Twitter Heron
Streaming at Scale

SANJEEV KULKARNI

@SANJEEVRK

#TwitterHeron

Sanjeev Kulkarni, Nikunj Bhagat, Maosong Fu, Vikas Kedigehalli,
Christopher Kellogg, Sailesh Mittal, Jignesh Patel, Siddarth Taneja

TALK OUTLINE

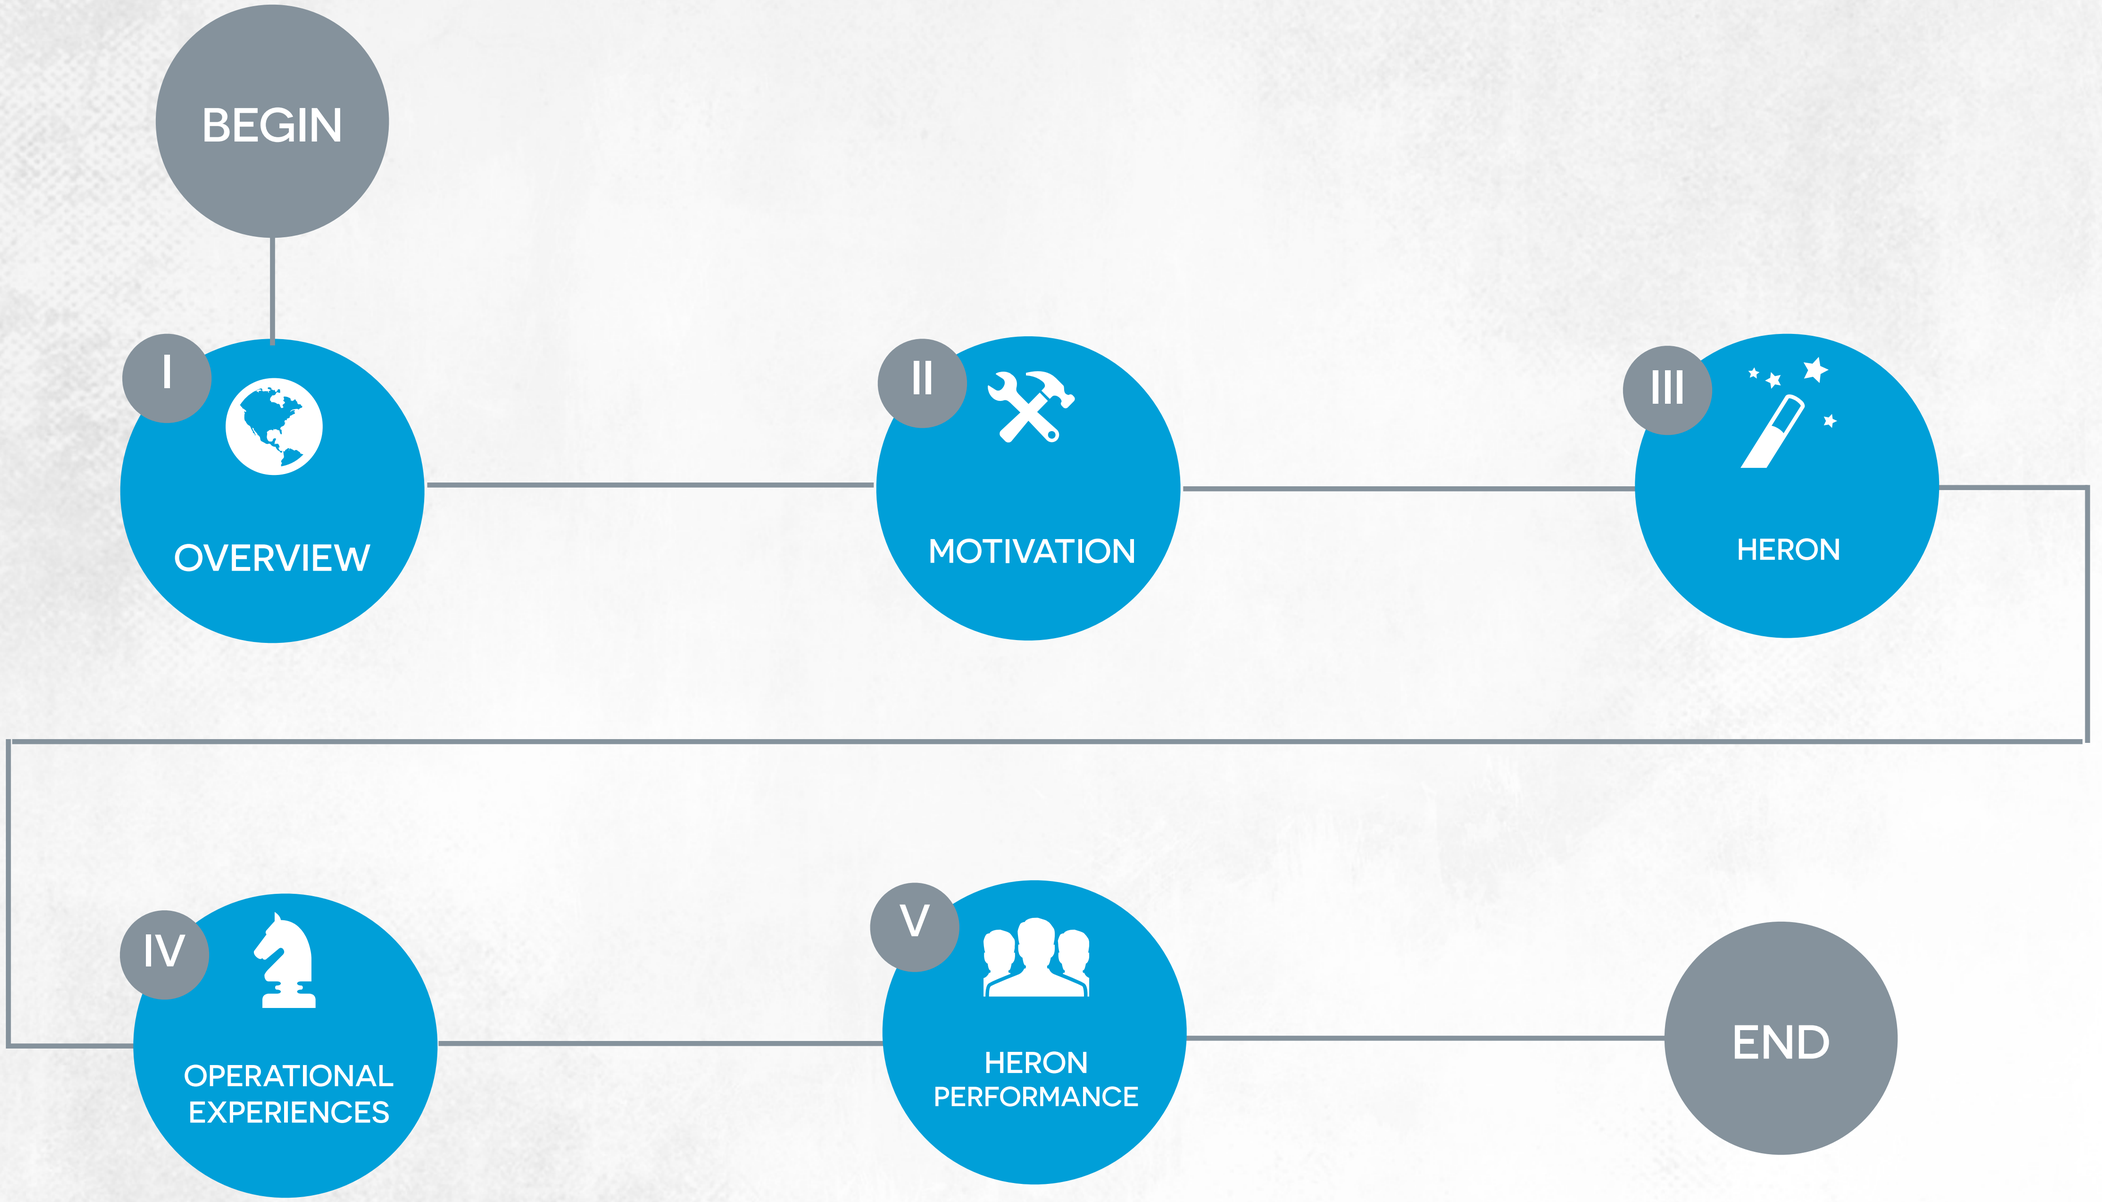

OVERVIEW

TWITTER IS REAL TIME

REAL TIME TRENDS

Emerging break out trends in Twitter (in the form #hashtags)

REAL TIME CONVERSATIONS

Real time sports conversations related with a topic (recent goal or touchdown)

REAL TIME RECOMMENDATIONS

Real time product recommendations based on your behavior & profile

REAL TIME SEARCH

Real time search of tweets

ANALYZING BILLIONS OF EVENTS IN REAL TIME IS A CHALLENGE!

TWITTER STORM

Streaming platform for analyzing **realtime data** as they arrive, so you can react to data **as it happens**.

