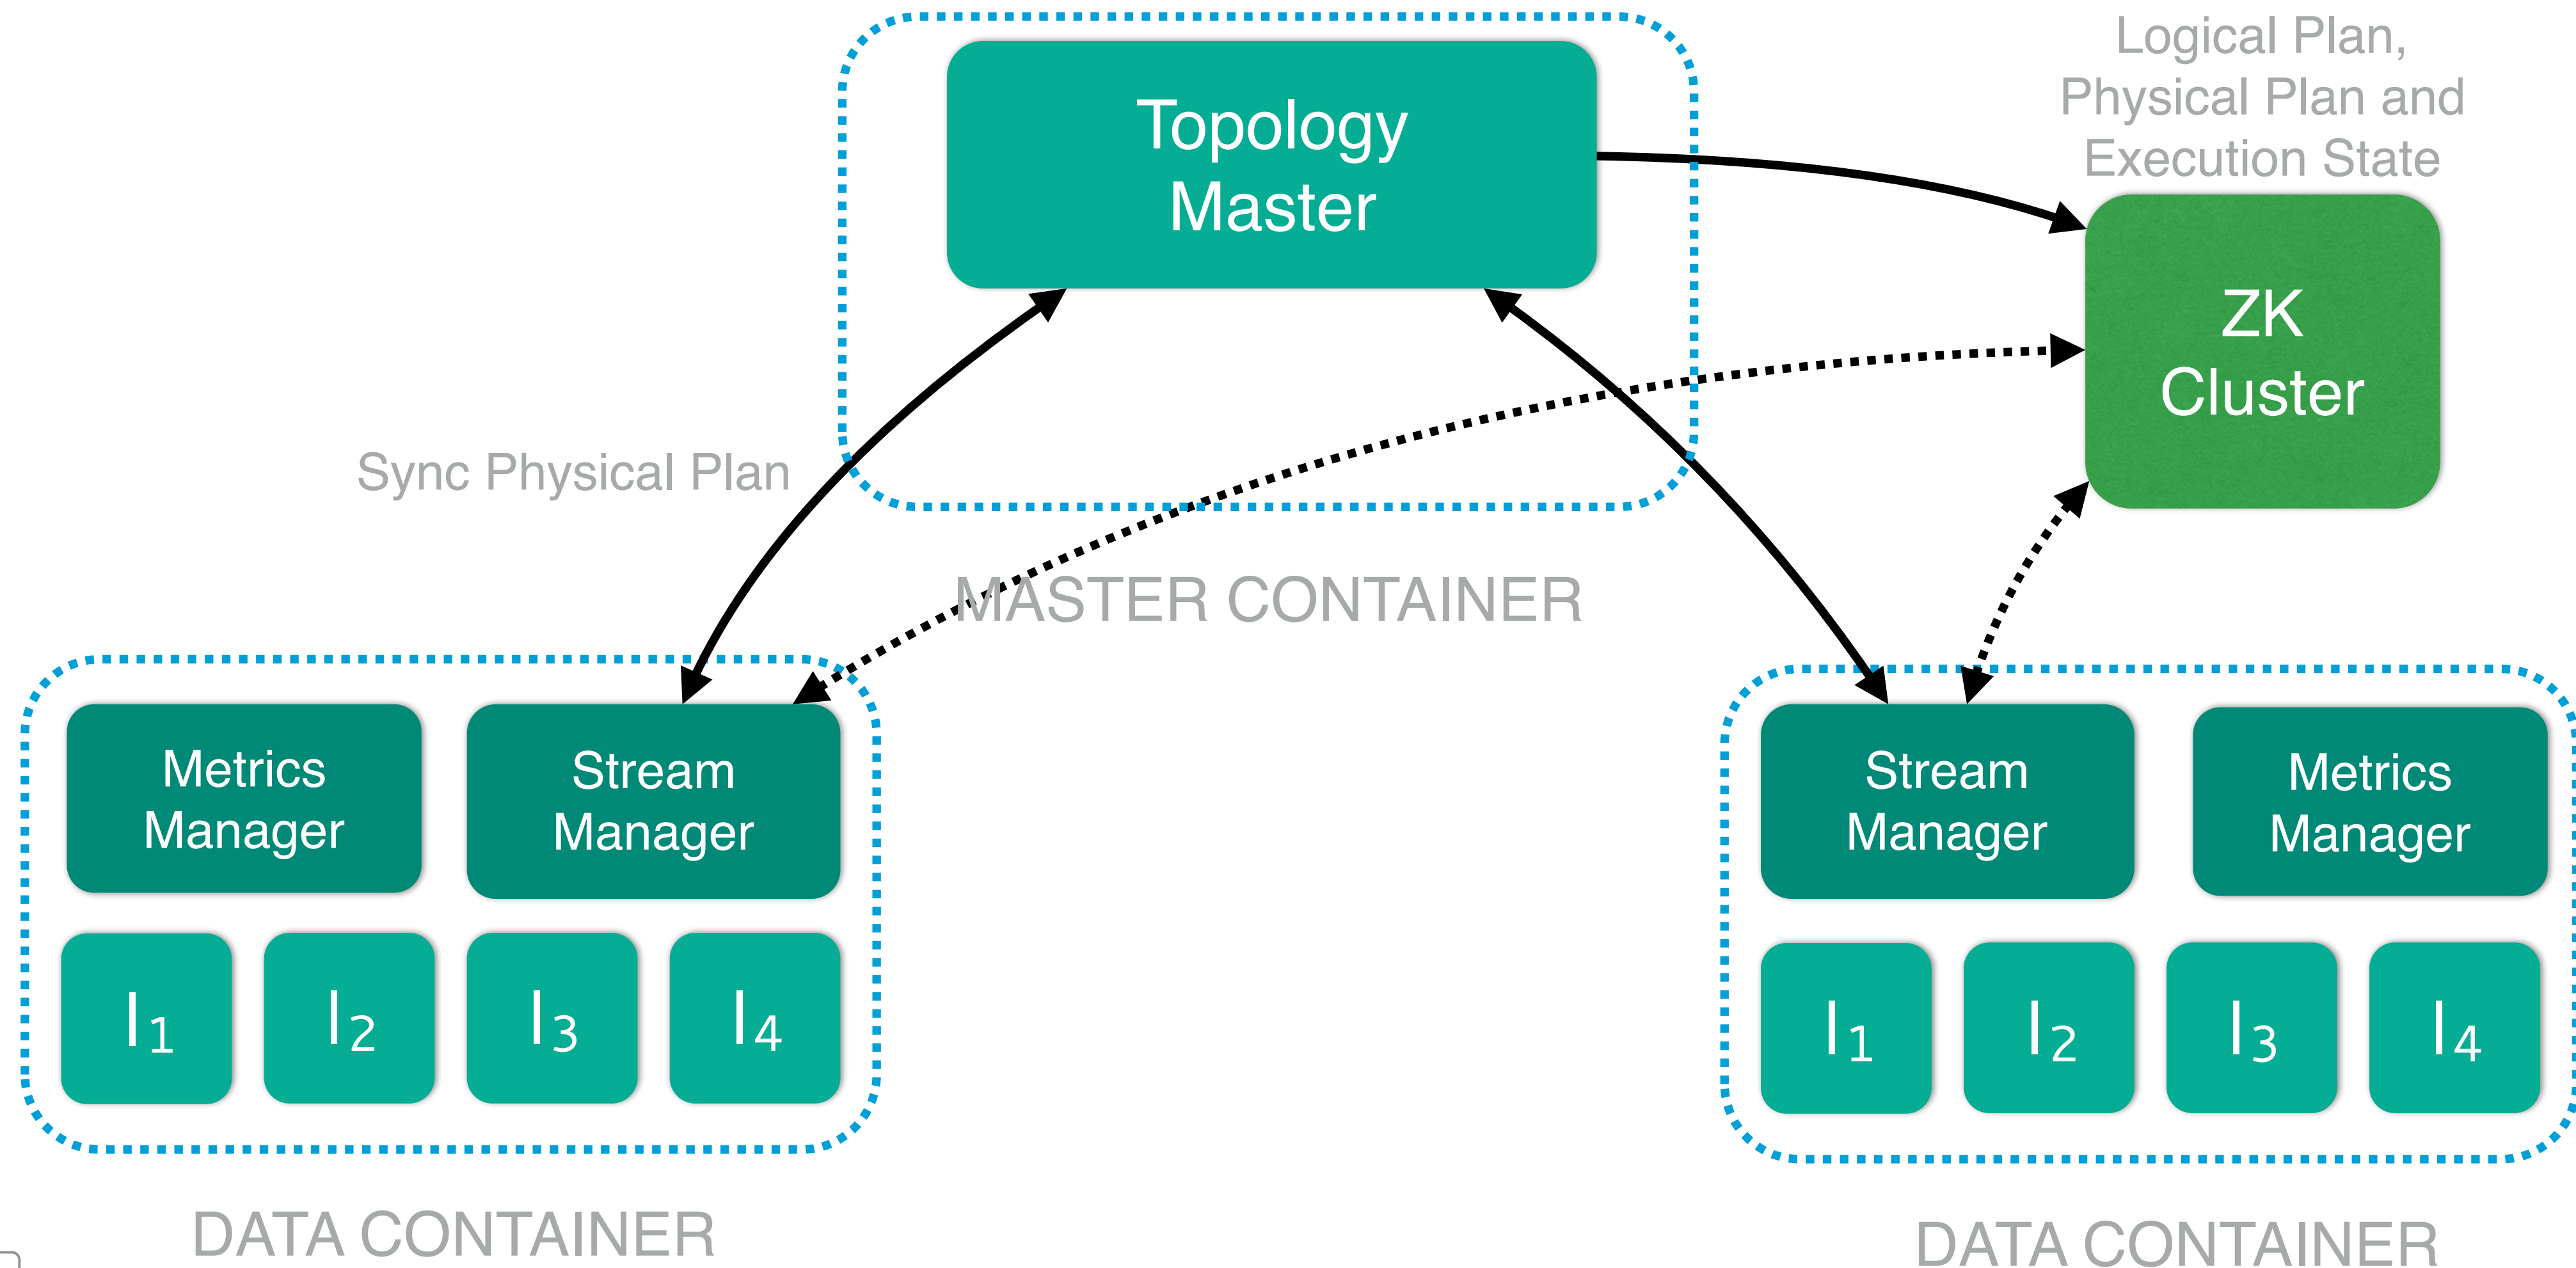
Heron Topology - Physical Execution



Heron



Heron Topology Components



Heron @Twitter

> 500 Real
Time Jobs

500 Billions Events/Day
