



# STREAMING IN PRACTICE

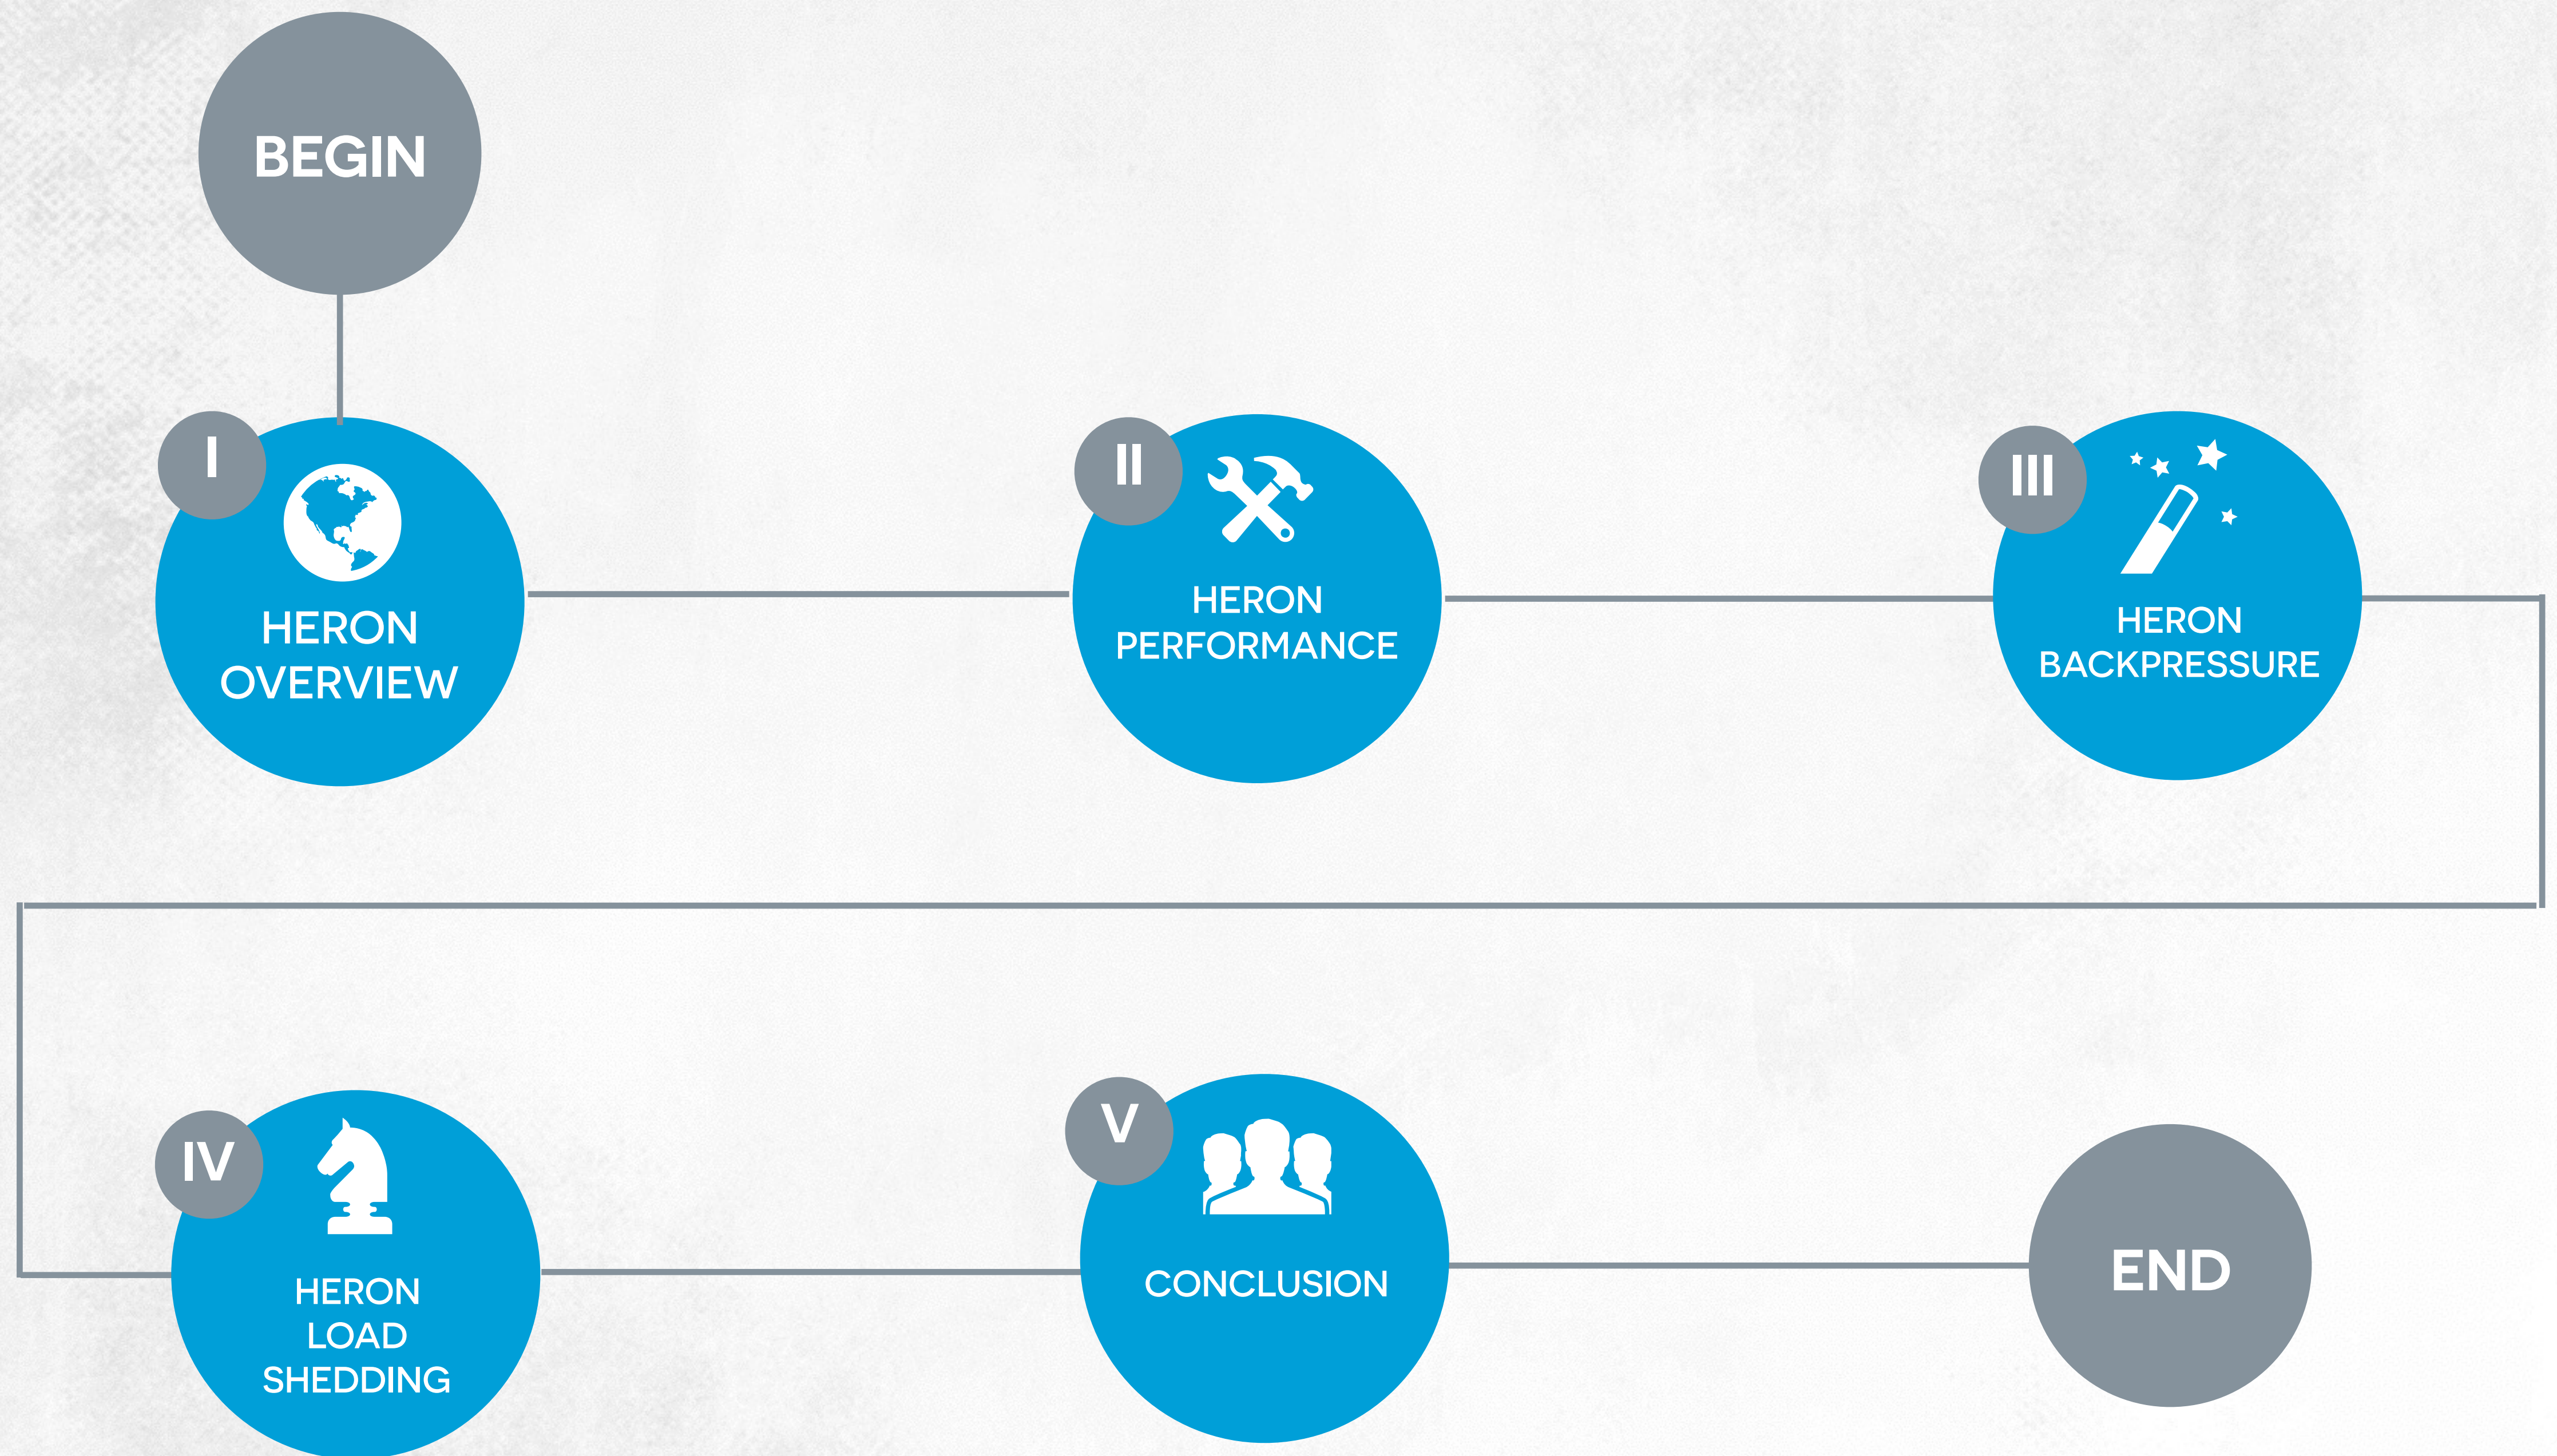