GUARANTEED MESSAGE PROCESSING

HORIZONTAL SCALABILITY

ROBUST FAULT TOLERANCE

CONCISE CODE- FOCUS ON LOGIC

STORM TERMINOLOGY

TOPOLOGY

Directed acyclic graph

Vertices=computation, and edges=streams of data tuples

SPOUTS

Sources of data tuples for the topology

Examples - Kafka/Kestrel/MySQL/Postgres

BOLTS

Process incoming tuples and emit outgoing tuples

Examples - filtering/aggregation/join/arbitrary function

STORM TOPOLOGY

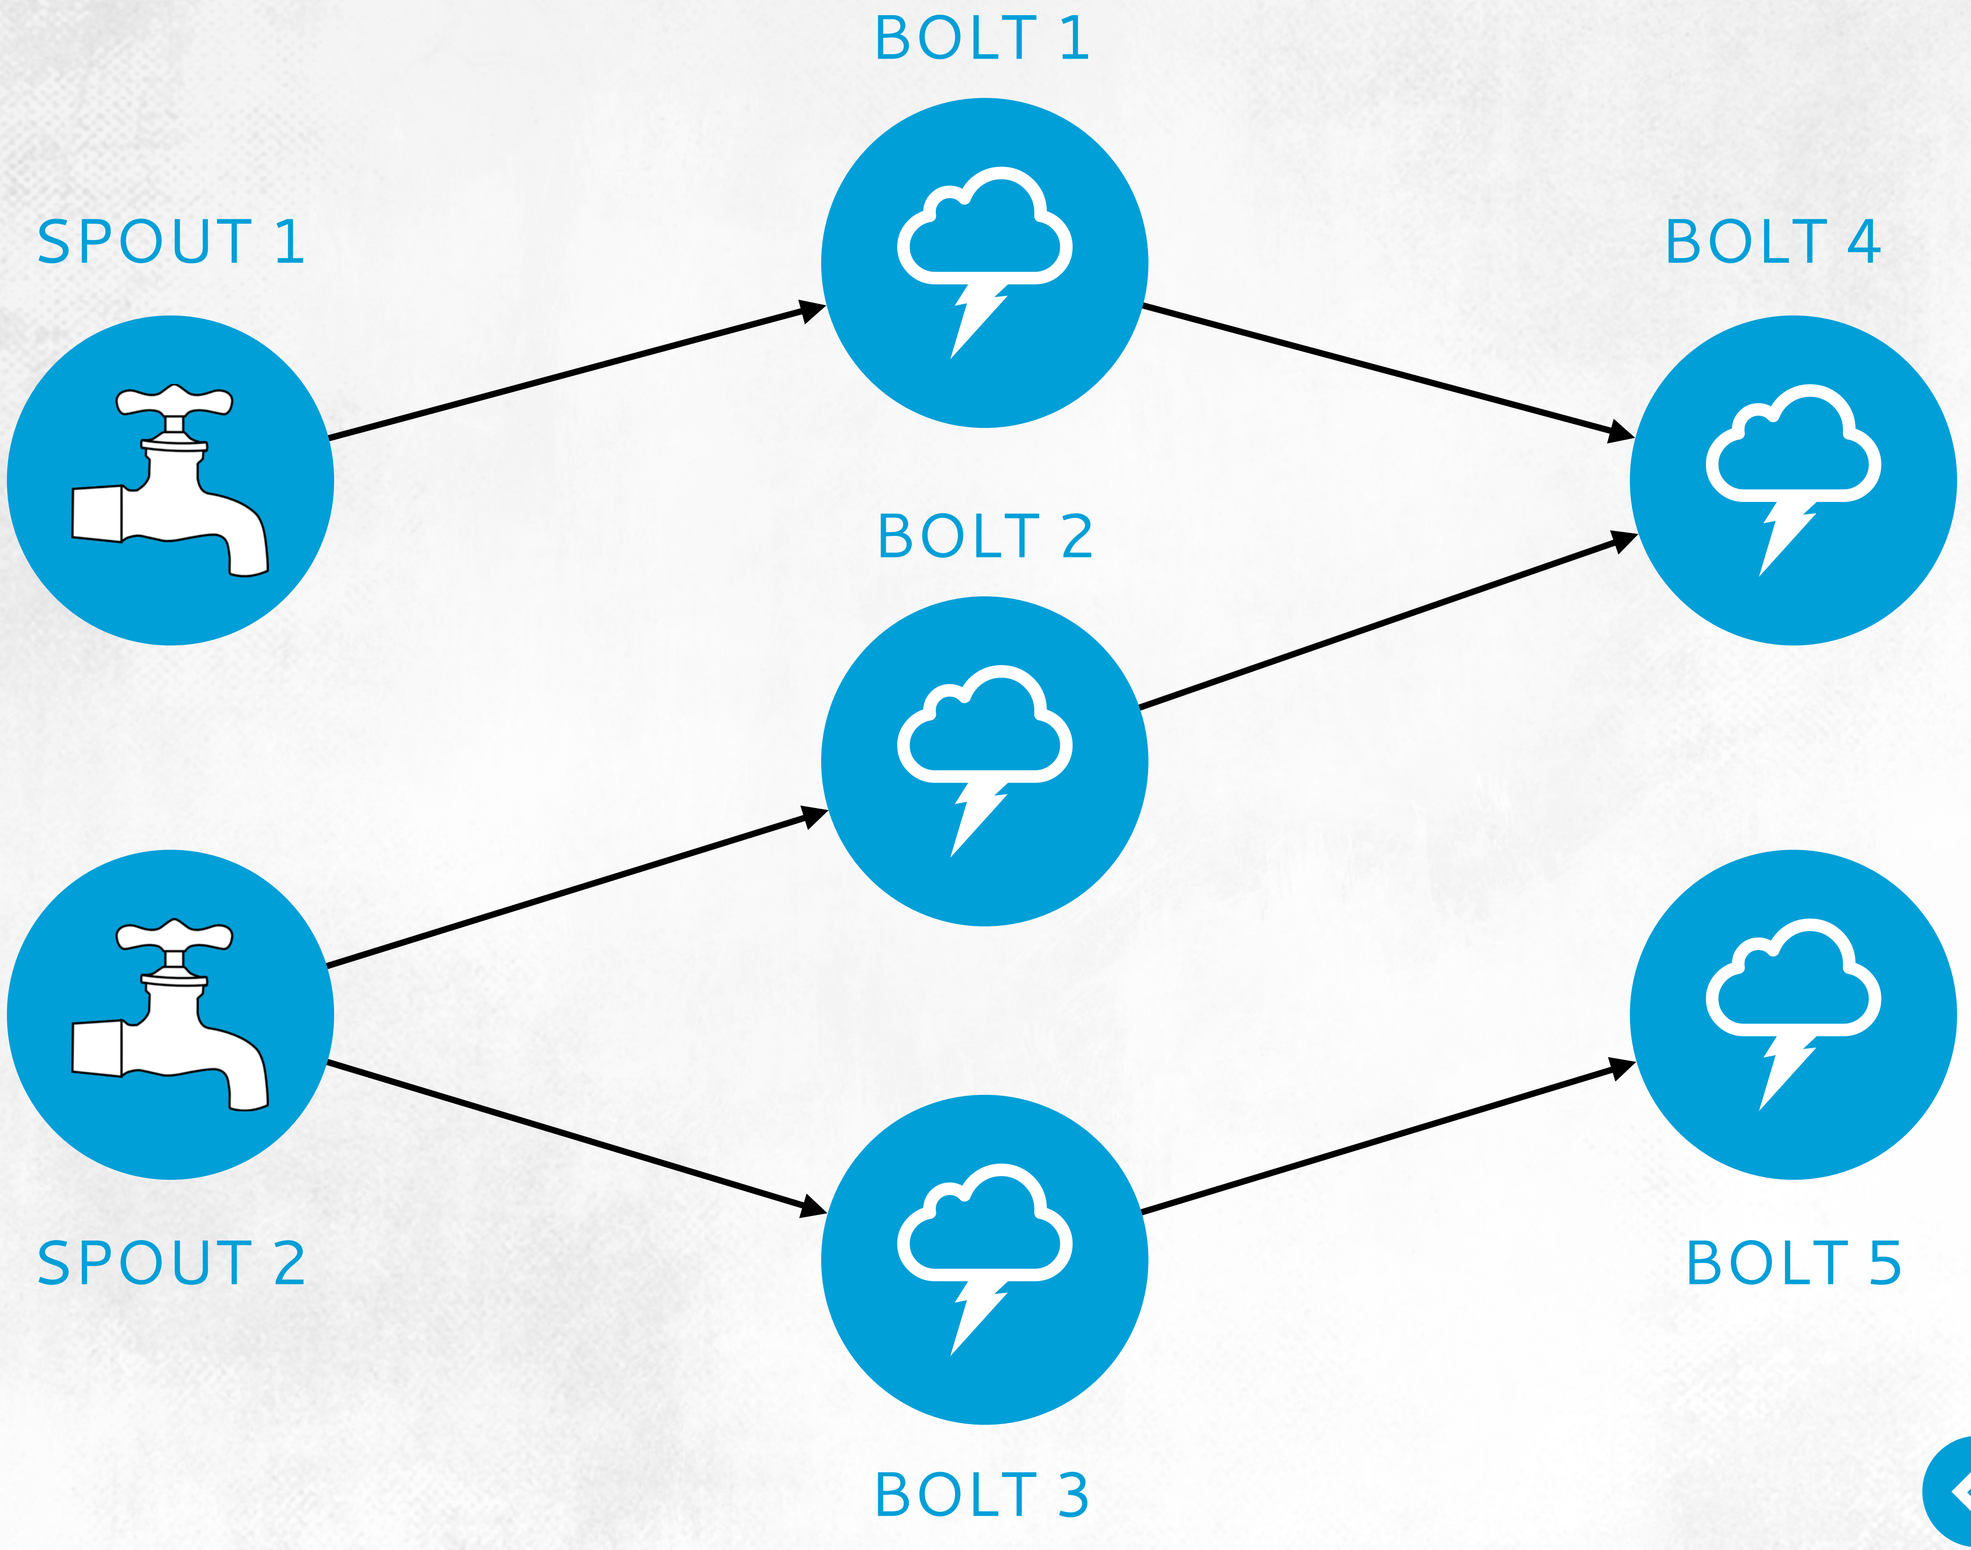

MOTIVATION

STORM ARCHITECTURE

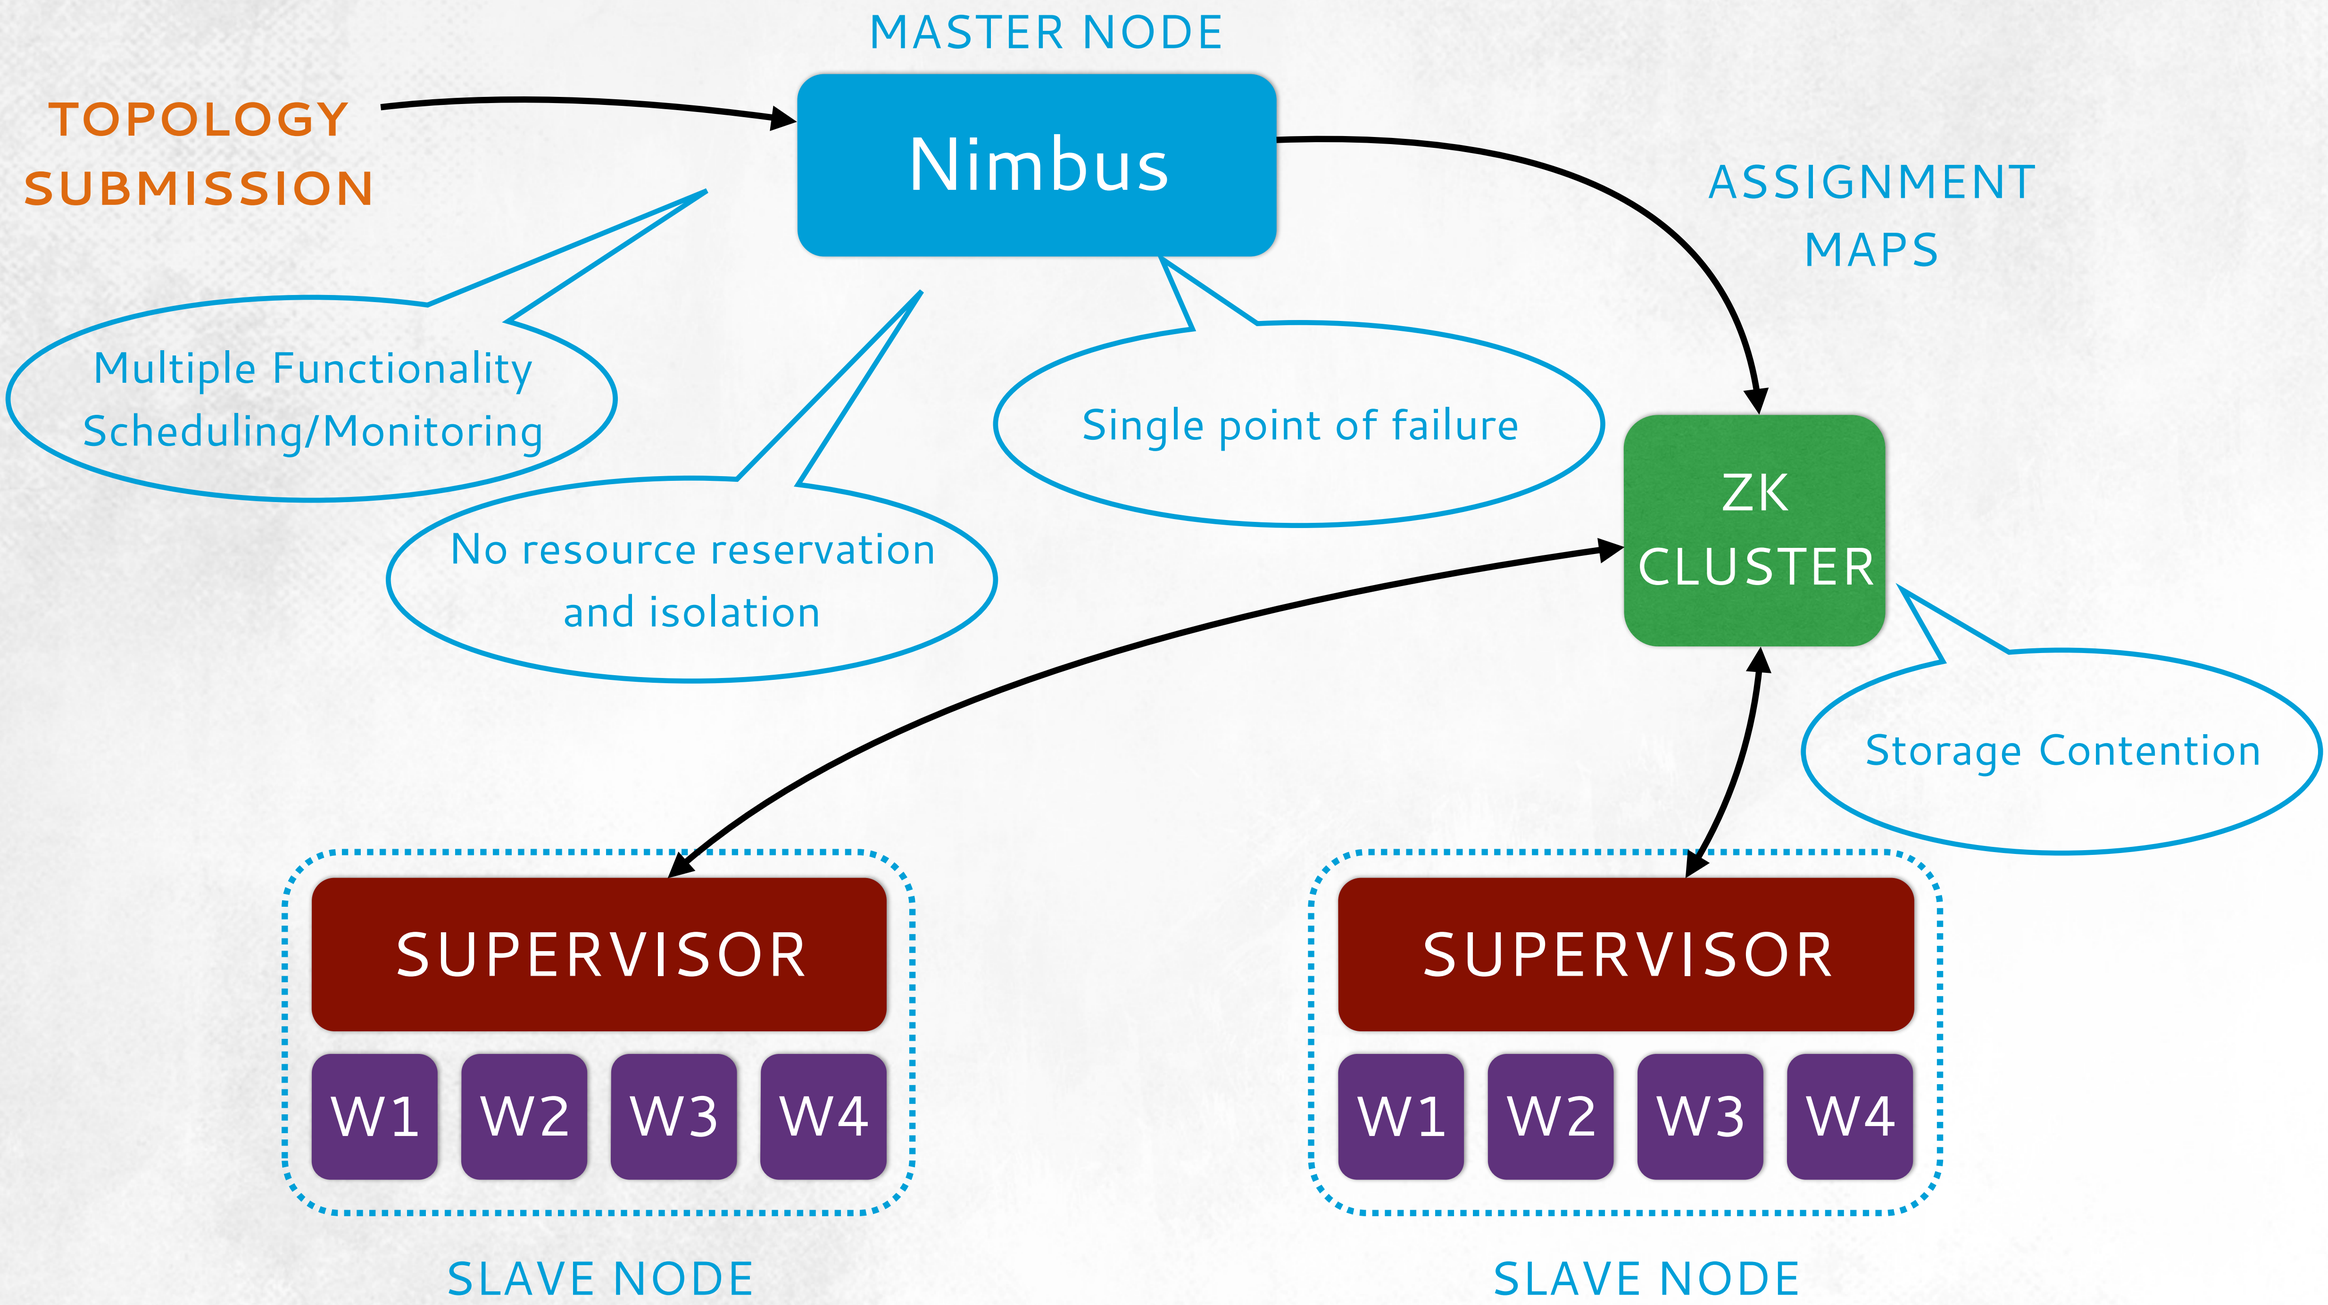

STORM WORKER

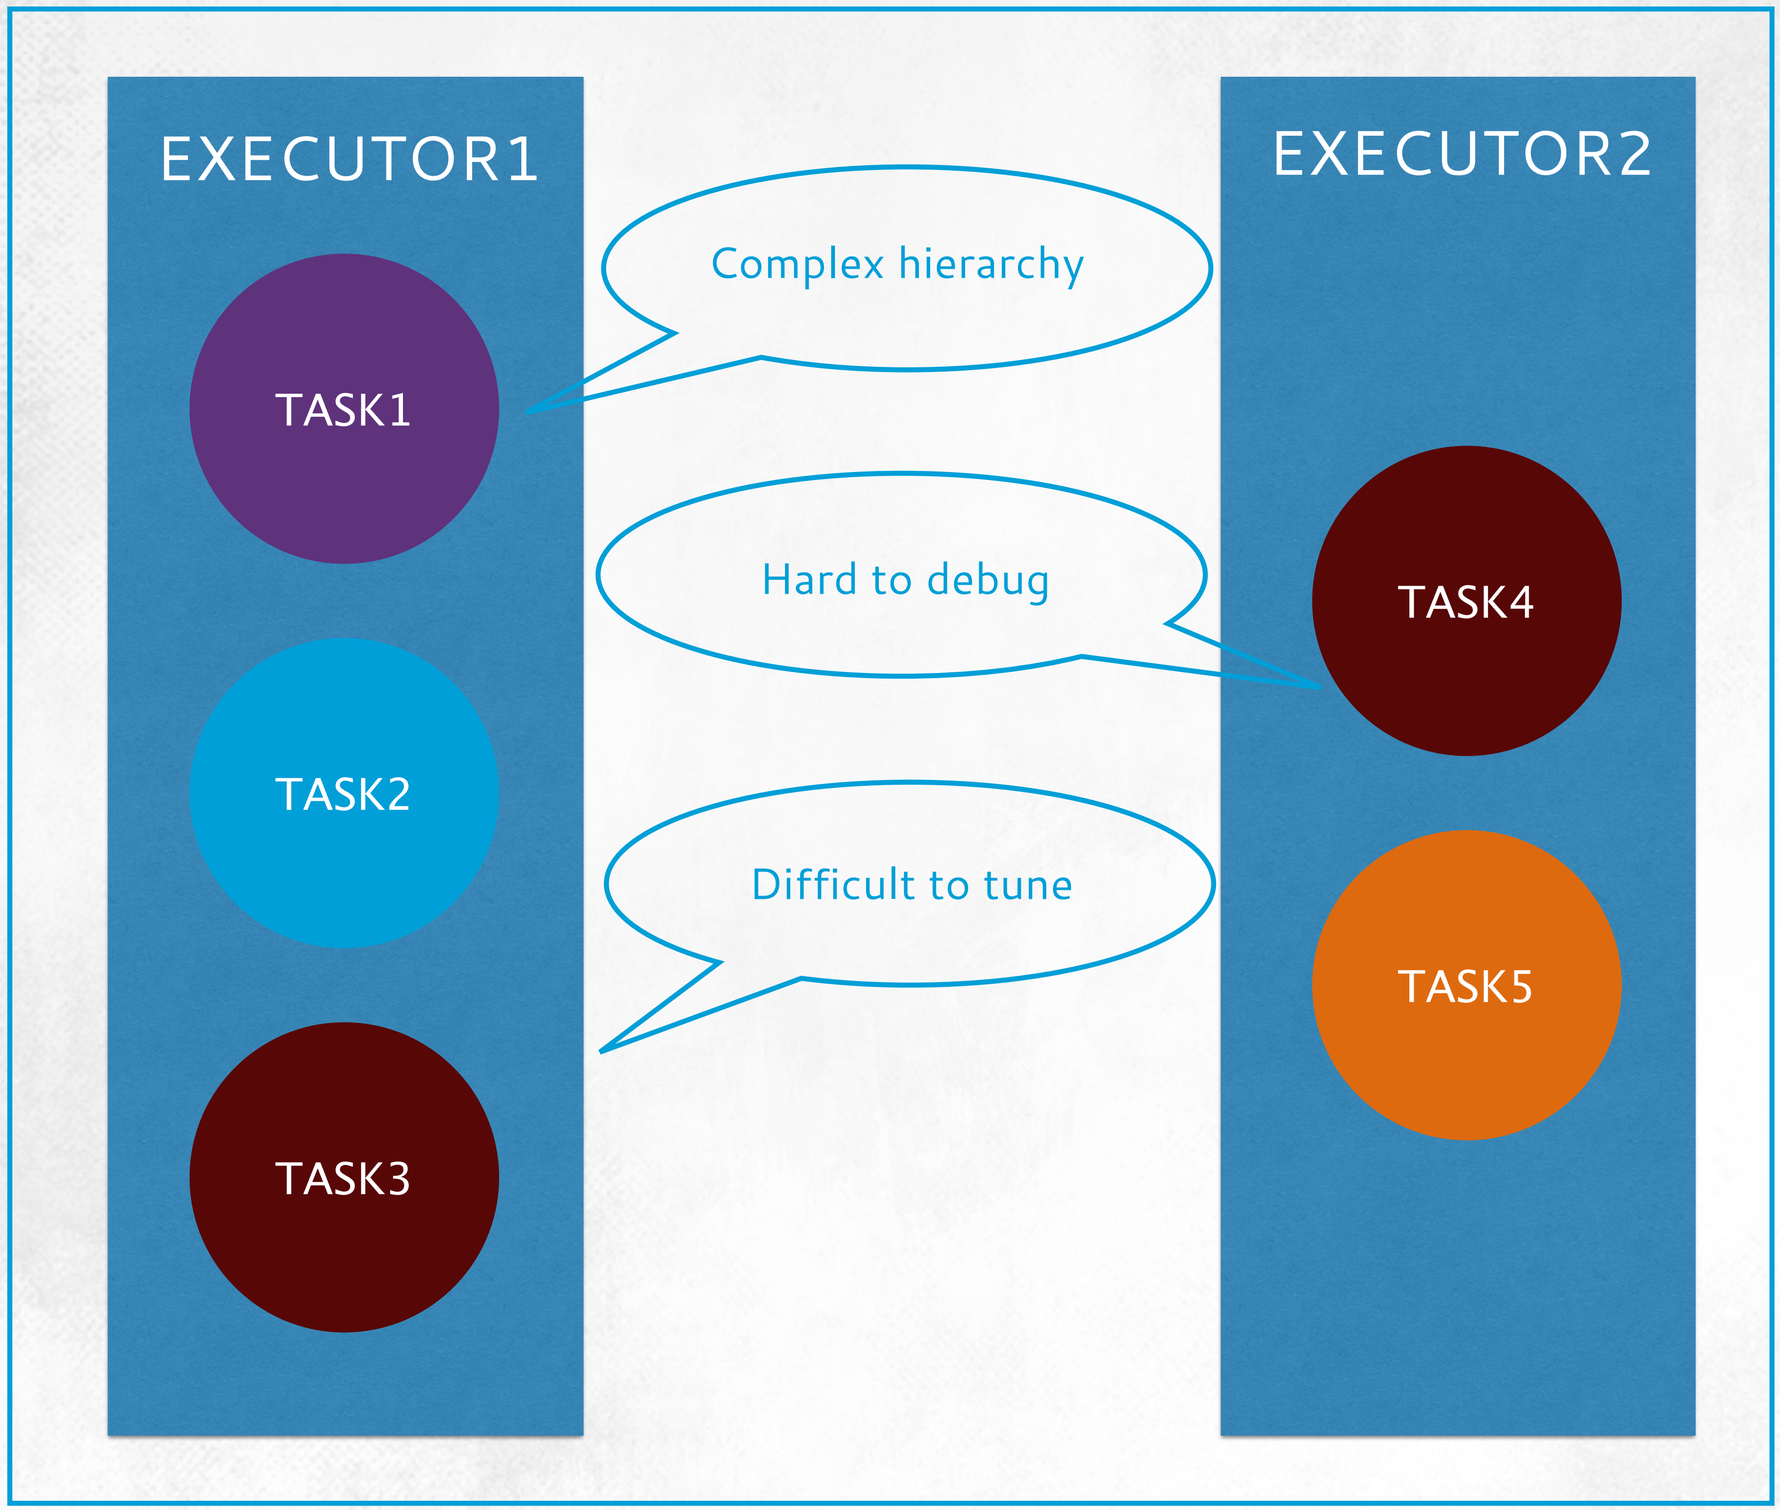

DATA FLOW IN STORM WORKERS

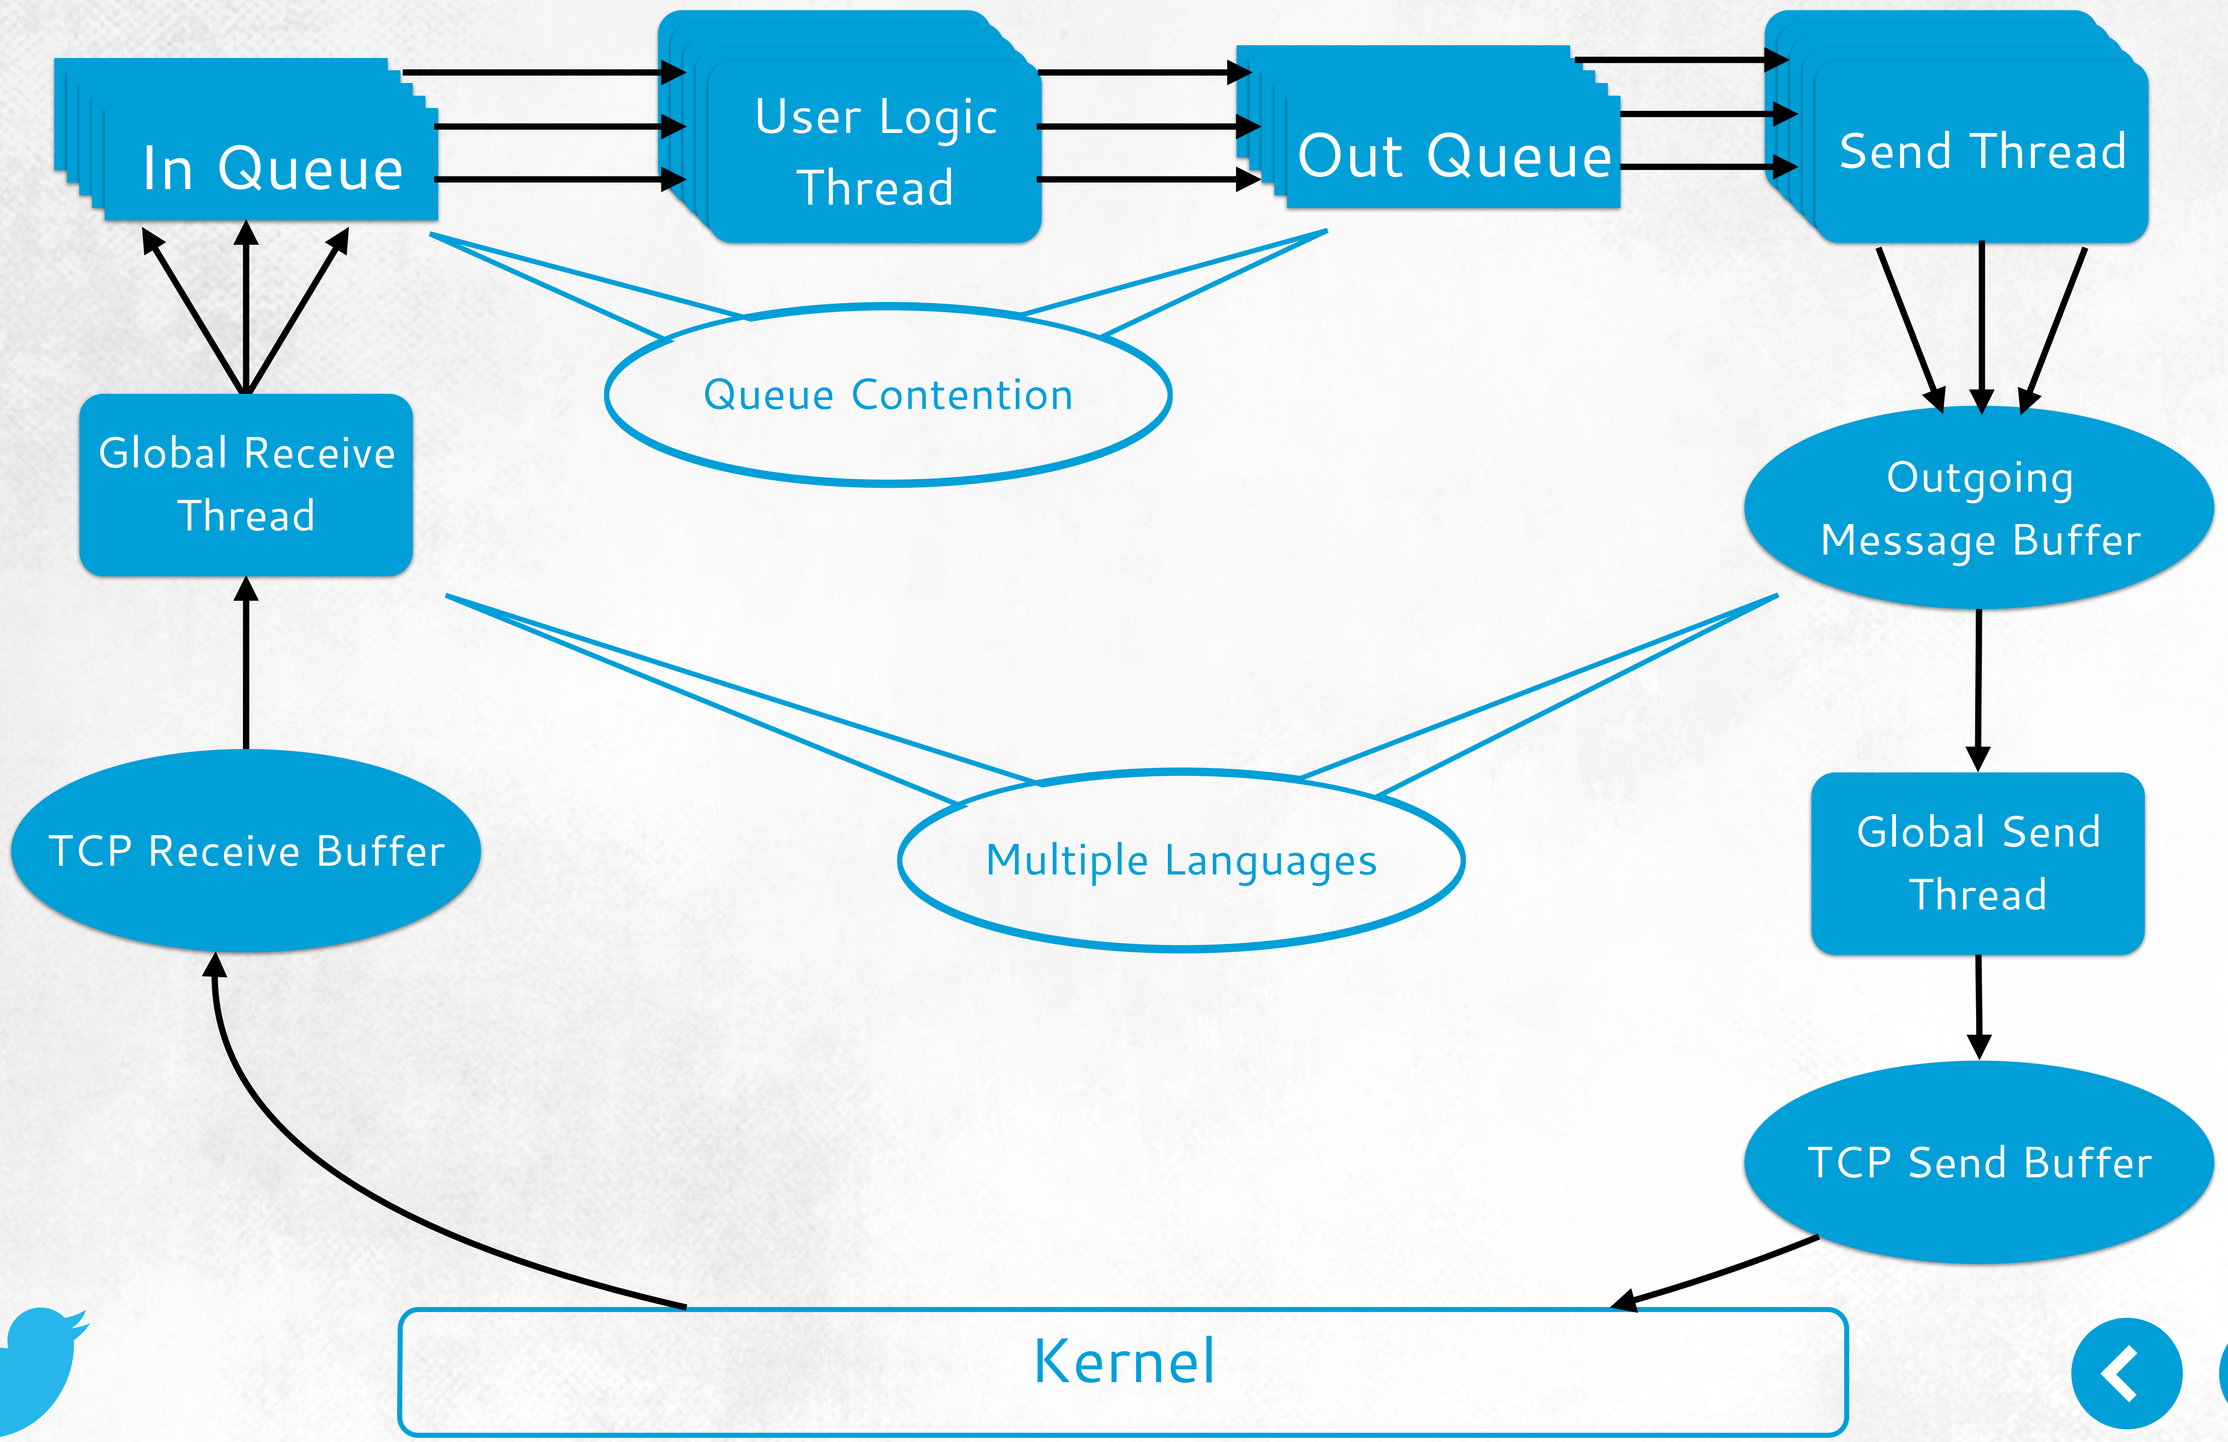

STORM ISSUES