PROCESSED

25-200
MS
latency

Common Operational Issues




01

Slow Hosts



02


Network Issues




03

Data Skew

Slow Hosts



Memory Parity Errors



Impeding Disk Failures

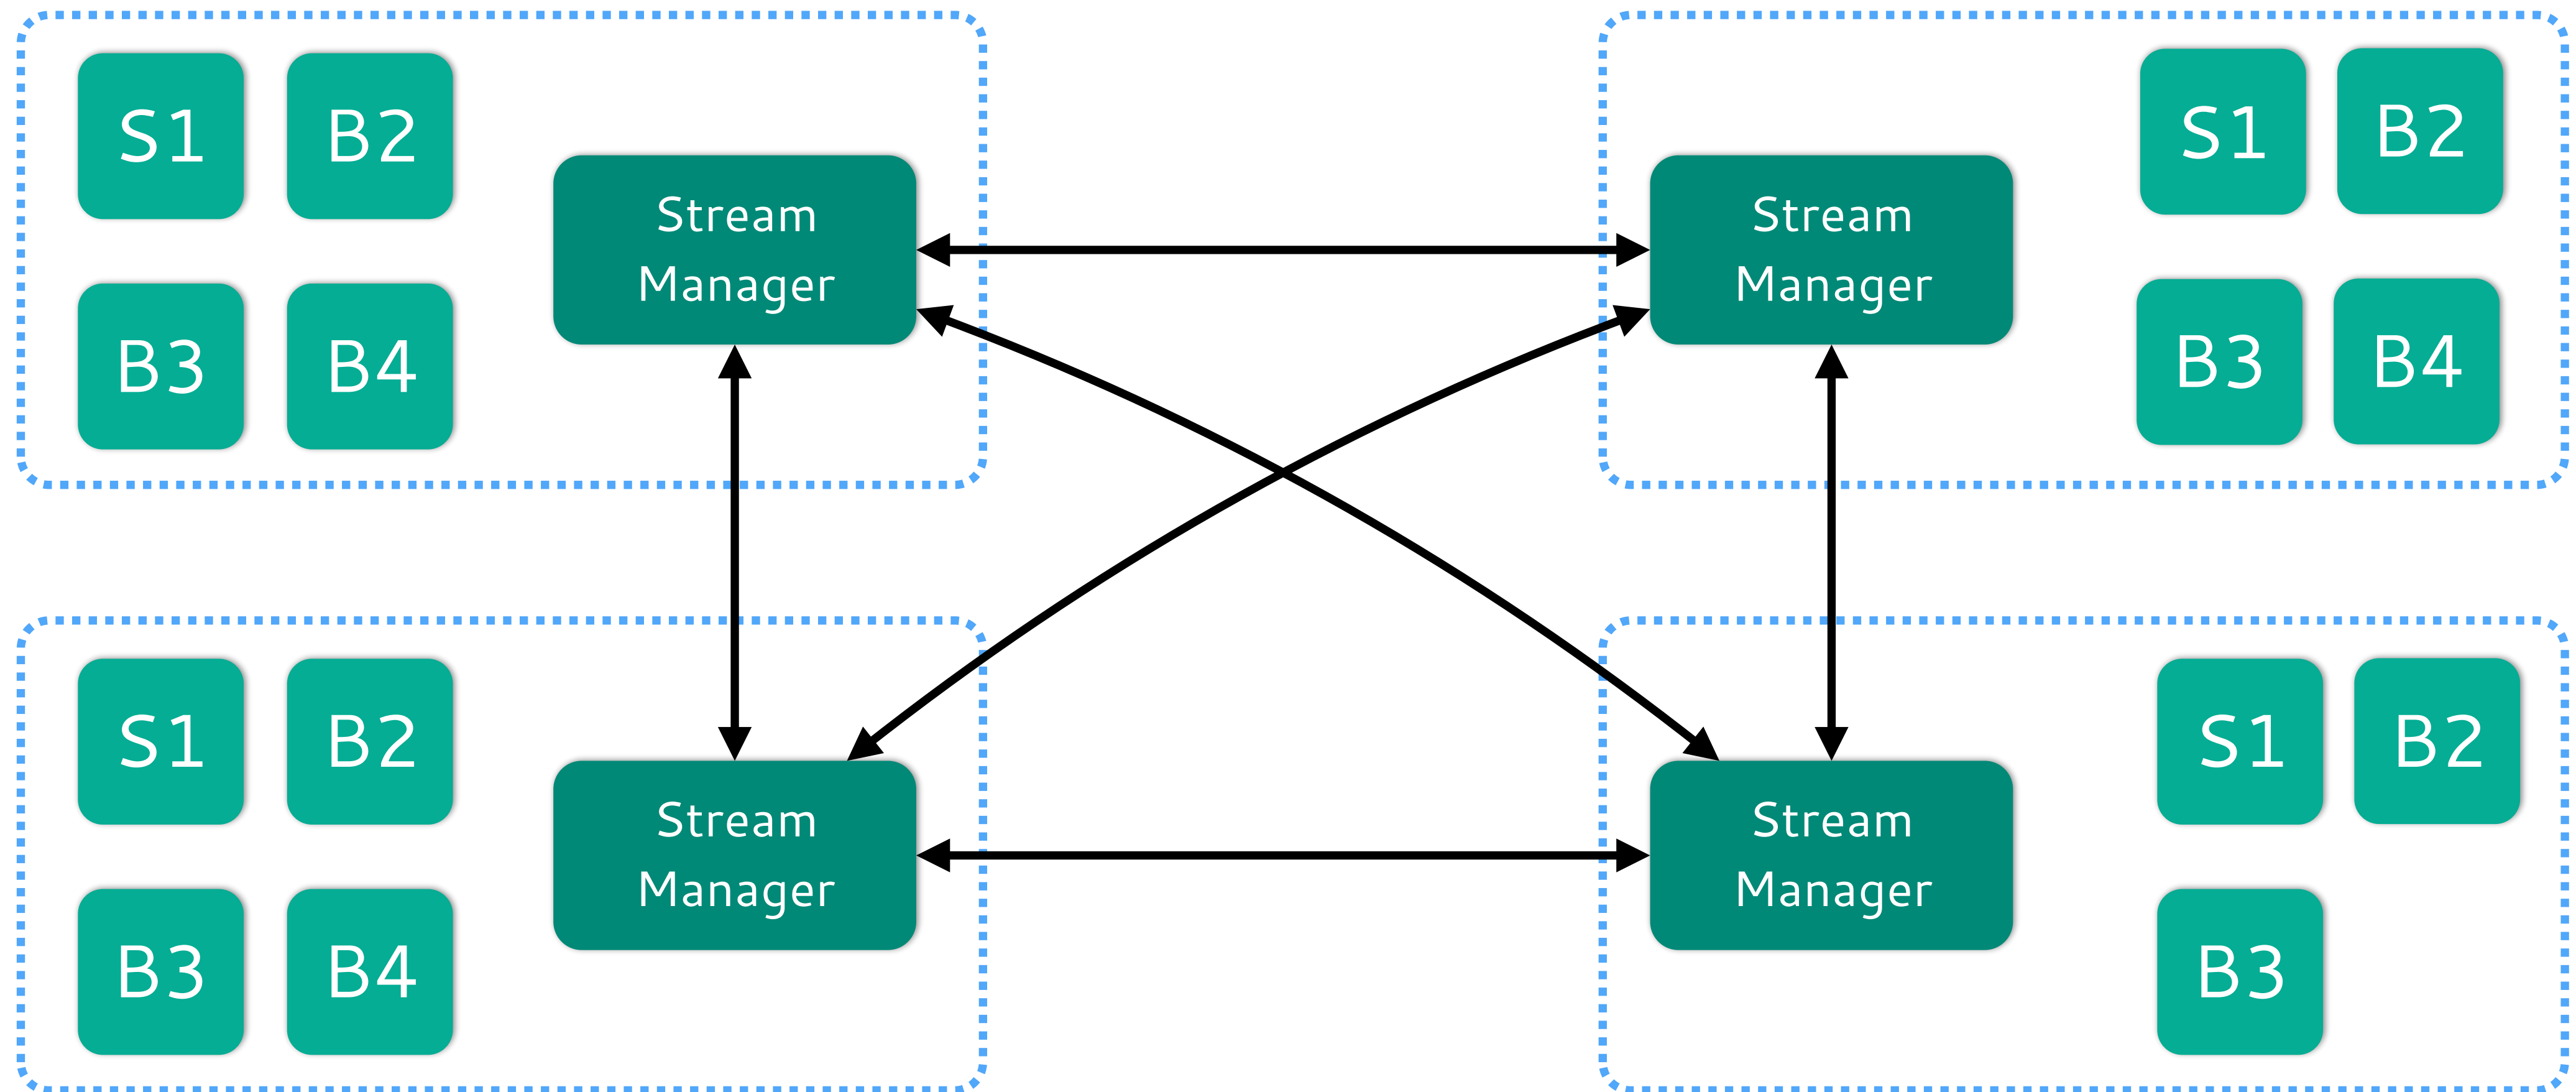


Lower GHZ

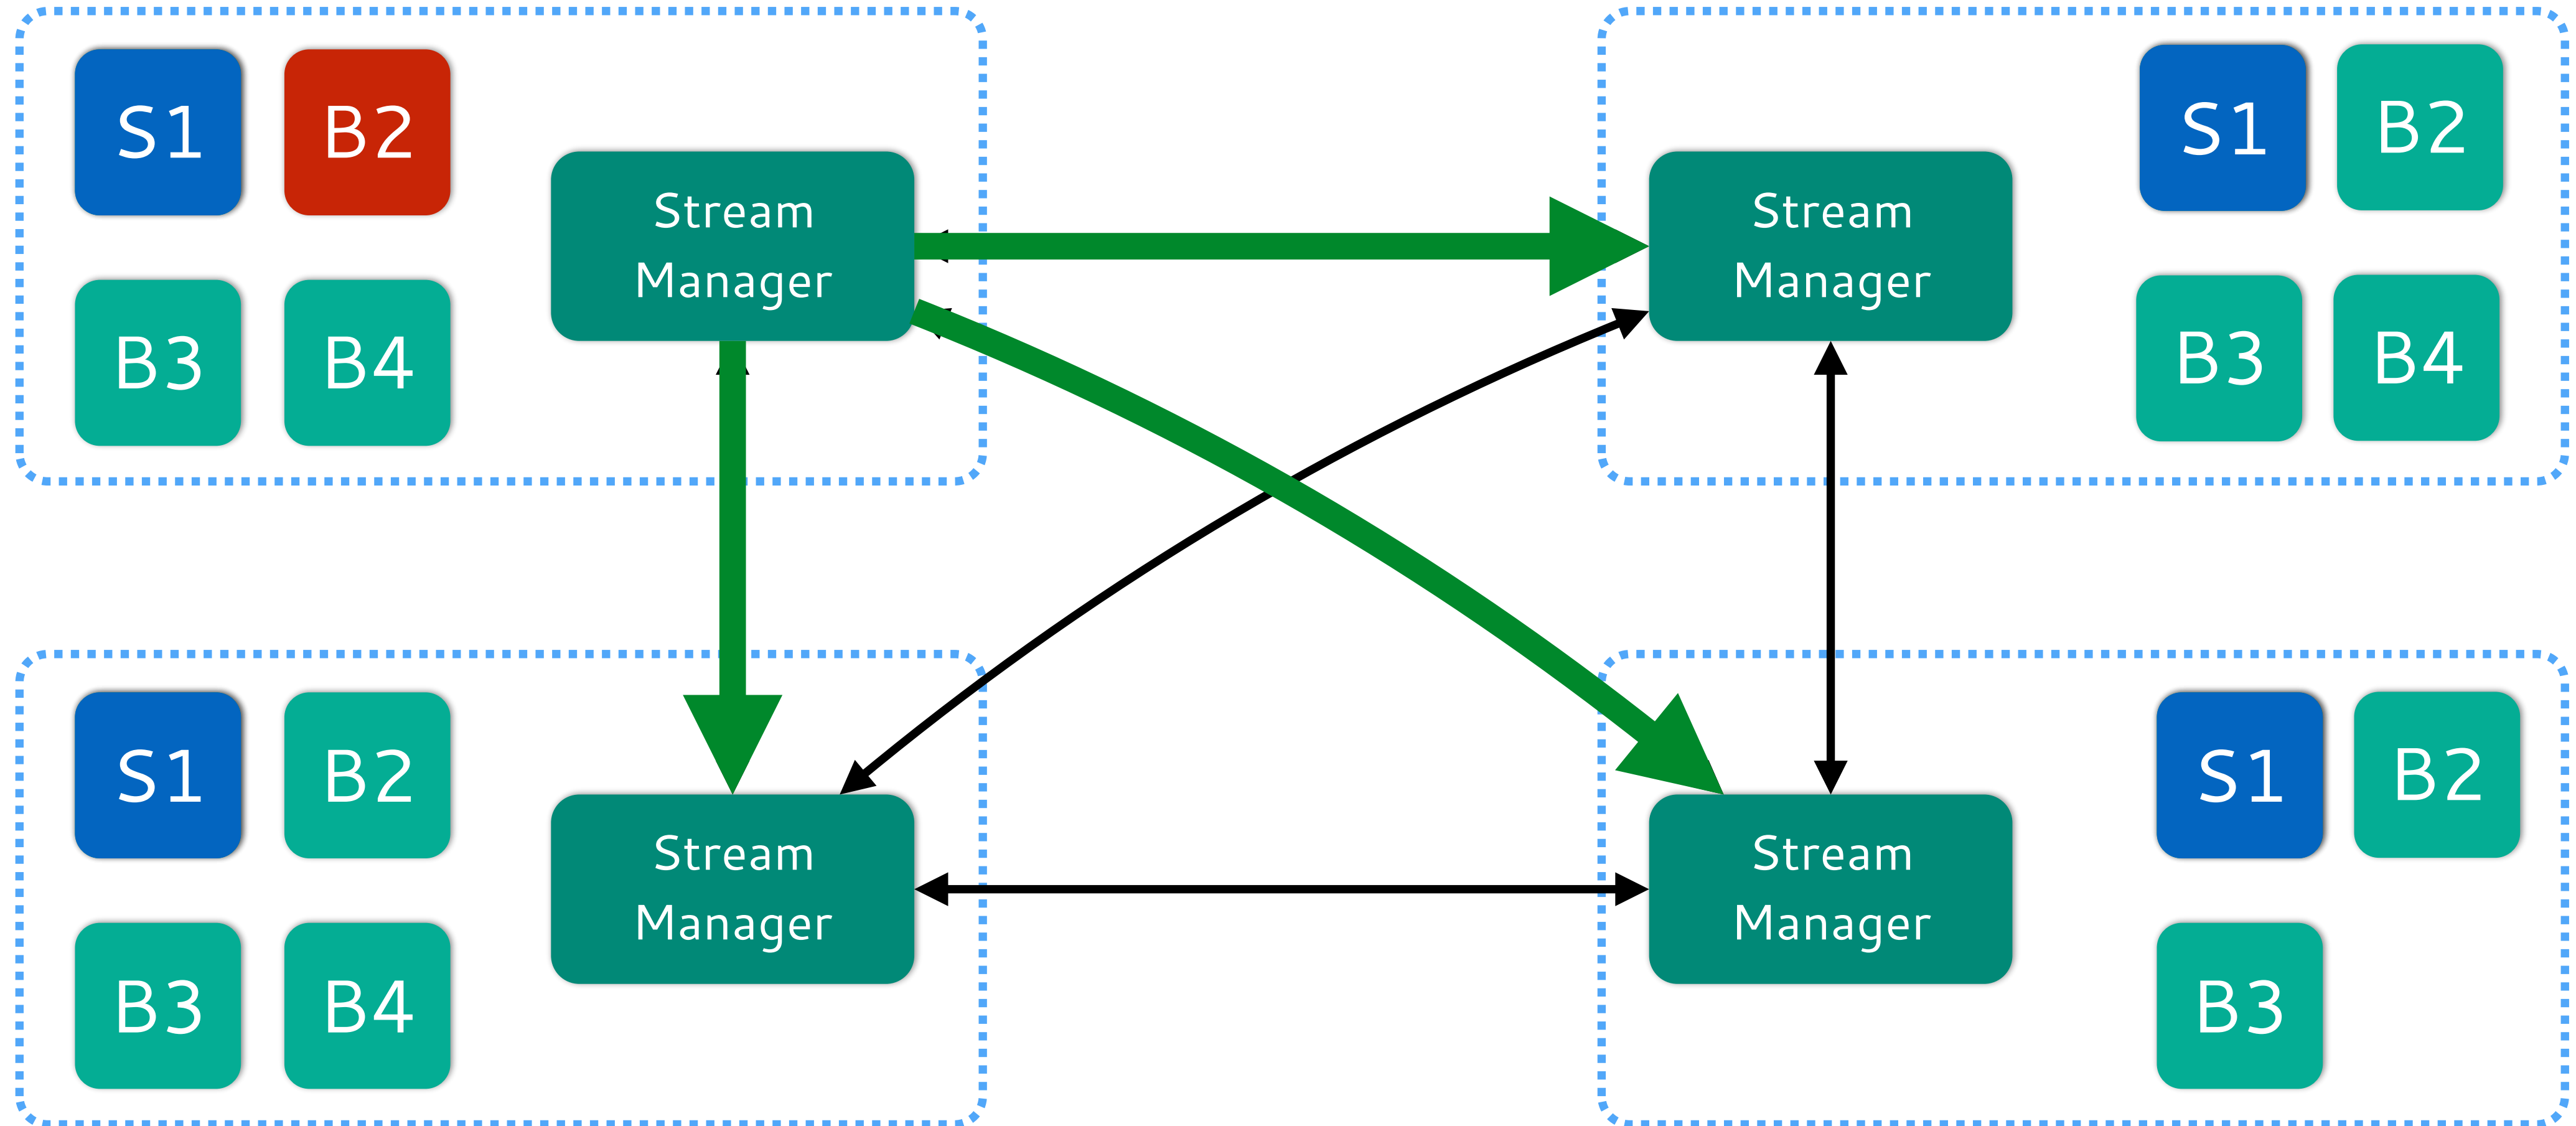
Heron Backpressure



Stream Manager




Spout Backpressure




Can we do better?