KARTHIK RAMASAMY

@KARTHIKZ

#TwitterHeron



# TALK OUTLINE







# HERON OVERVIEW



# STORM/HERON **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

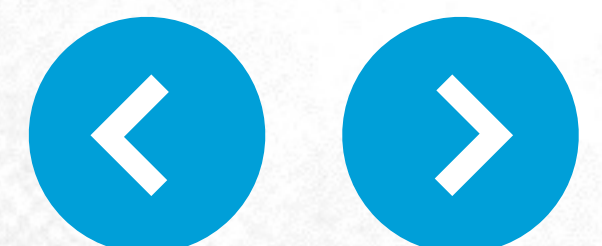
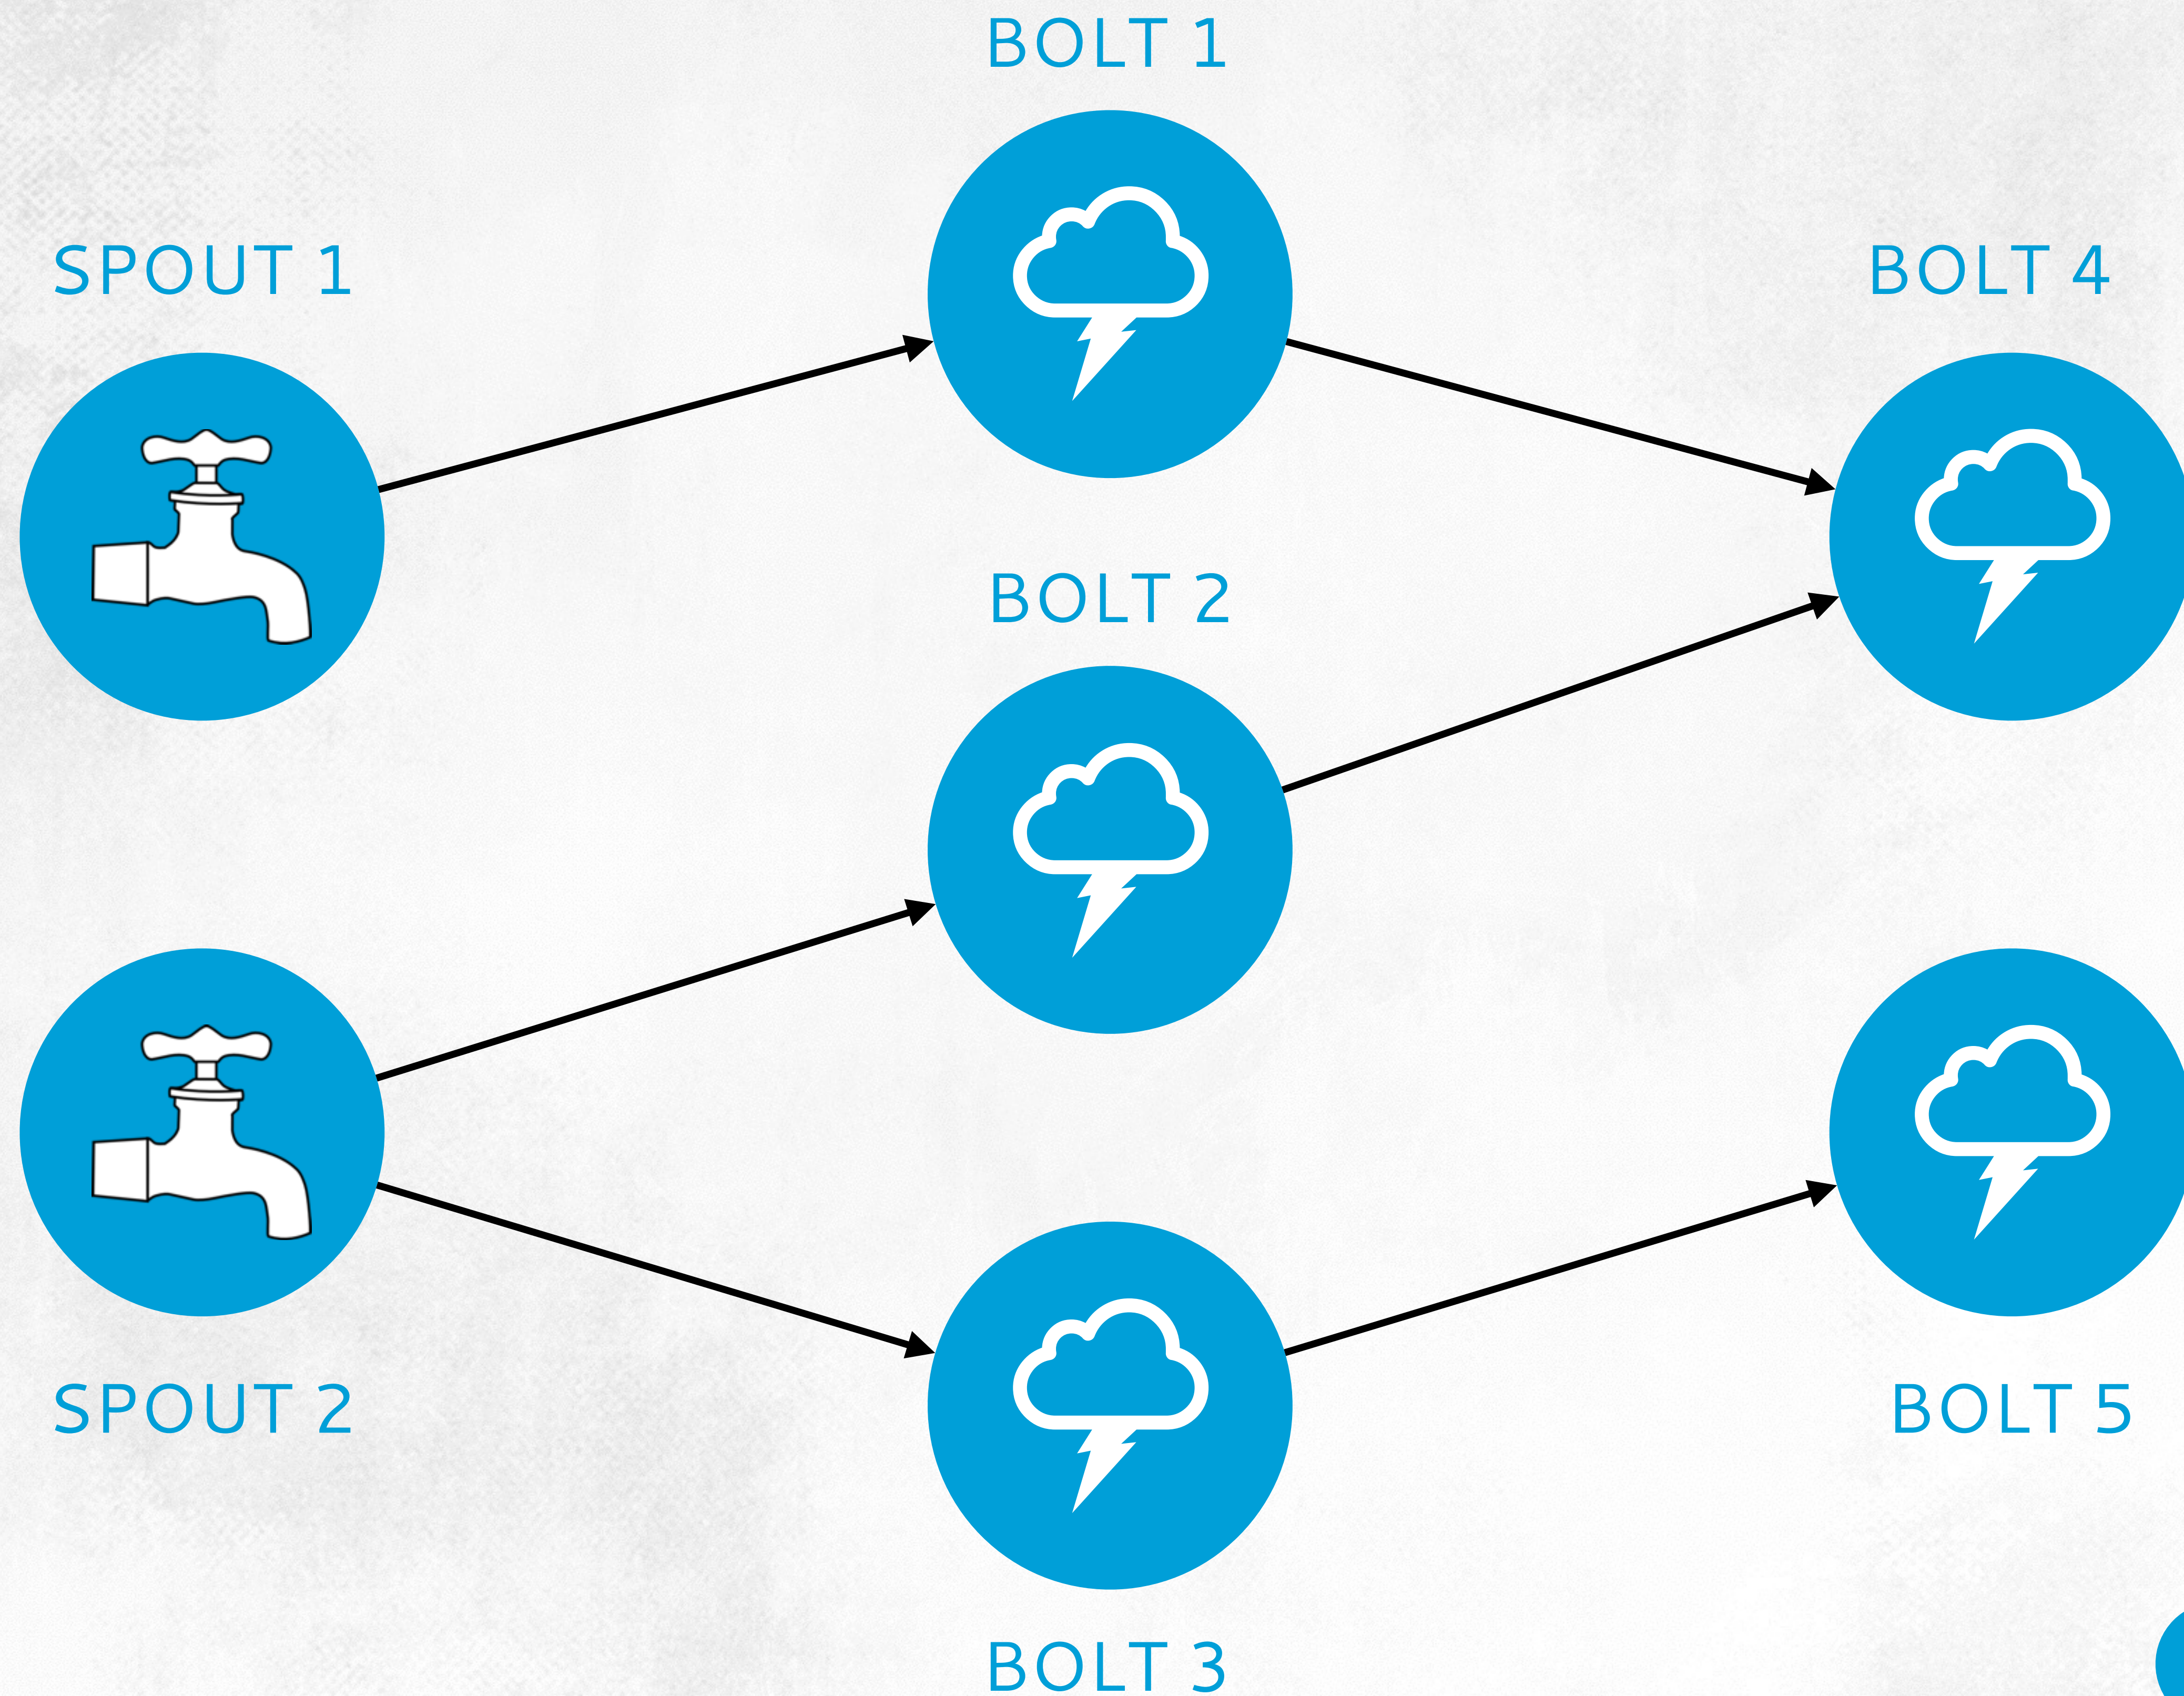
Process incoming tuples and emit outgoing tuples