LACK OF BACK PRESSURE

Drops tuples unpredictably

EFFICIENCY

Serialization program consumes 75 cores at 30% CPU

Topology consumes 600 cores at 20-30% CPU

NO BATCHING

Tuple oriented system - implicit batching by 0MQ

HERON

HERON DESIGN GOALS

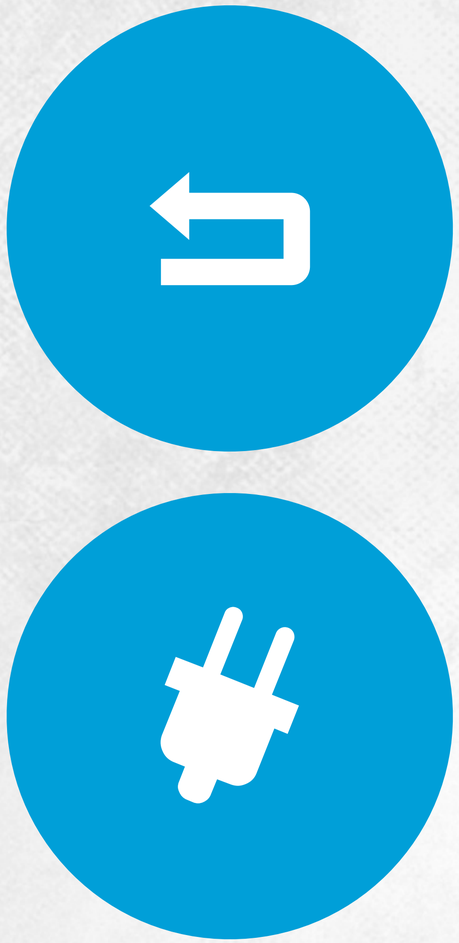

FULLY API COMPATIBLE WITH STORM

Topology/Spouts/Bolts