Network



Network Slowness



Network Partitioning

Network Slowness



Delays processing

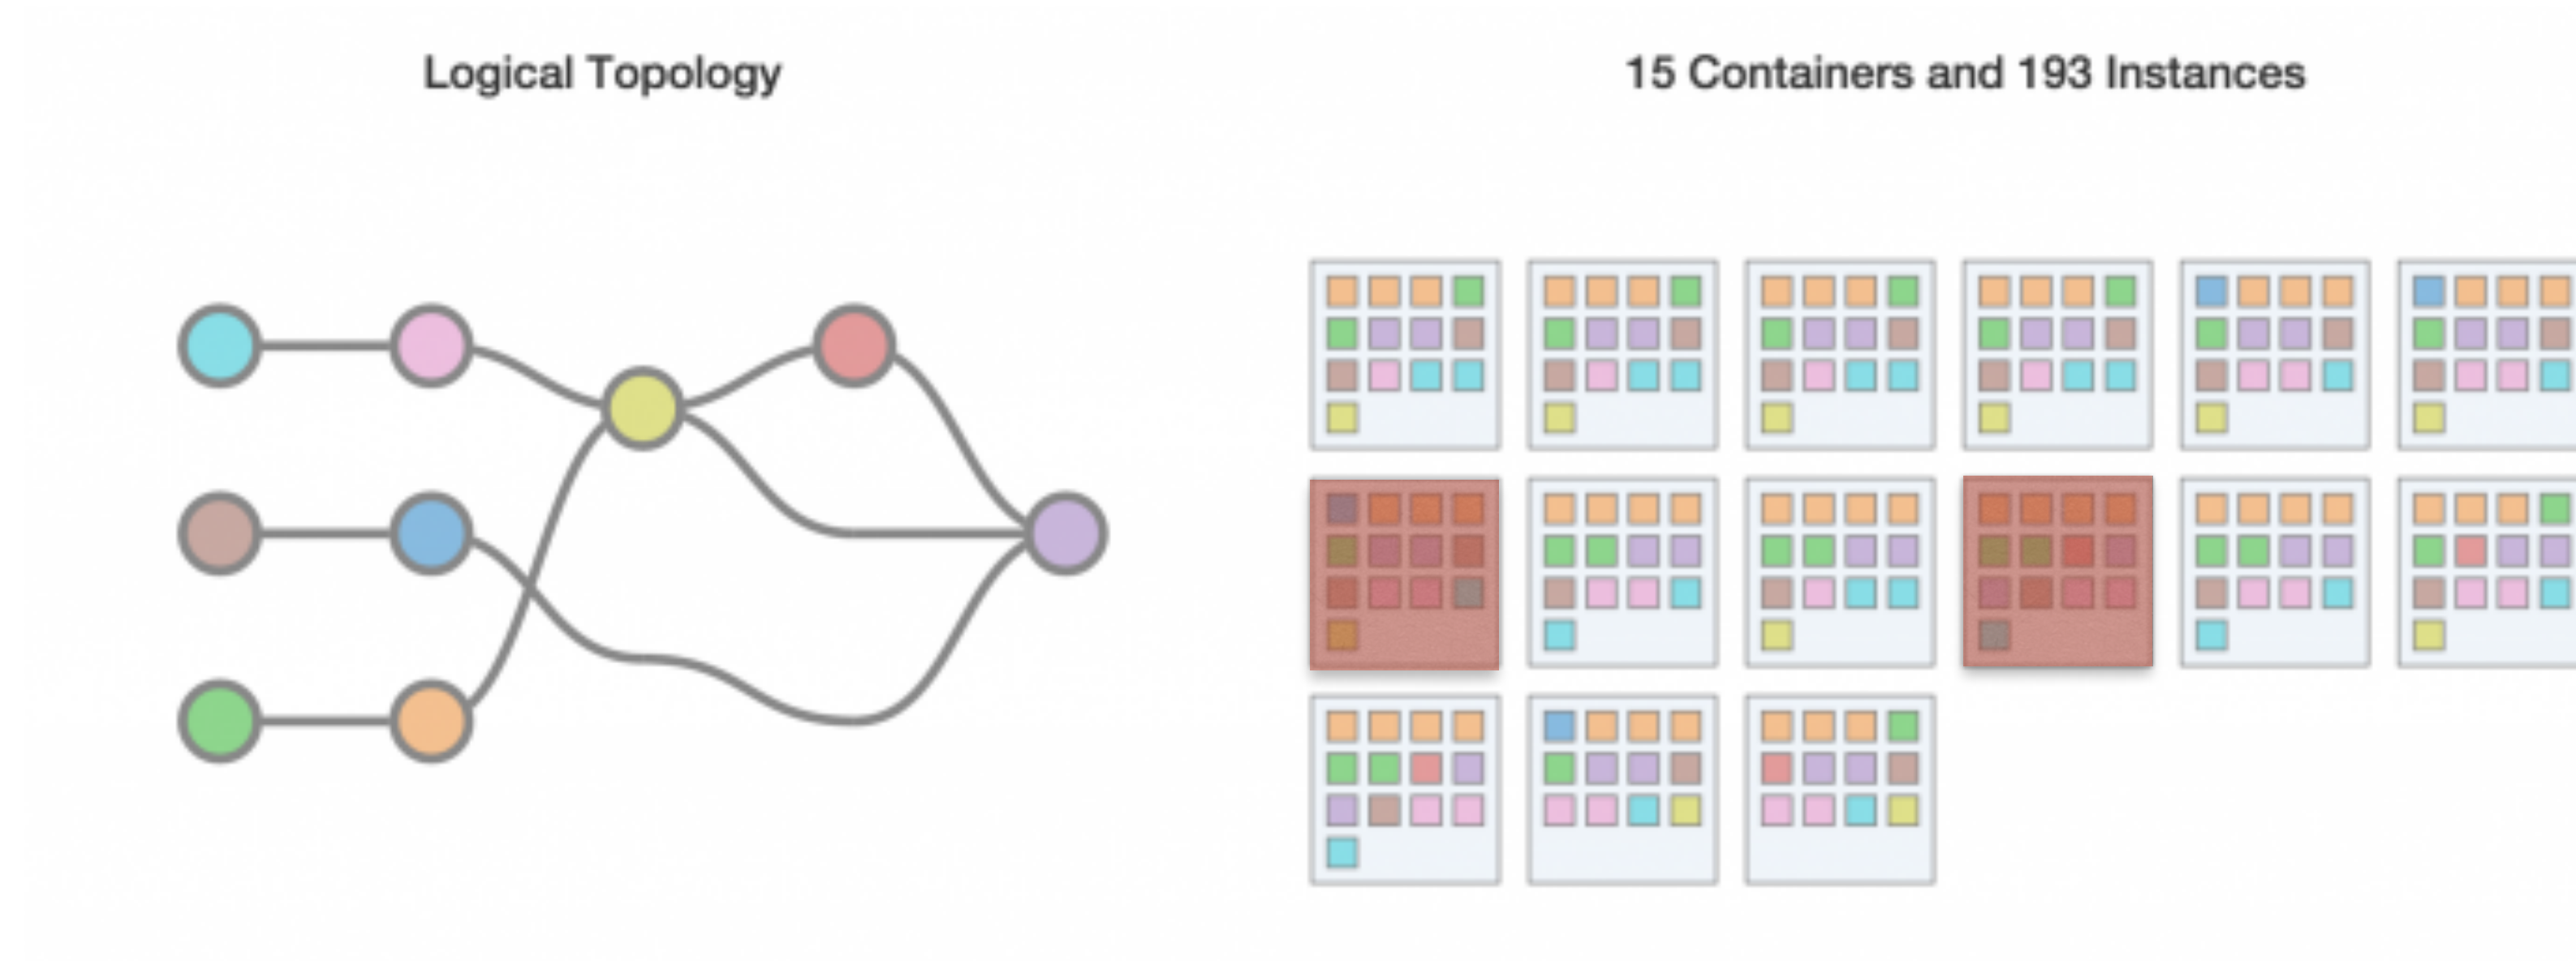


Data is accumulating







Timeliness of results is affected

Network Slowness





Can we do better?