TASK ISOLATION

Ease of debug ability/resource isolation/profiling

USE OF MAIN STREAM LANGUAGES

C++/JAVA/Python

HERON ARCHITECTURE

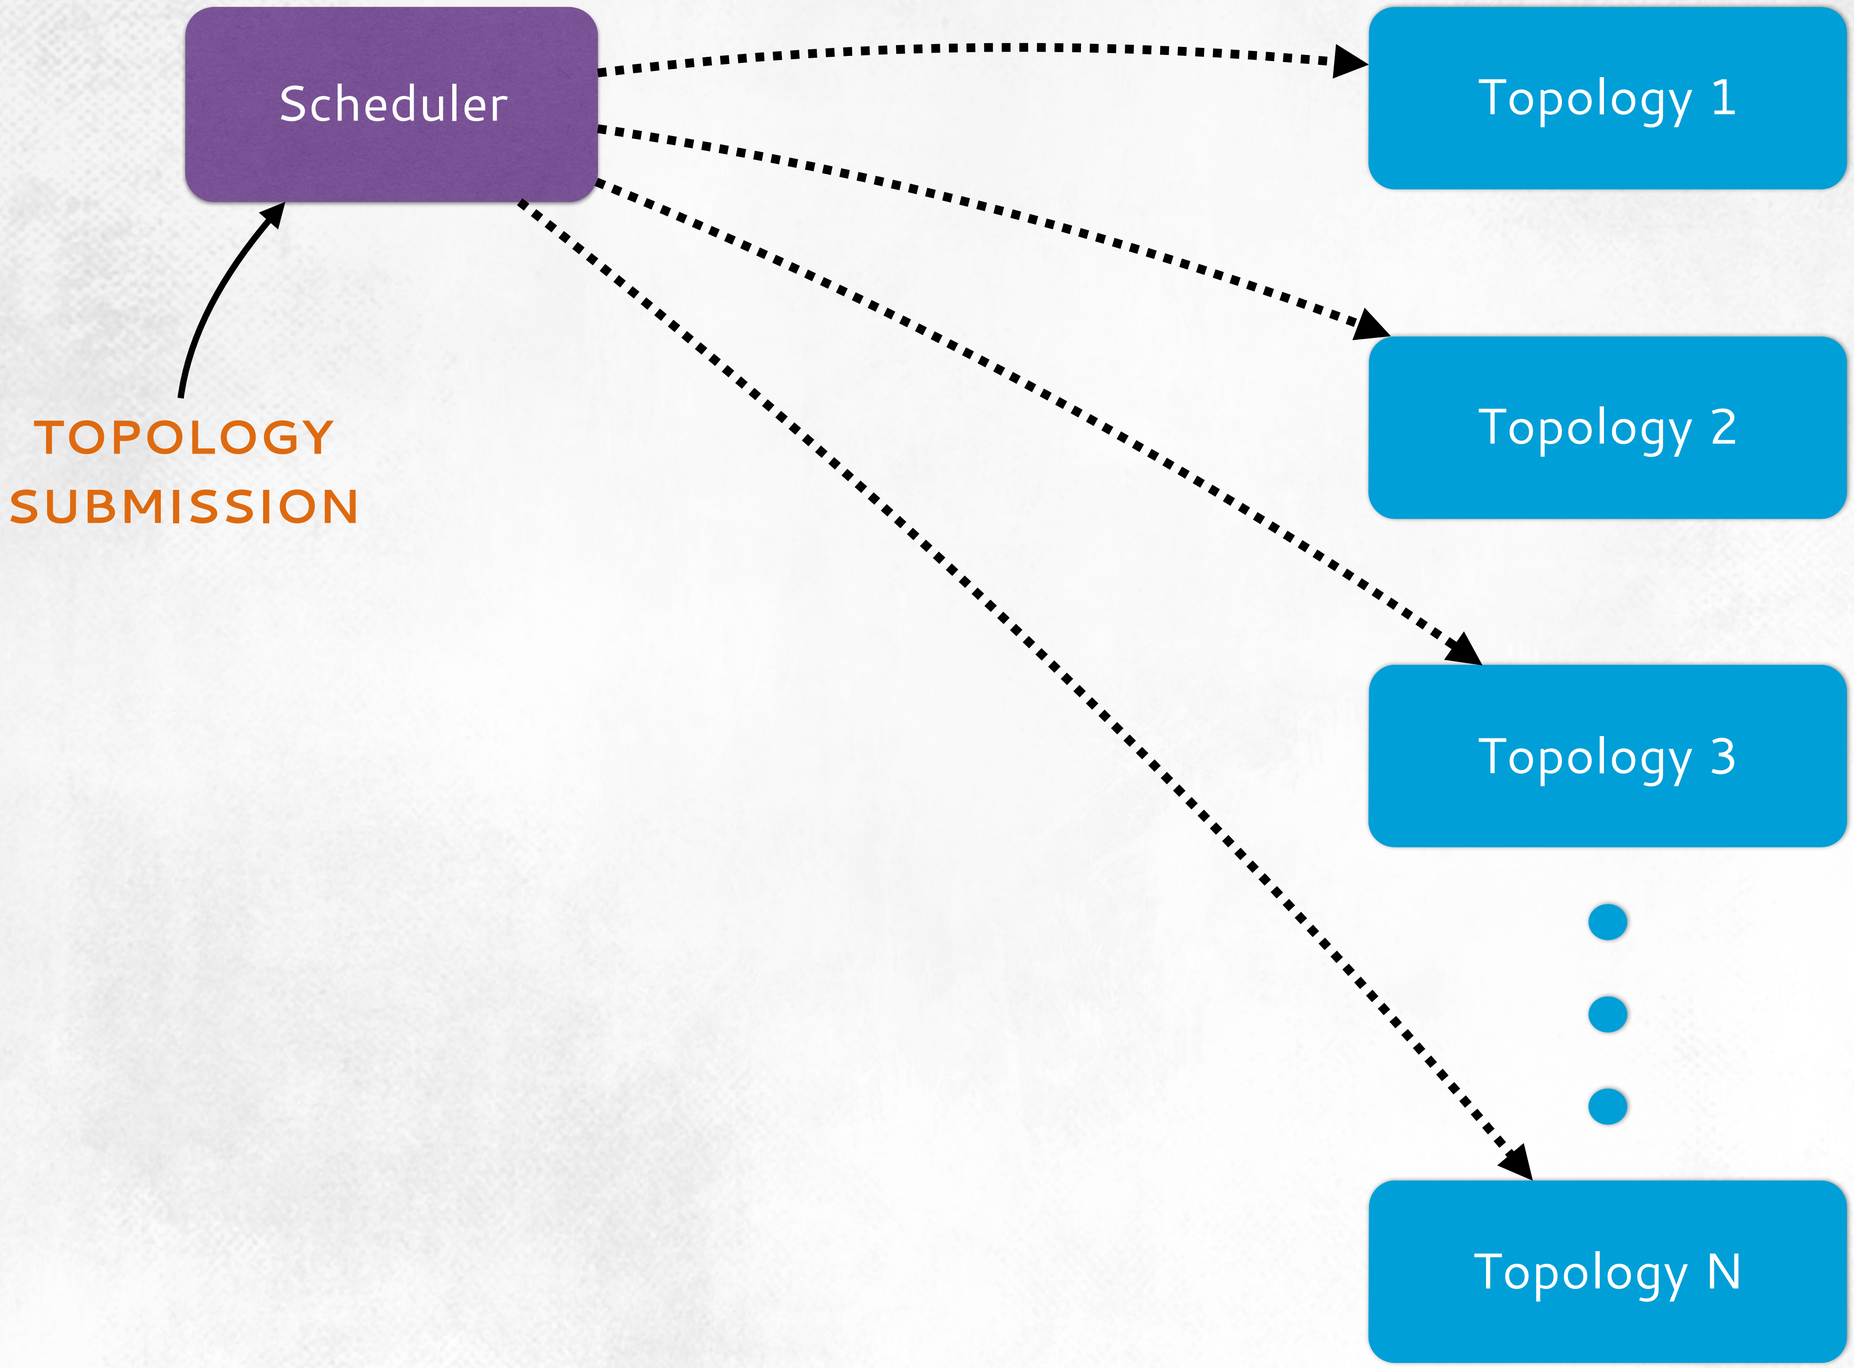

TOPOLOGY ARCHITECTURE

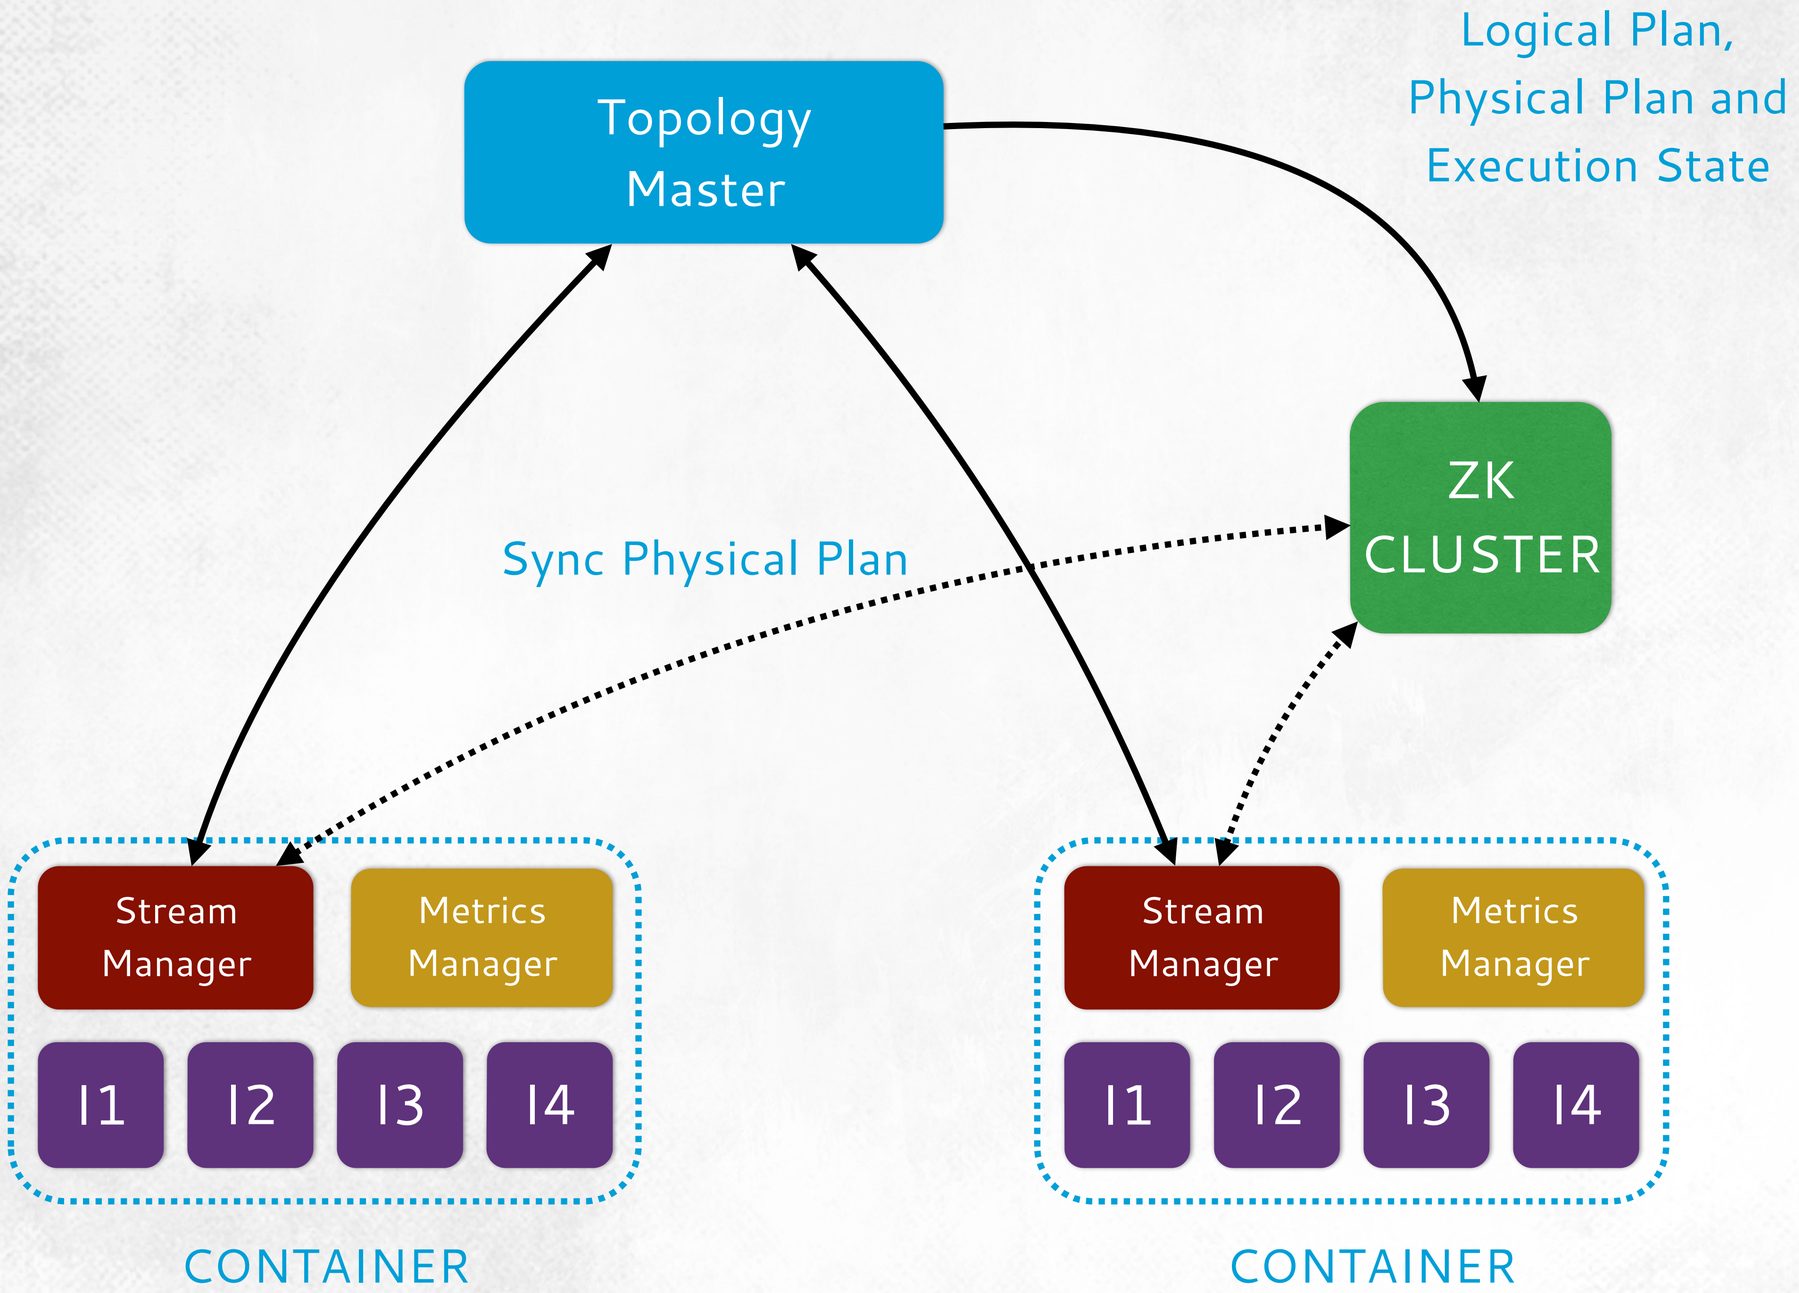

TOPOLOGY MASTER

Solely responsible for the entire topology

ASSIGNS ROLE

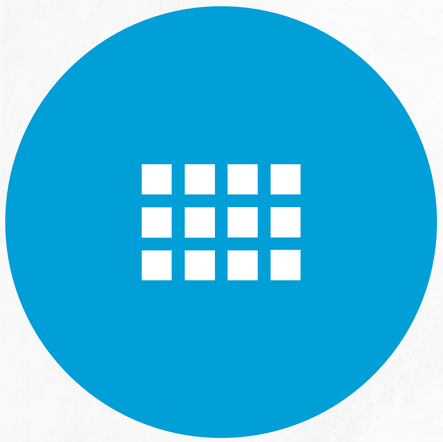

MONITORING

METRICS

STREAM MANAGER

Routing Engine

ROUTES TUPLES

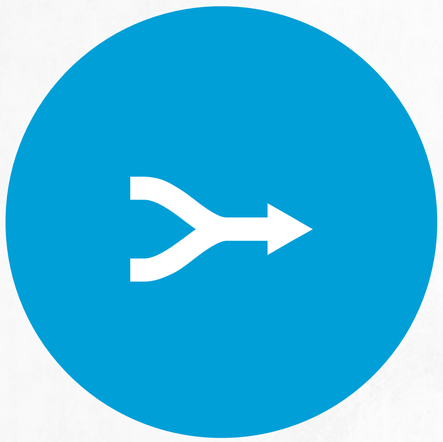

BACK PRESSURE

ACK MGMT

STREAM MANAGER

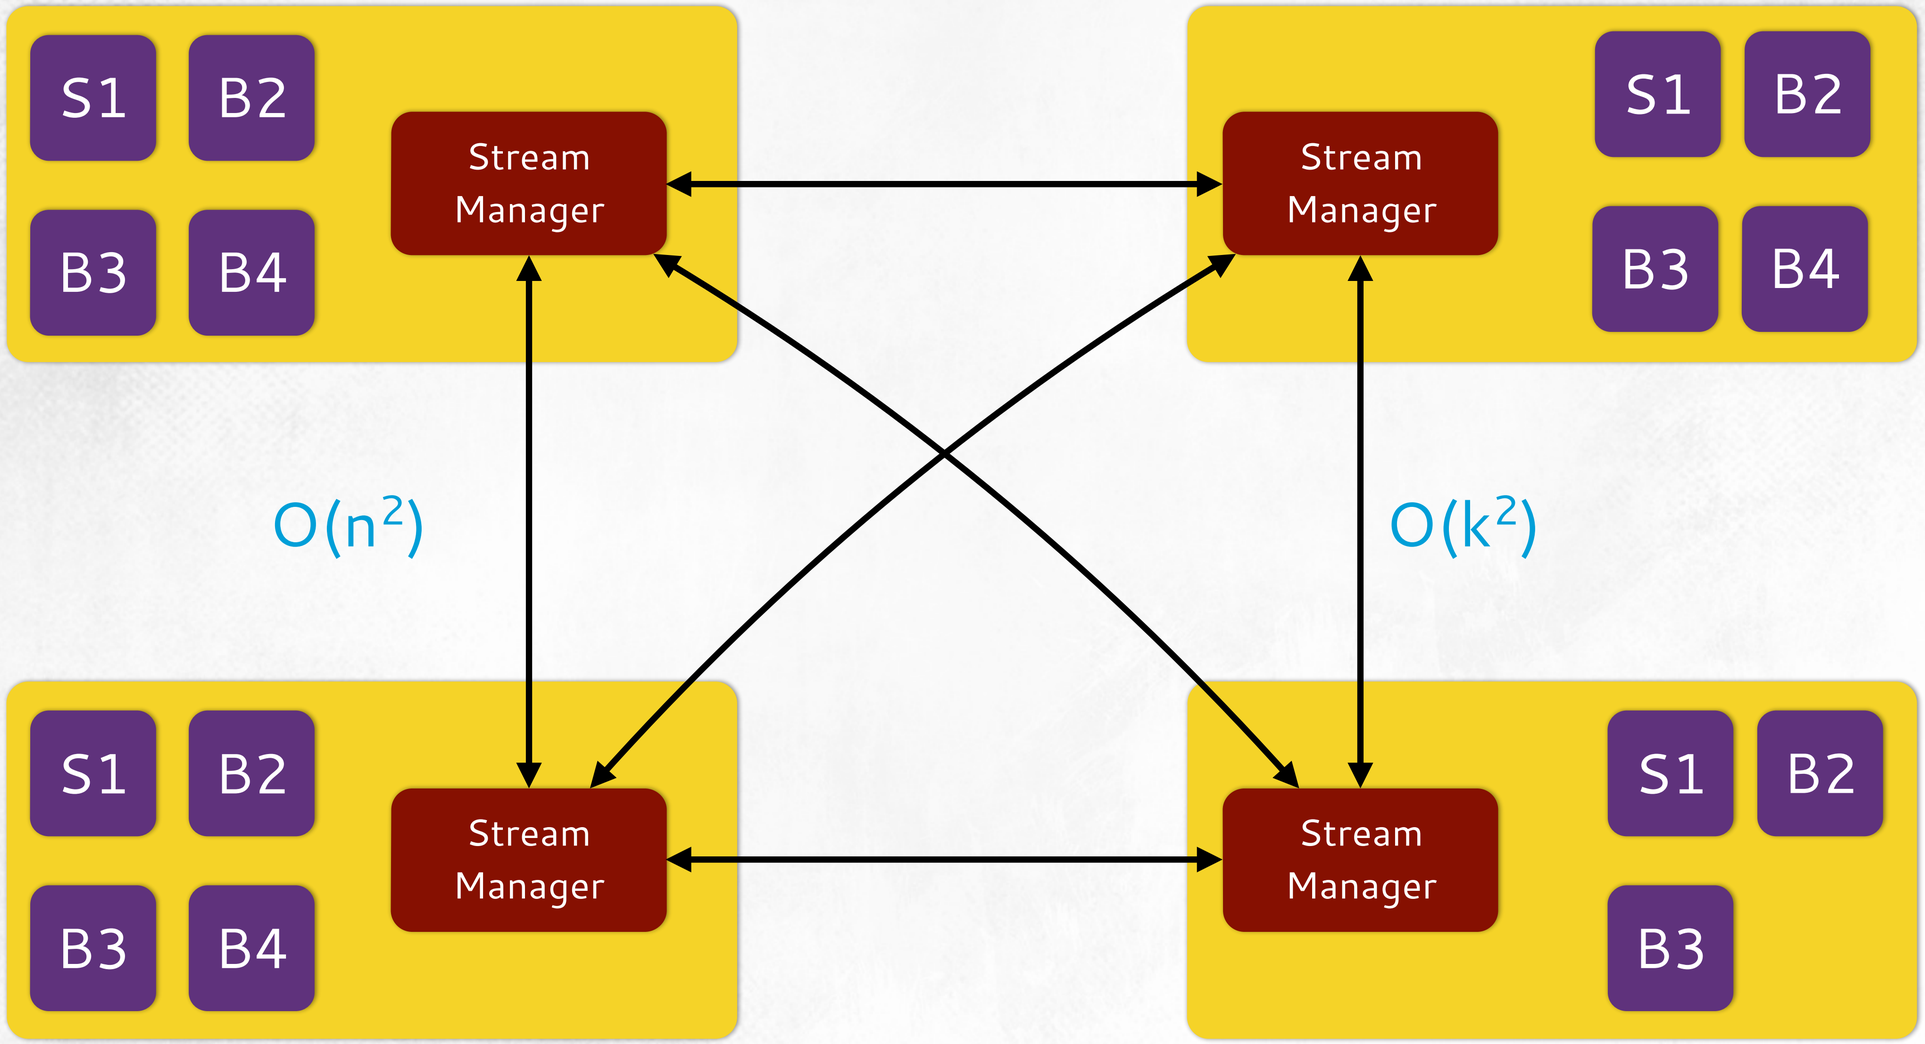

STREAM MANAGER

tcp back pressure

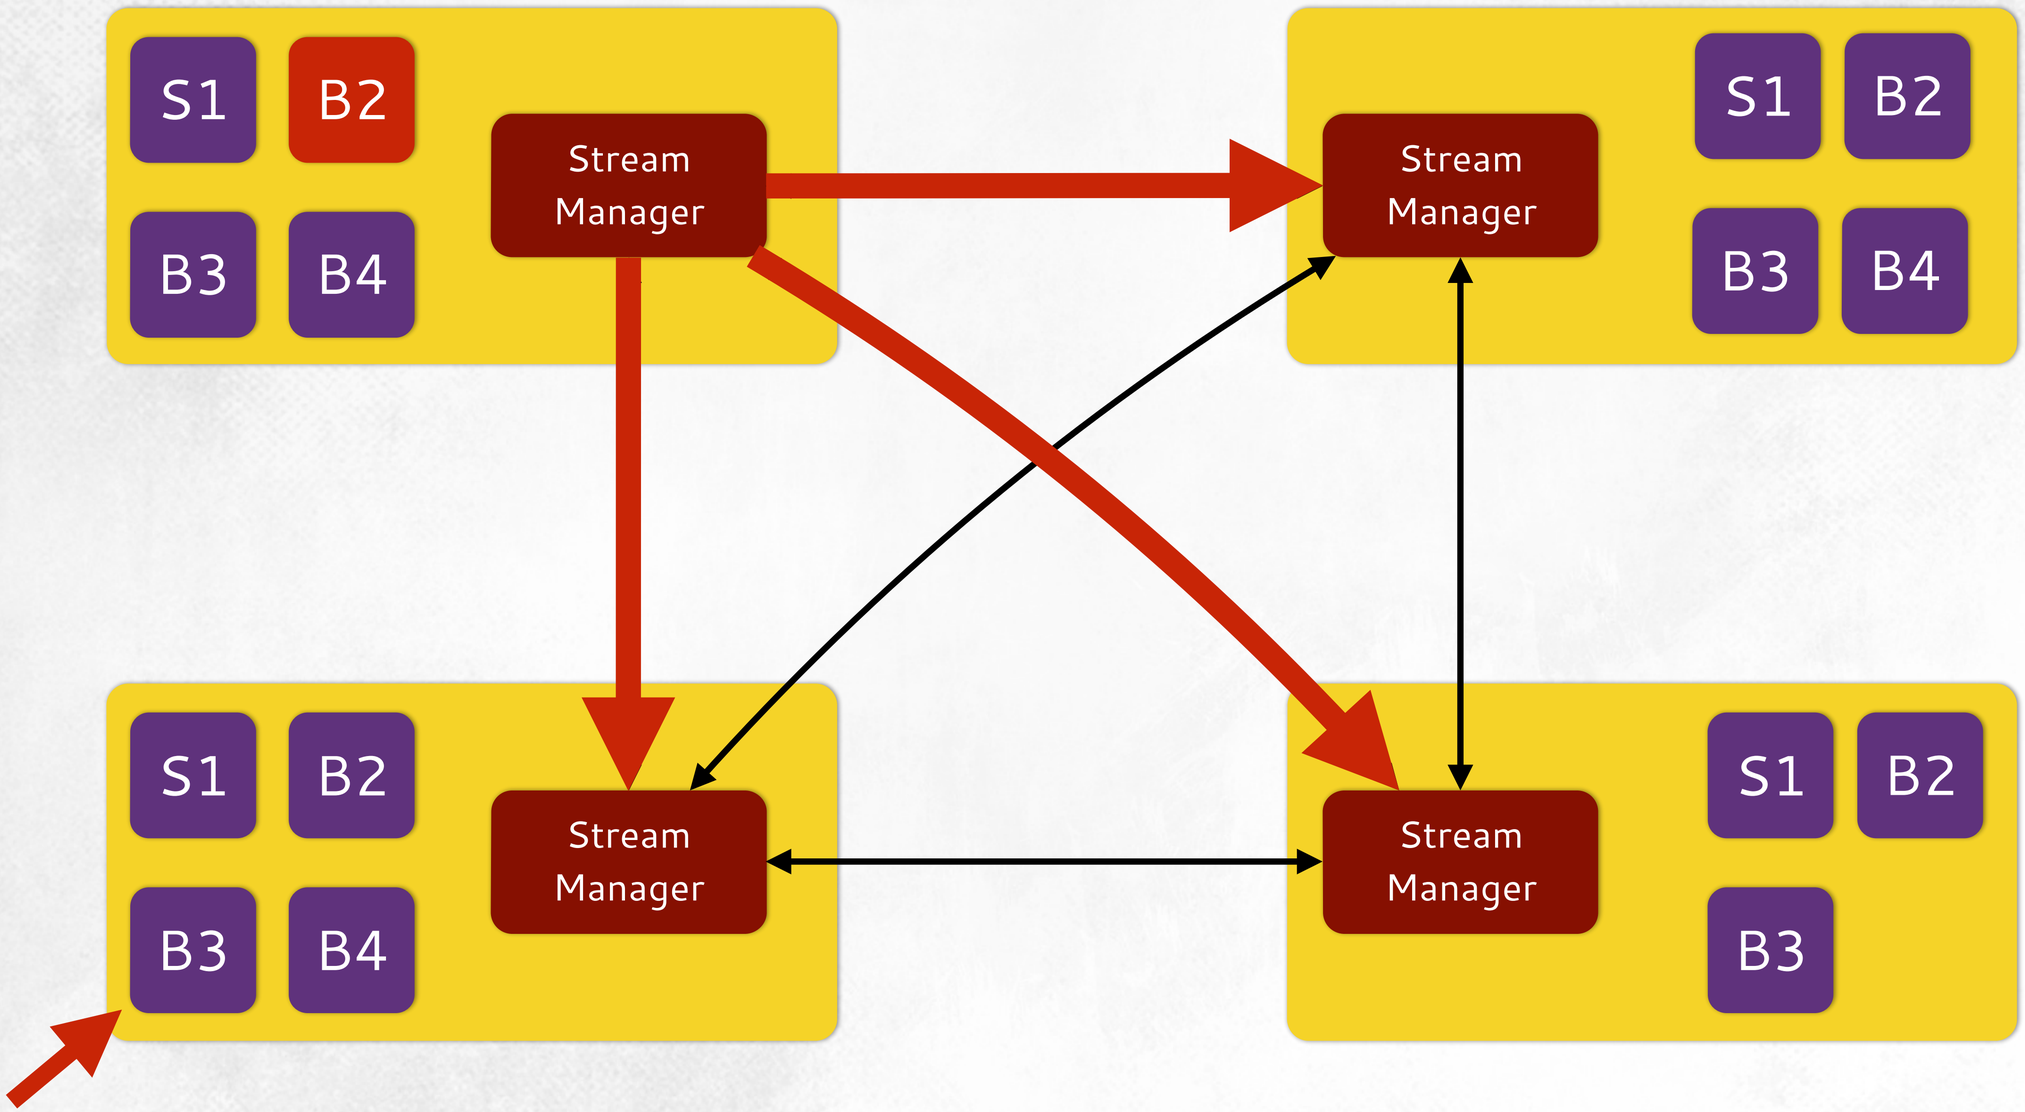

SLOWS UPSTREAM AND DOWNSTREAM INSTANCES

HERON INSTANCE

Does the real work!