Network Partitioning



Network Partitioning



New Master Container

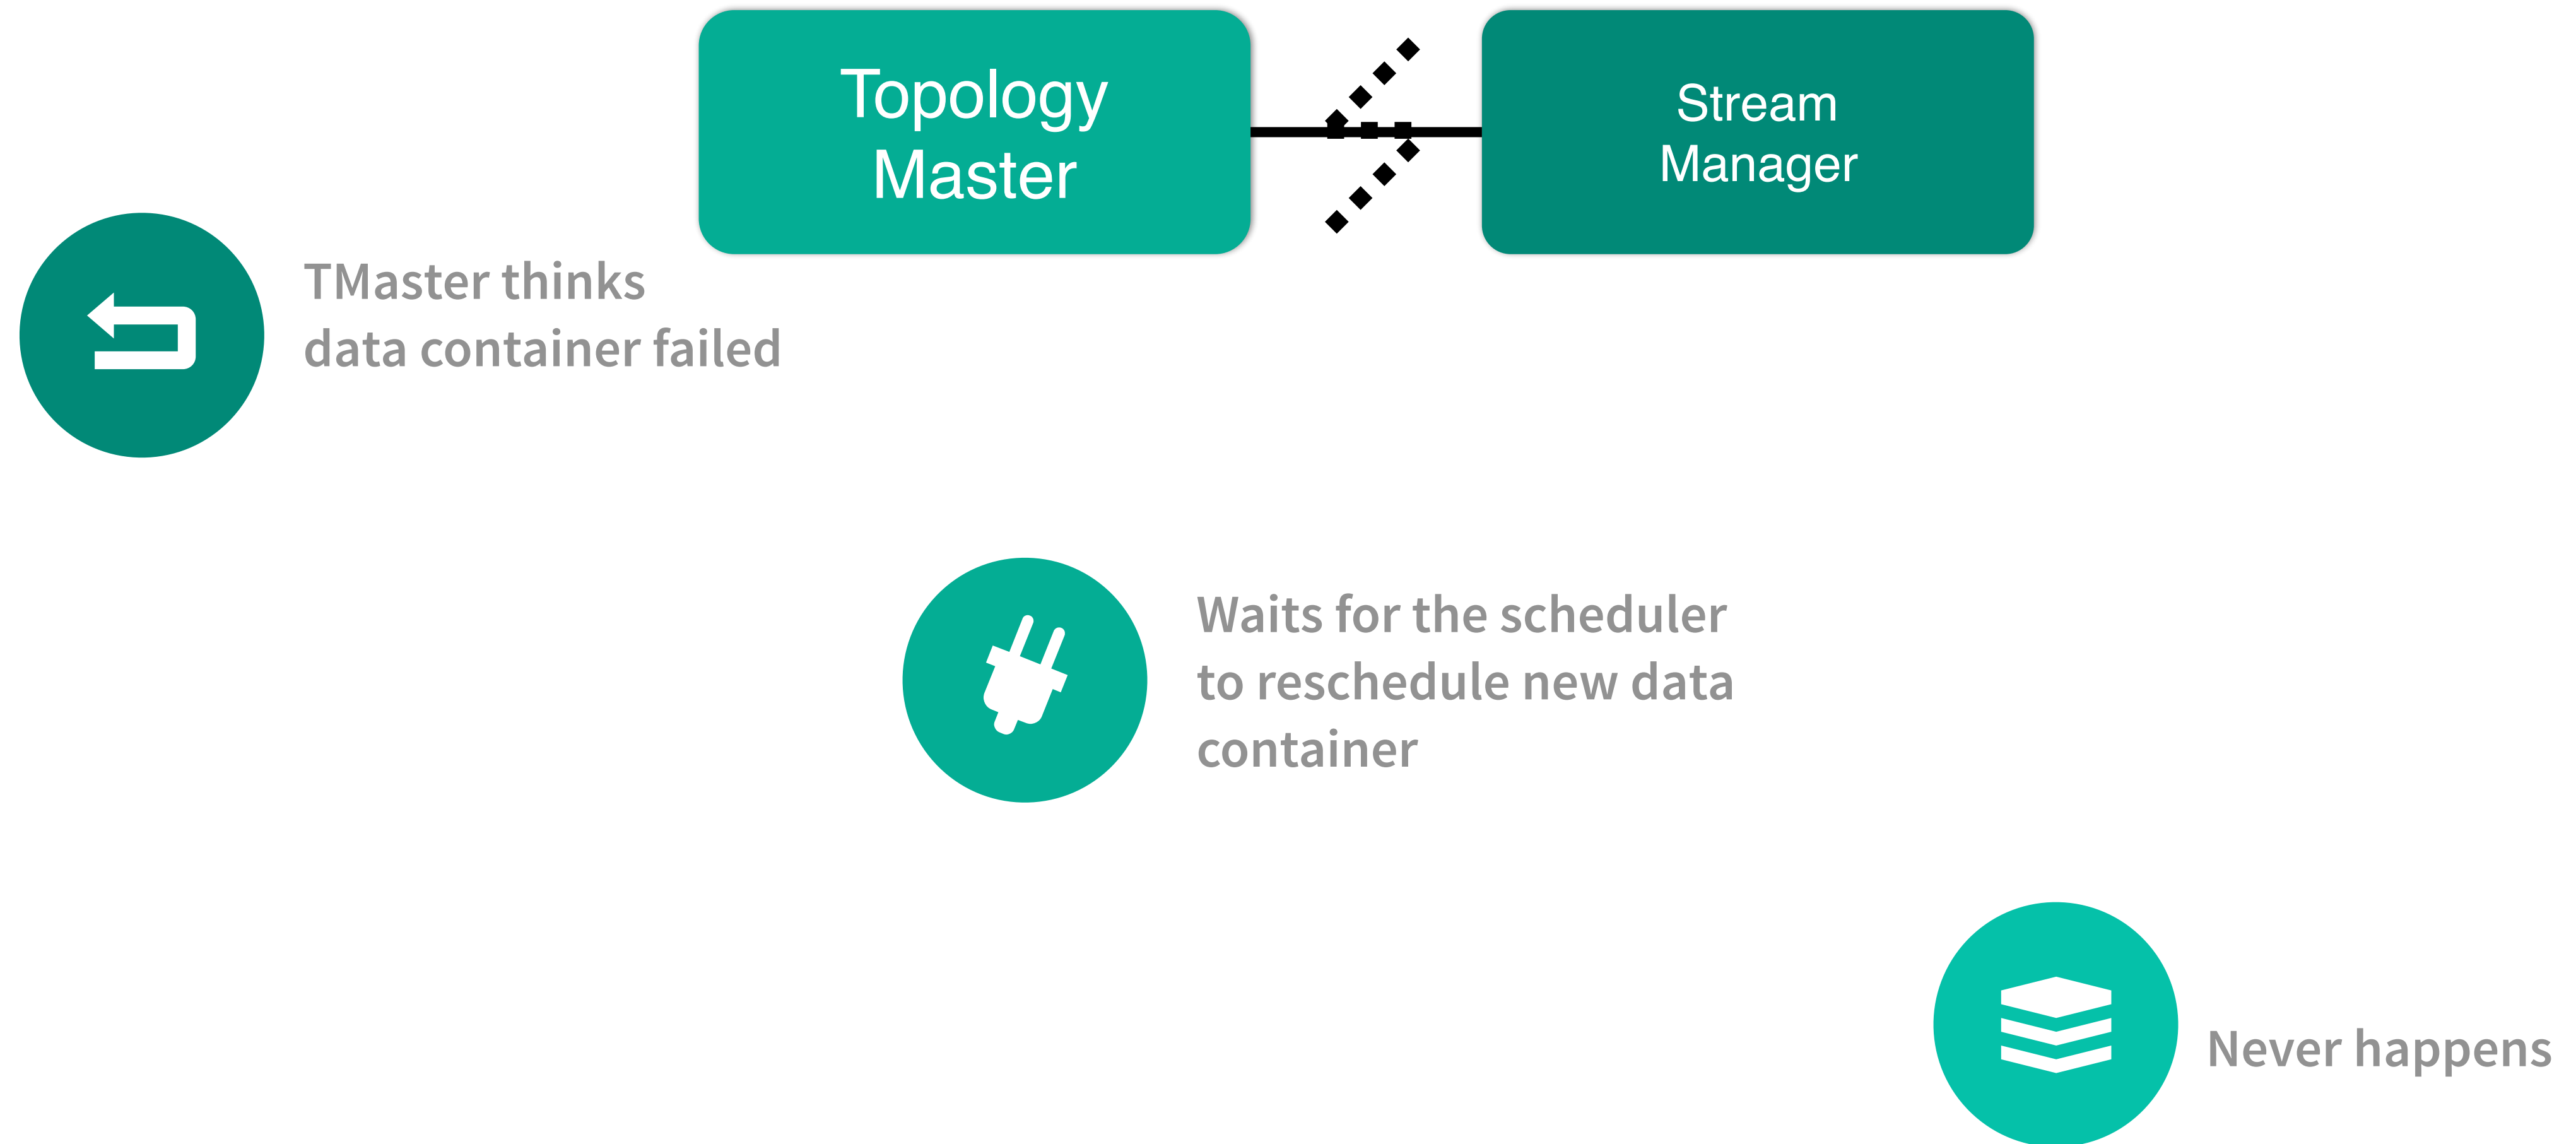


Acquiring Mastership
in ZooKeeper fails





Master Container Dies


Network Partitioning



Network Partitioning



Cannot exchange data





Data accumulates




Chaos ensues!

Can we do better?