RUNS ONE TASK

EXPOSES API

COLLECTS METRICS

HERON INSTANCE

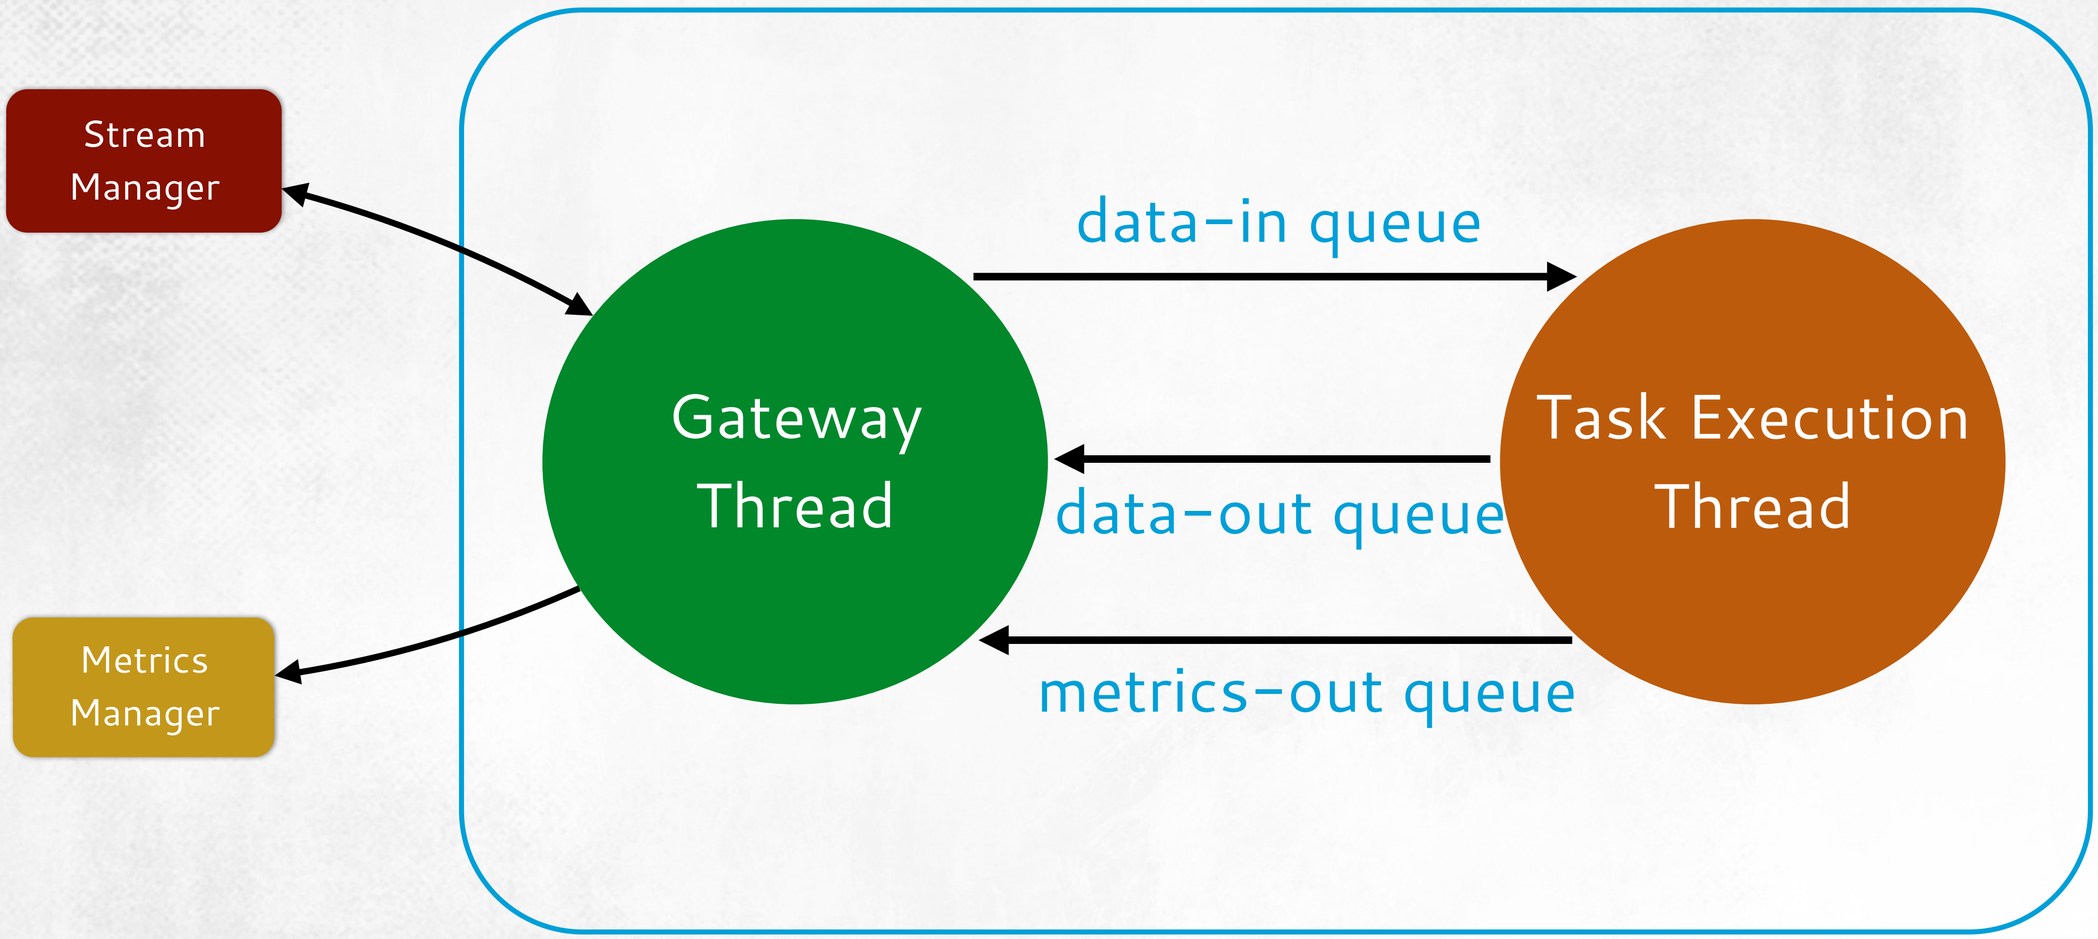

OPERATIONAL EXPERIENCES

HERON @TWITTER

STORM is decommissioned

Large amount of data produced every day

Large cluster

Several topologies deployed

Several billion messages every day

3x reduction in cores and memory

HERON PERFORMANCE

Word count topology - Acknowledgements enabled

Throughput

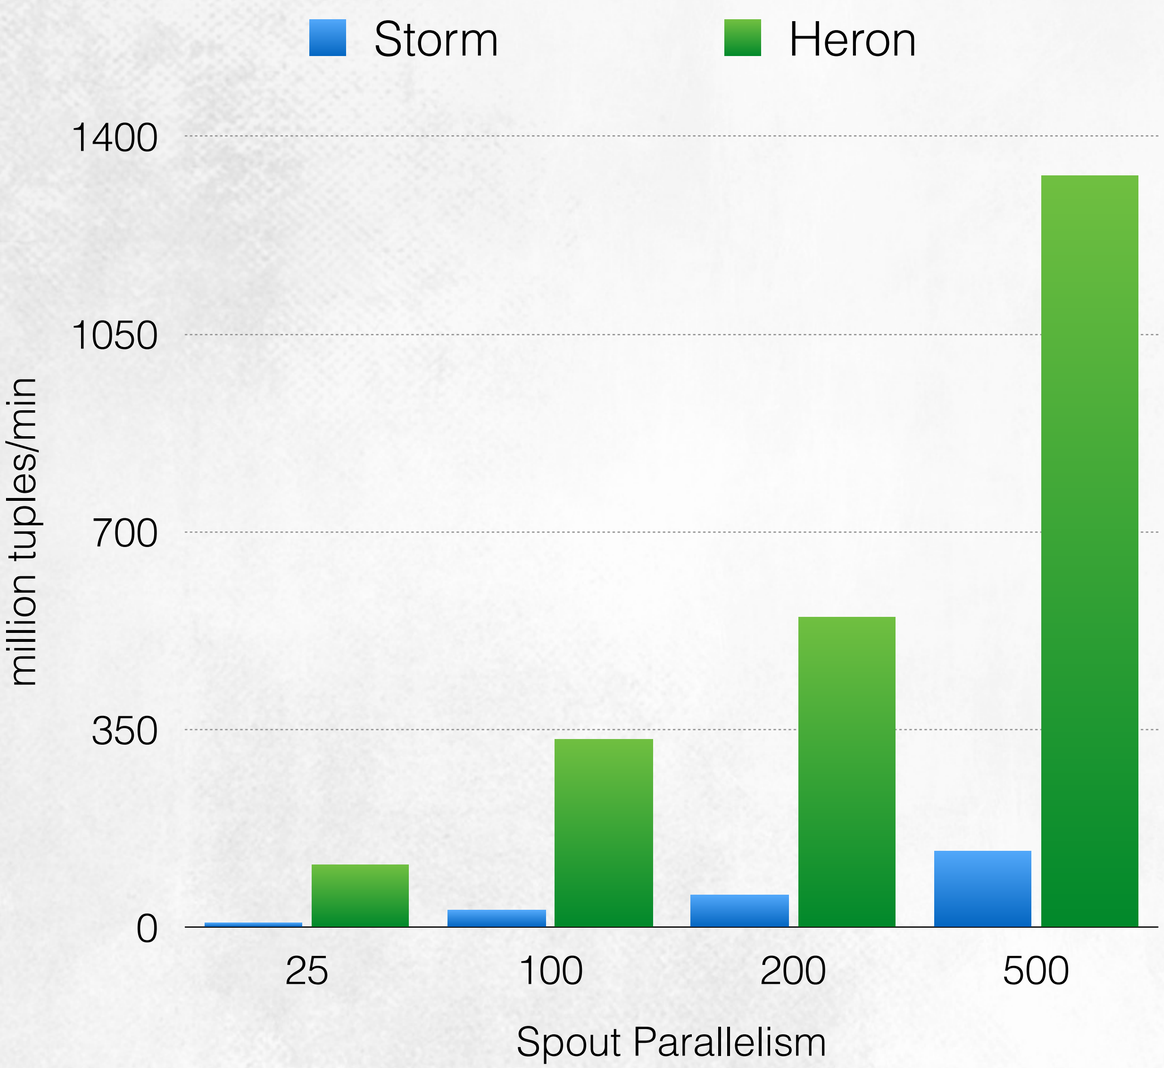

10-14x

Latency

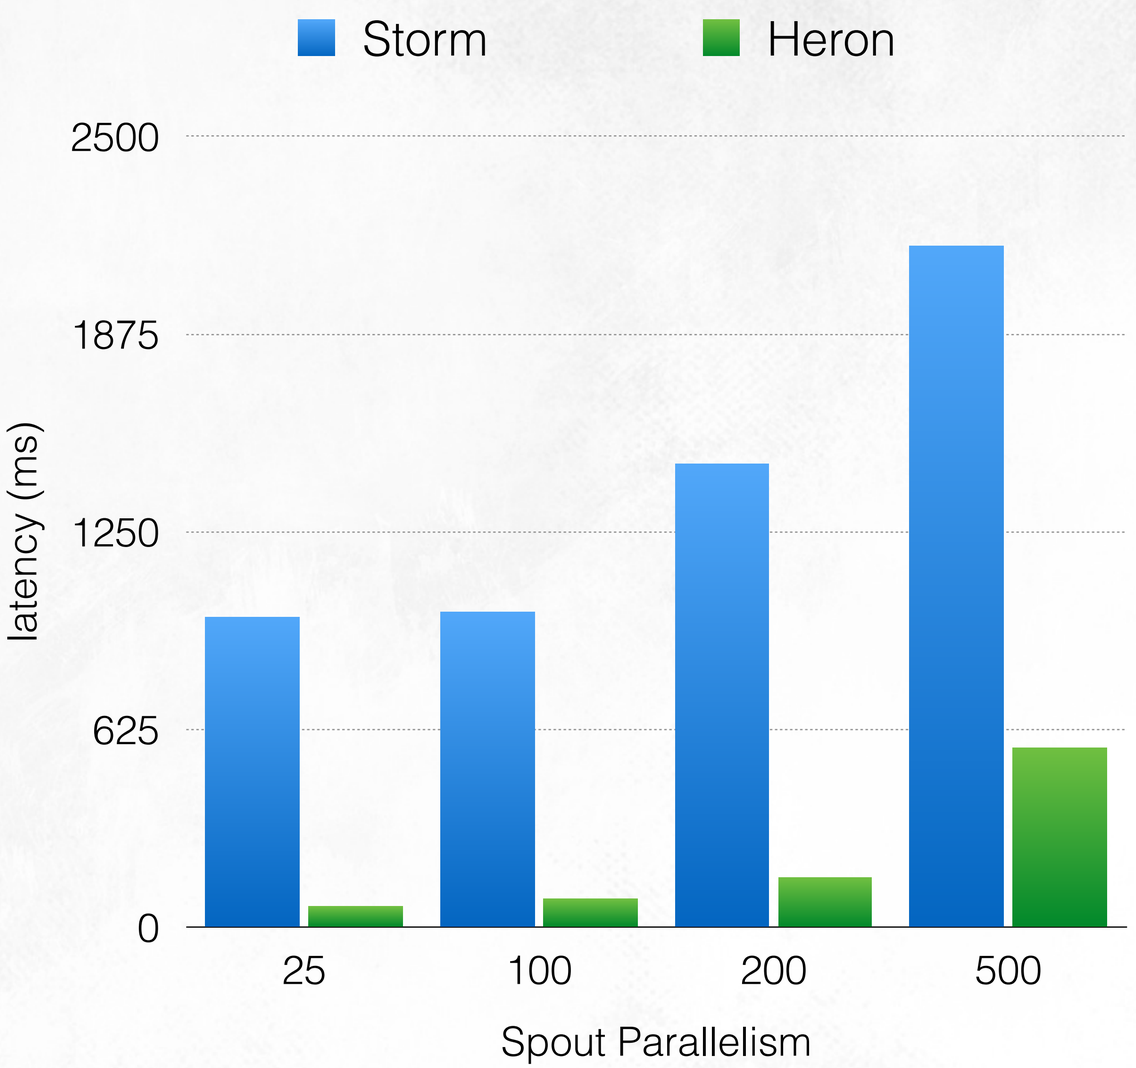

5-15x

HERON PERFORMANCE

Word count topology - CPU usage

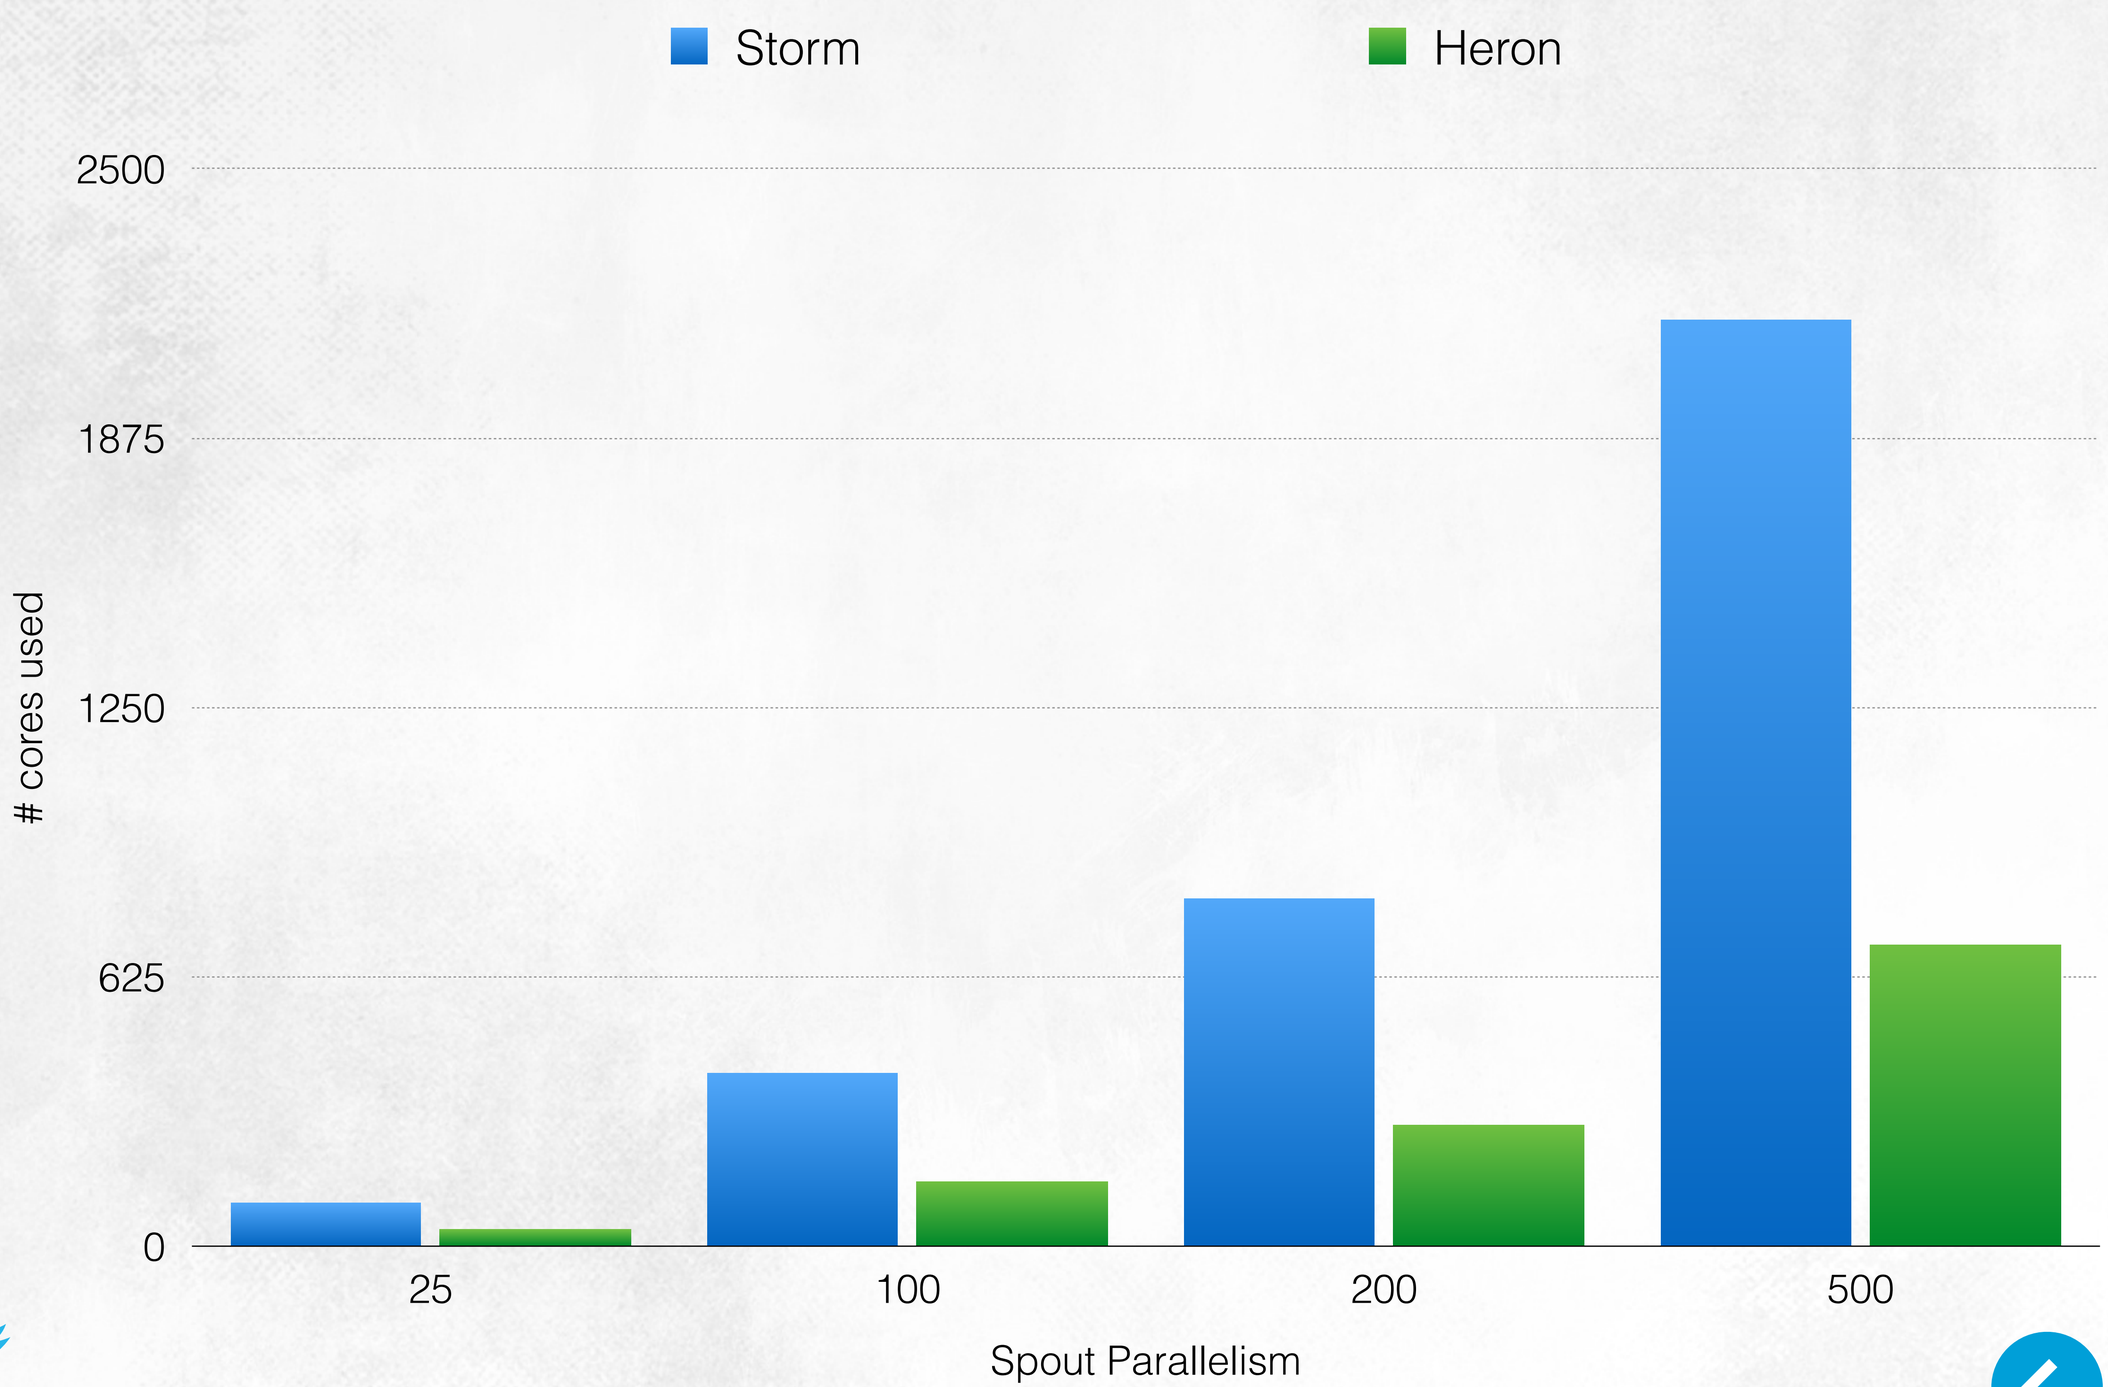

2-3x

CONCLUSION

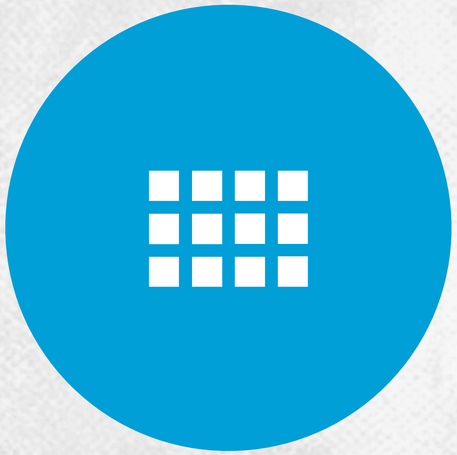

CONTAINER ARCHITECTURE

Use off the shelf schedulers

Simplified Operations

SIMPLIFIED/SEPARATE COMPONENTS

Easier to reason about behavior

Increases community collaboration

HIGH PERFORMANCE

3-5x increase in throughput

#ThankYou

FOR LISTENING