Network Partitioning



New data container spawned

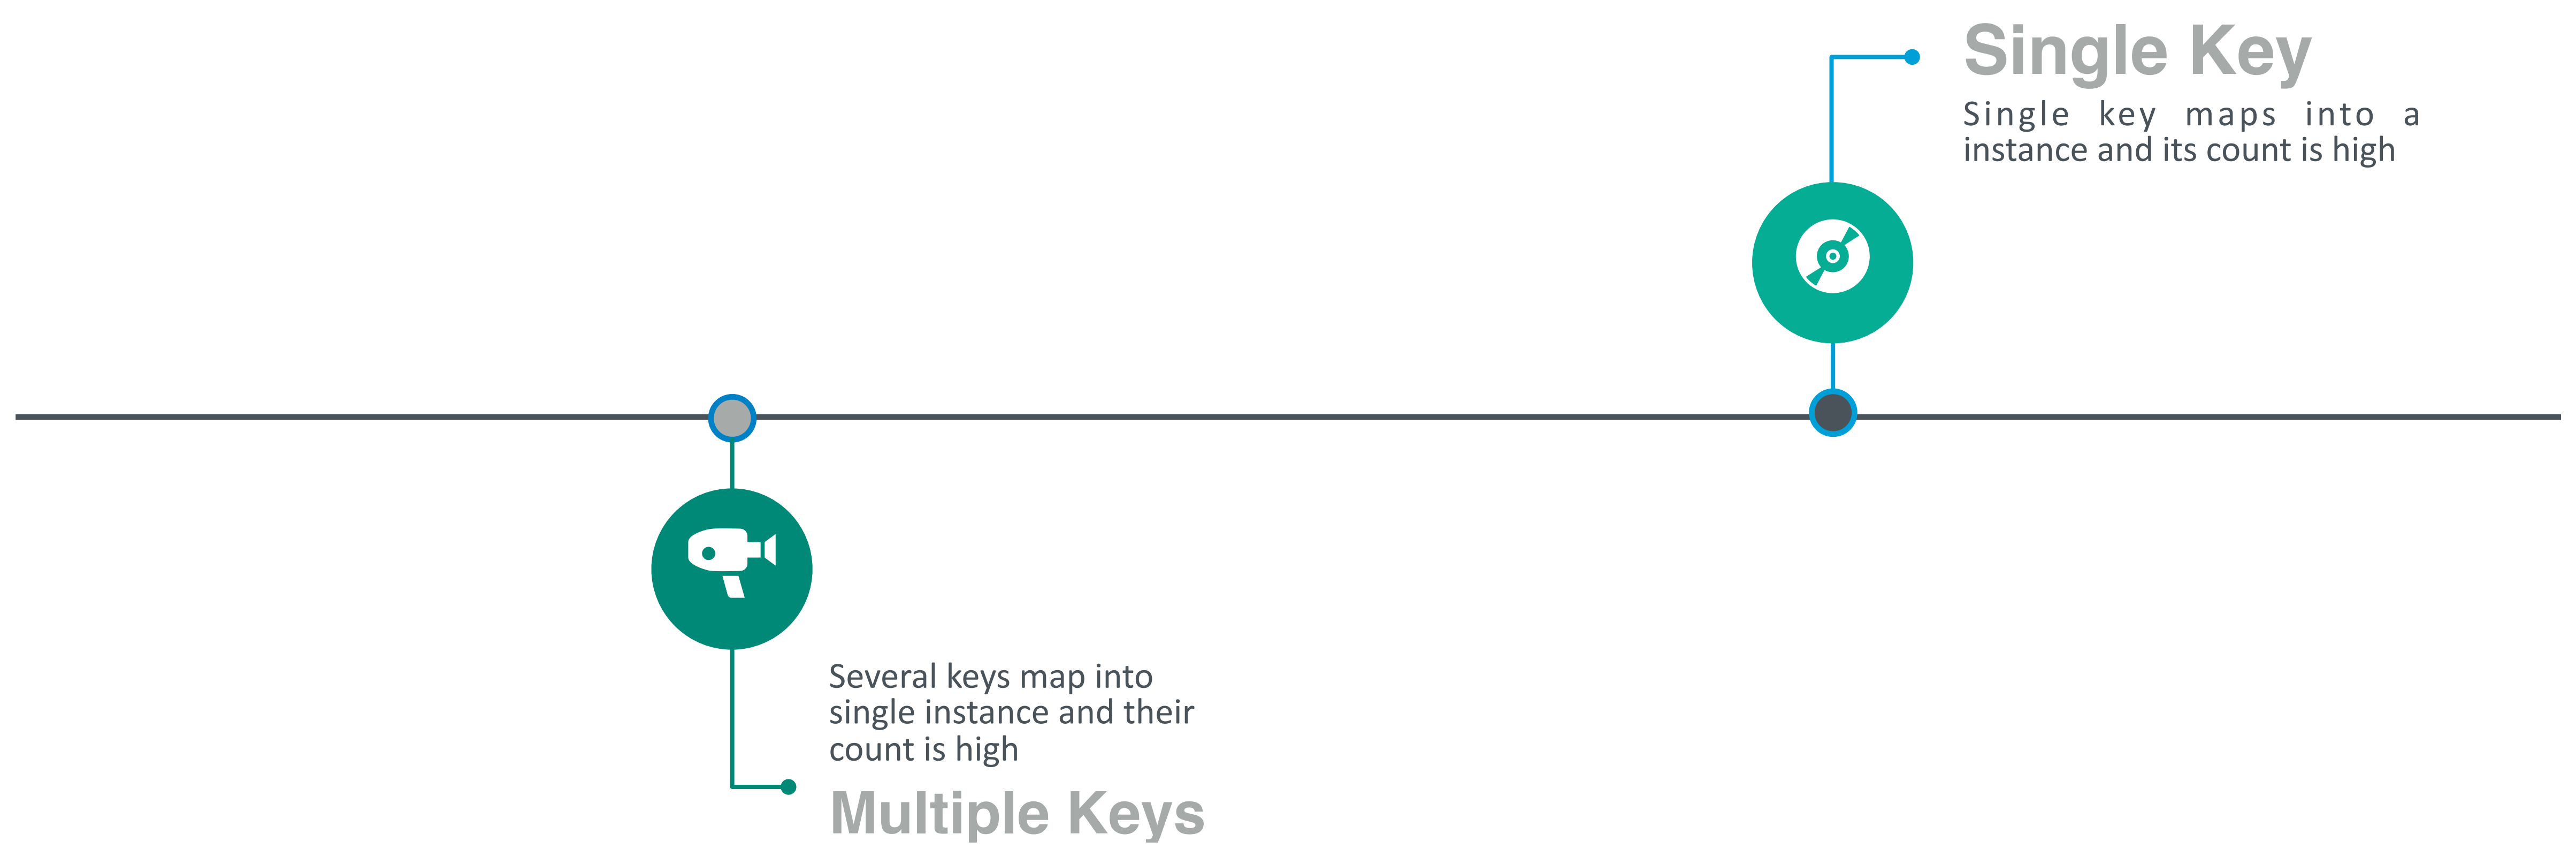


TMaster realizes two data containers report as the same

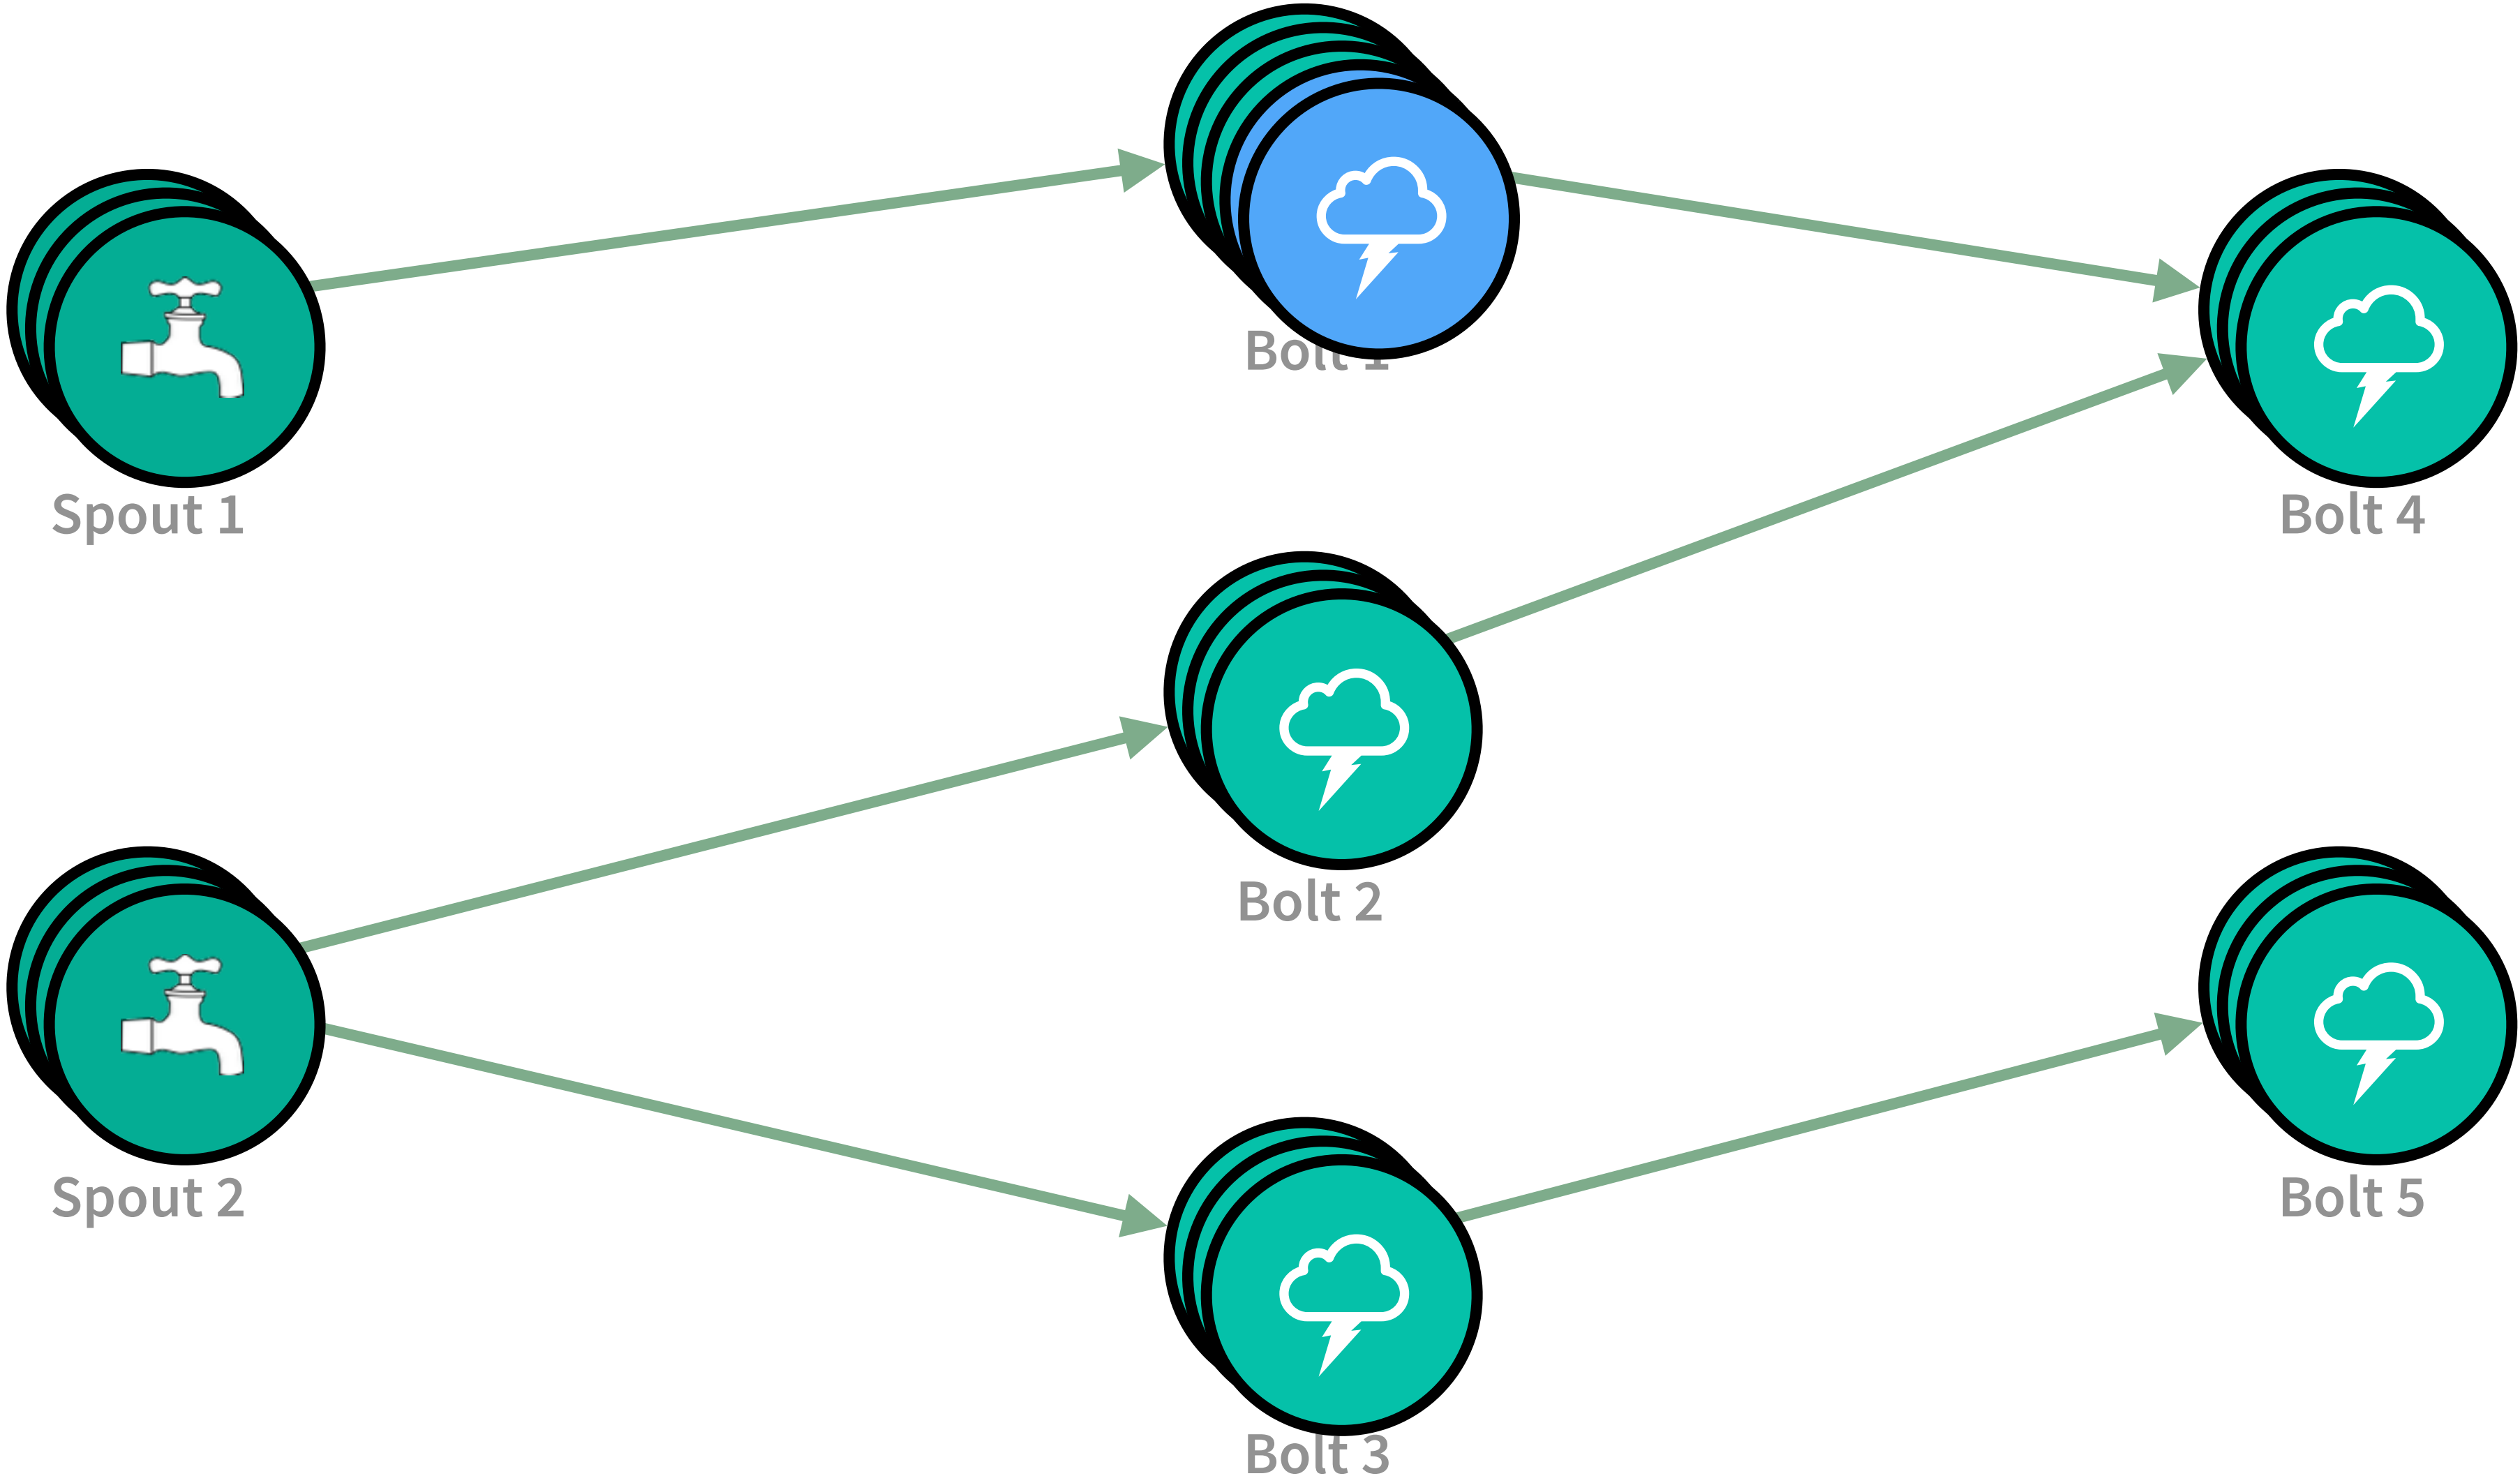


Does not accept the new one and eventually it dies

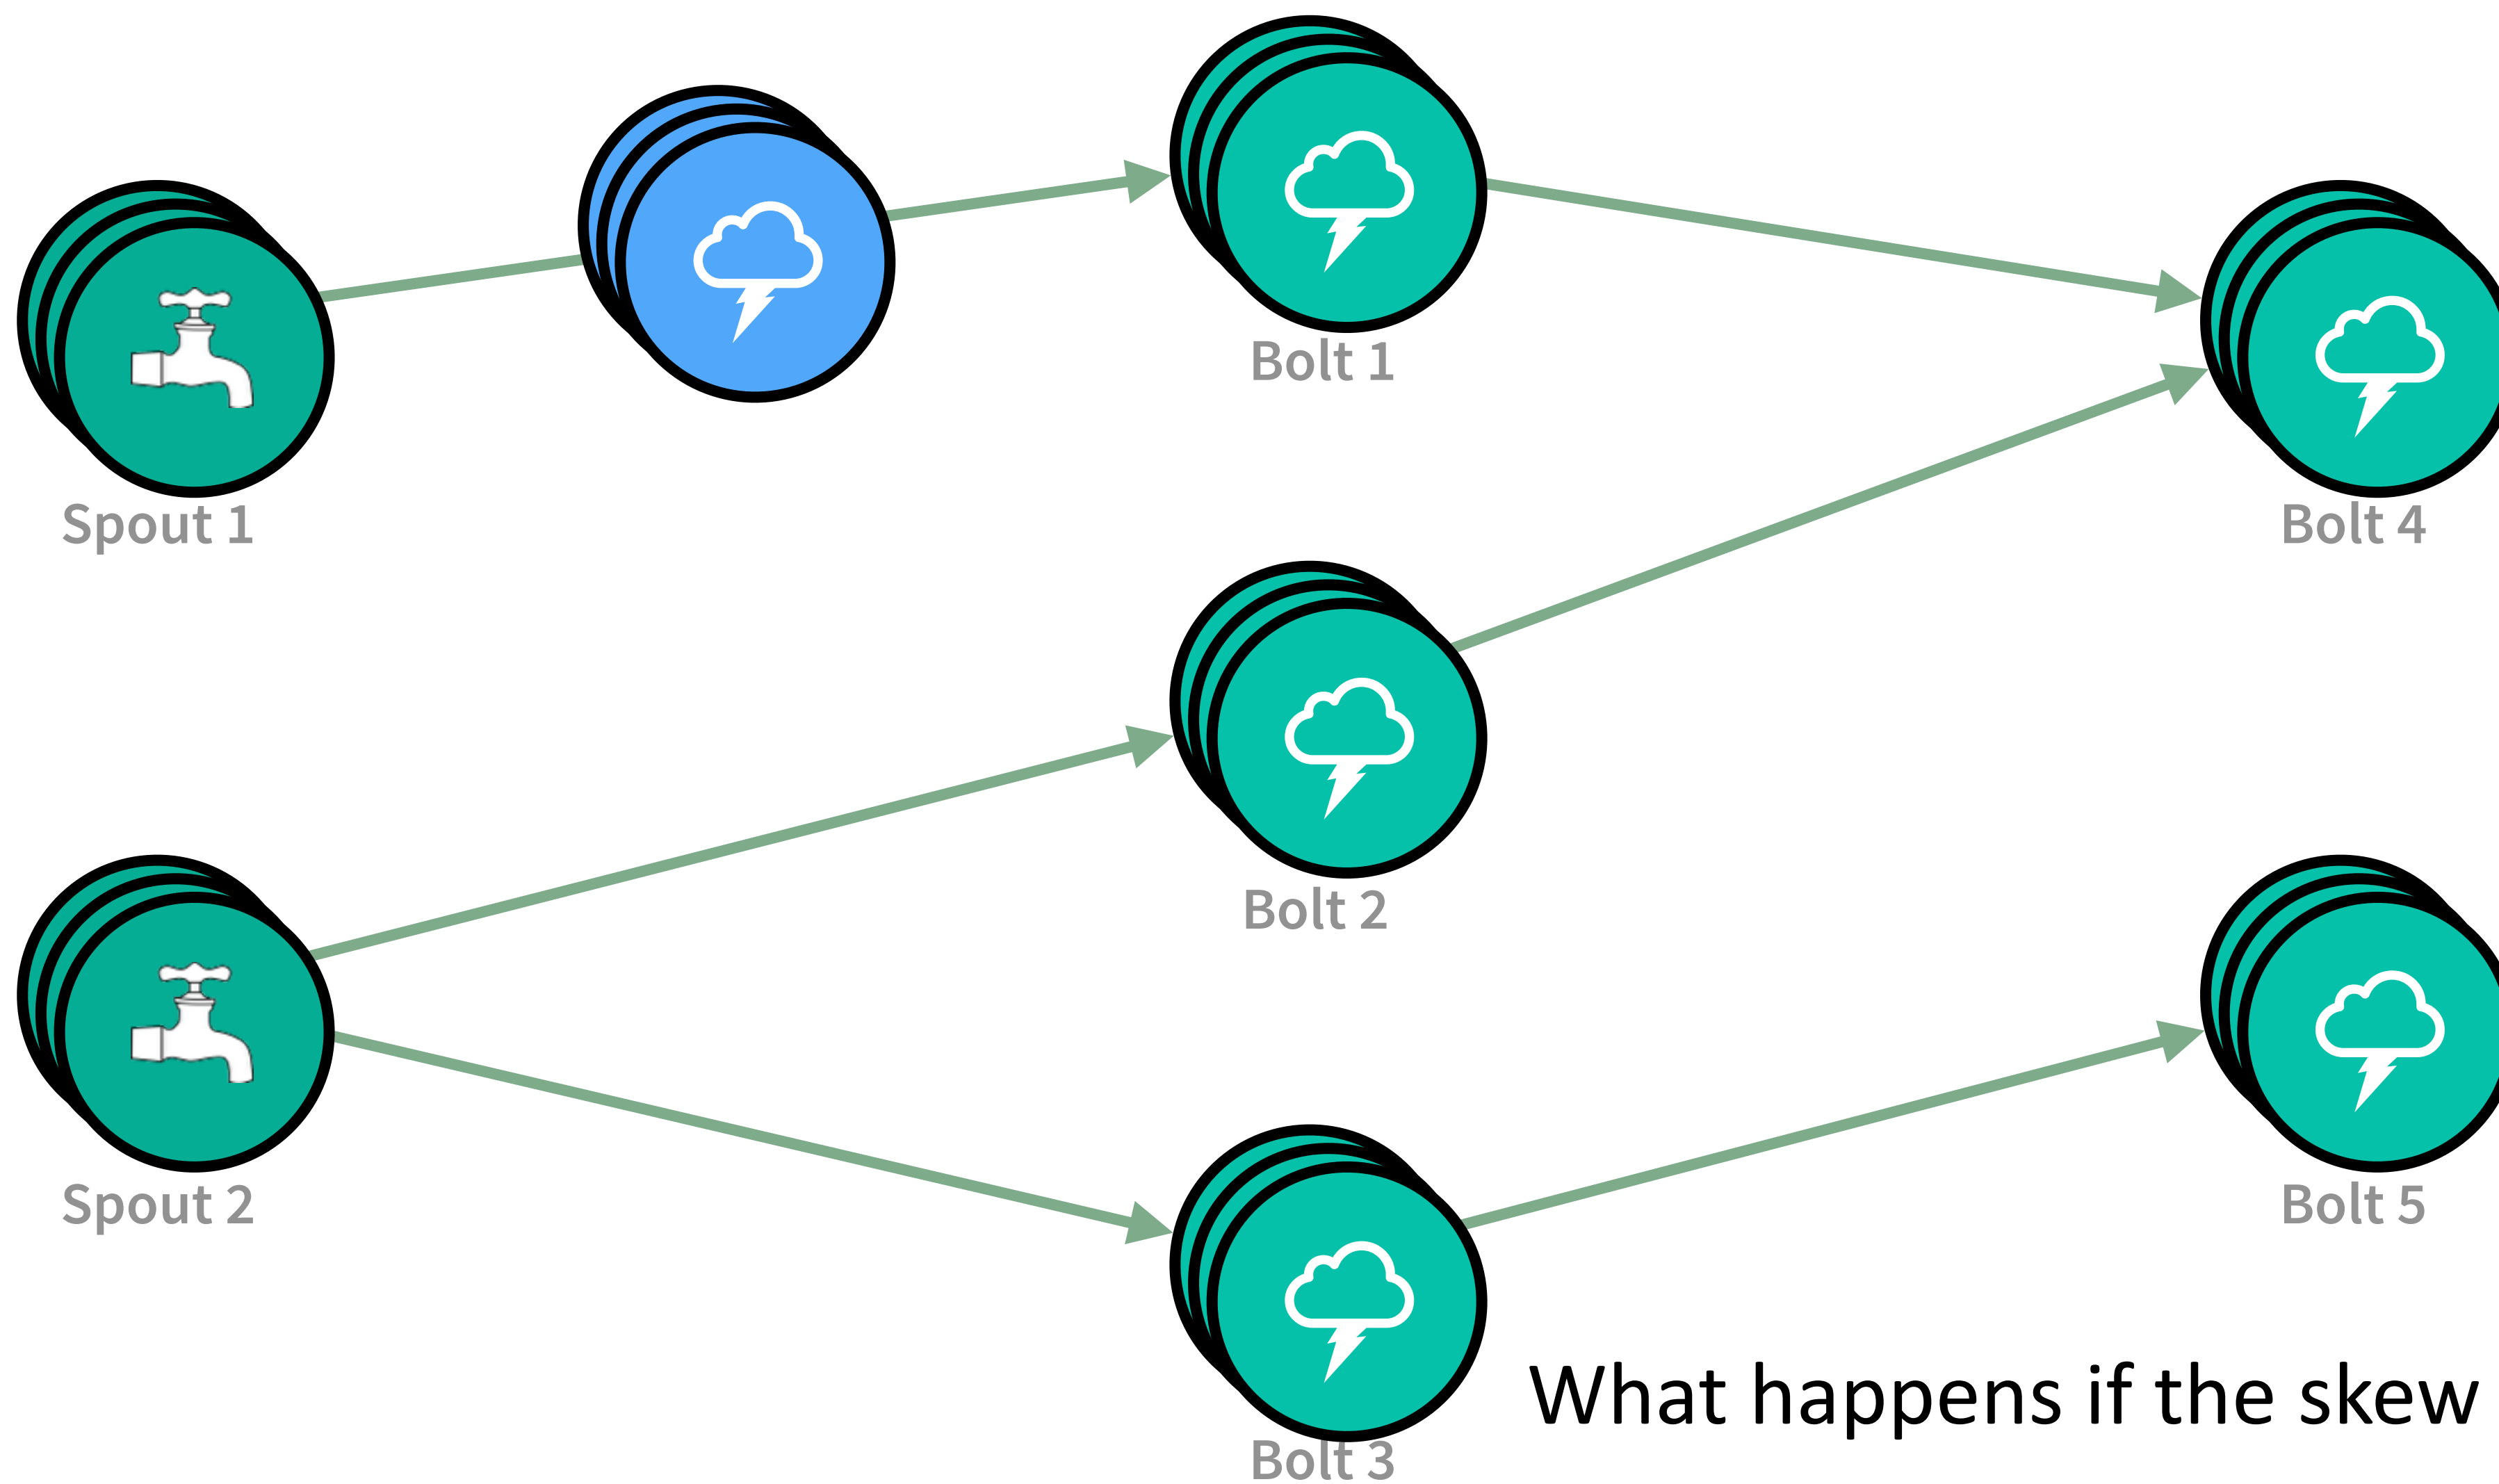
Data Skew



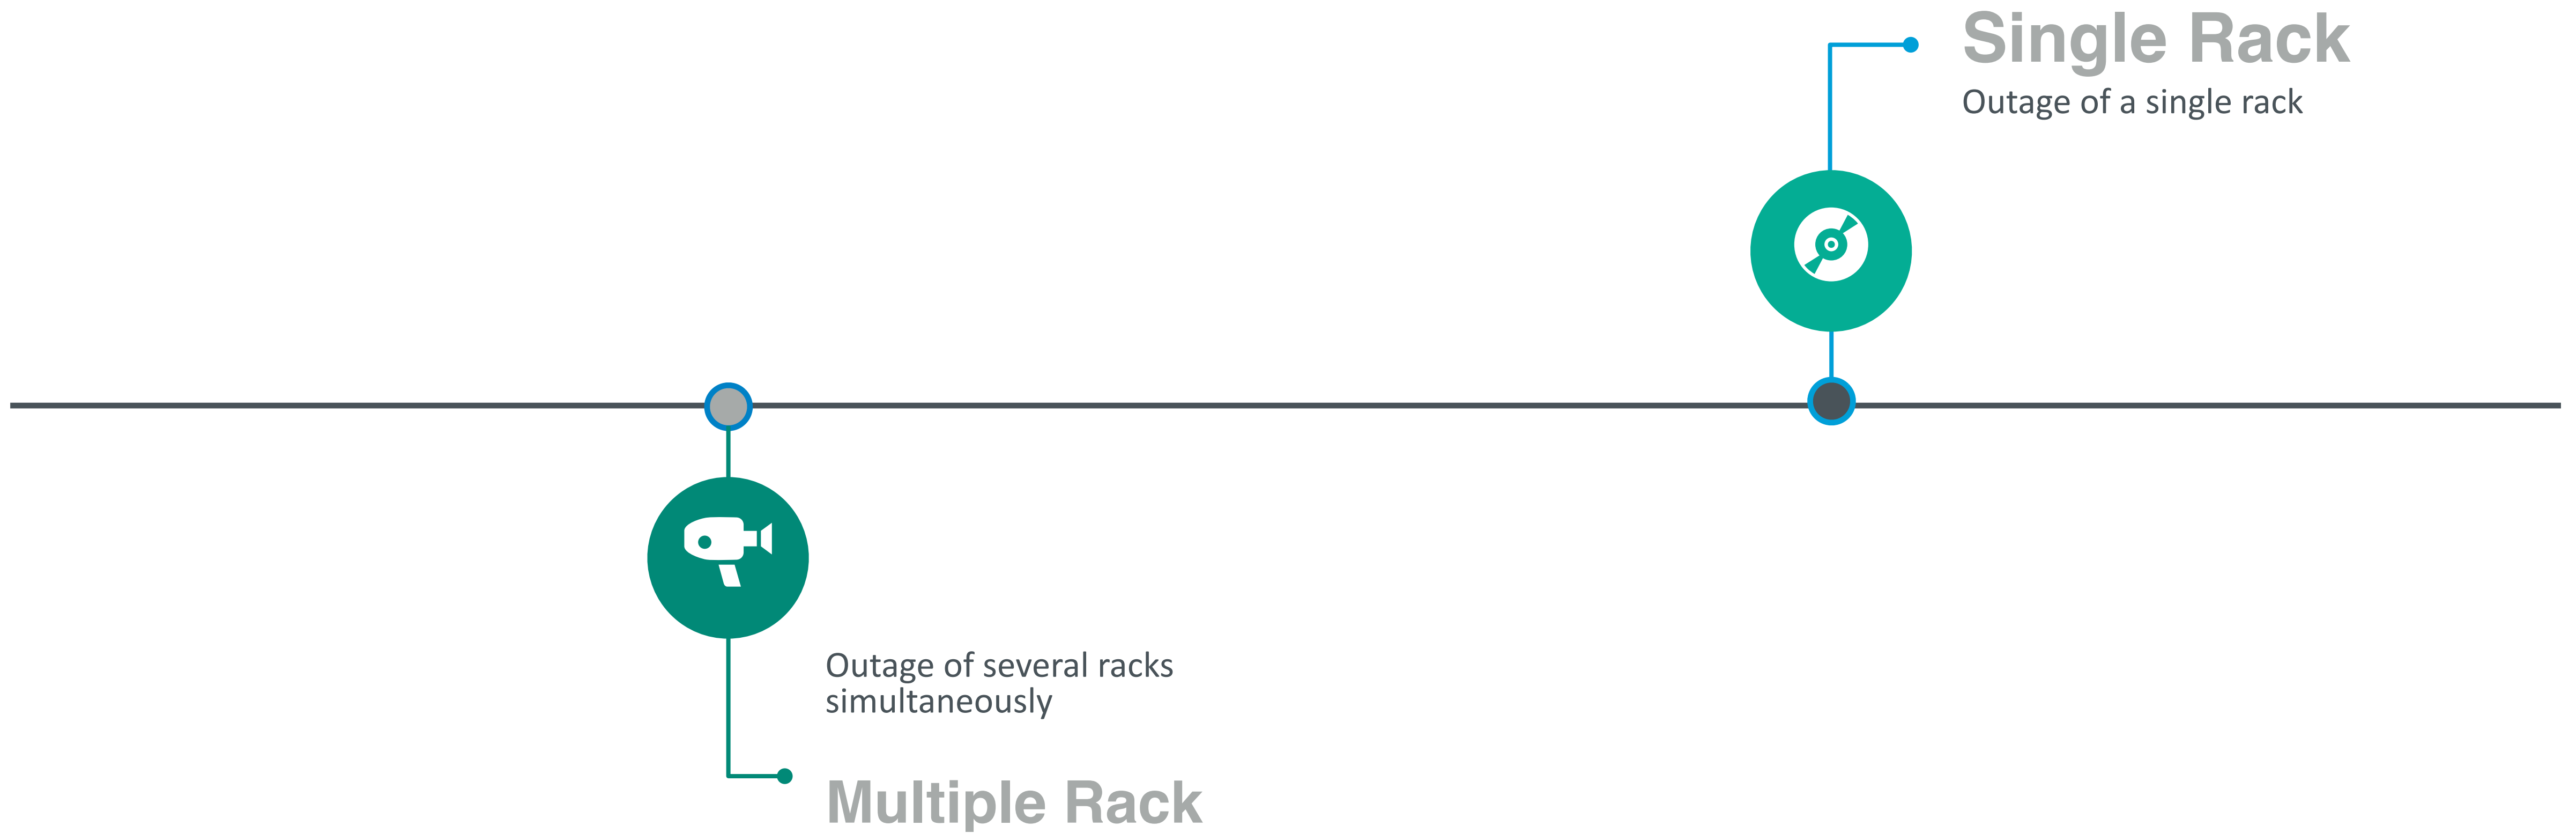
Data Skew - Multiple Keys



Data Skew - Single Key




Rack Failures




Rack Local vs Rack Diversity

Rack Local



Advantage of reduced network latency and high p2p bandwidth


Rack Diversity



Higher network latency and shared bandwidth


Single Rack Failure

Rack Local



Rack failure


If job is rack local entire job is down and needs to be rescheduled



Impact on real timeliness

Depends on the application

Single Rack Failure




Rack Diversity



Rack failure

If job is rack diverse impact is minimal



Impact on real timeliness

Is affected less - depending on how many containers are running

Conclusion



Self healing is important in Streaming systems



Key operational issues - Slow hosts, network and data skew



These requirements make the streaming systems more complex

Interested in Heron?

**HERON IS OPEN SOURCED
CONTRIBUTIONS ARE WELCOME!**

<https://github.com/twitter/heron>

<http://heronstreaming.io>

FOLLOW US @HERONSTREAMING

Any Questions ???

WHAT


WHY

WHERE

WHEN

WHO

HOW





Thanks For Listening

@karthikz

Detections and Resolutions

Detections

Proactive

Reactive

Affects real-time (depends on application tolerance)	Keeps up with real-time
Full manual intervention	N/A

Lazy

Instant

Resolutions