Examples – filtering/aggregation/join/arbitrary function





# STORM/HERON TOPOLOGY

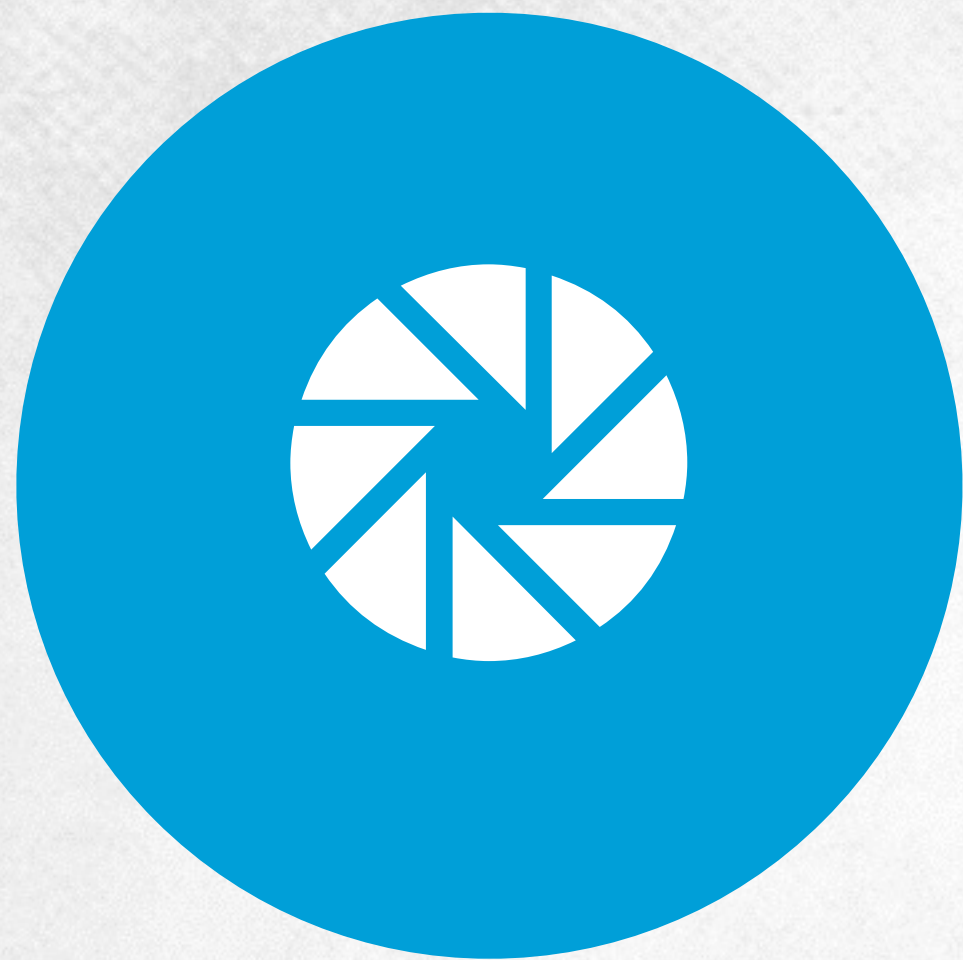




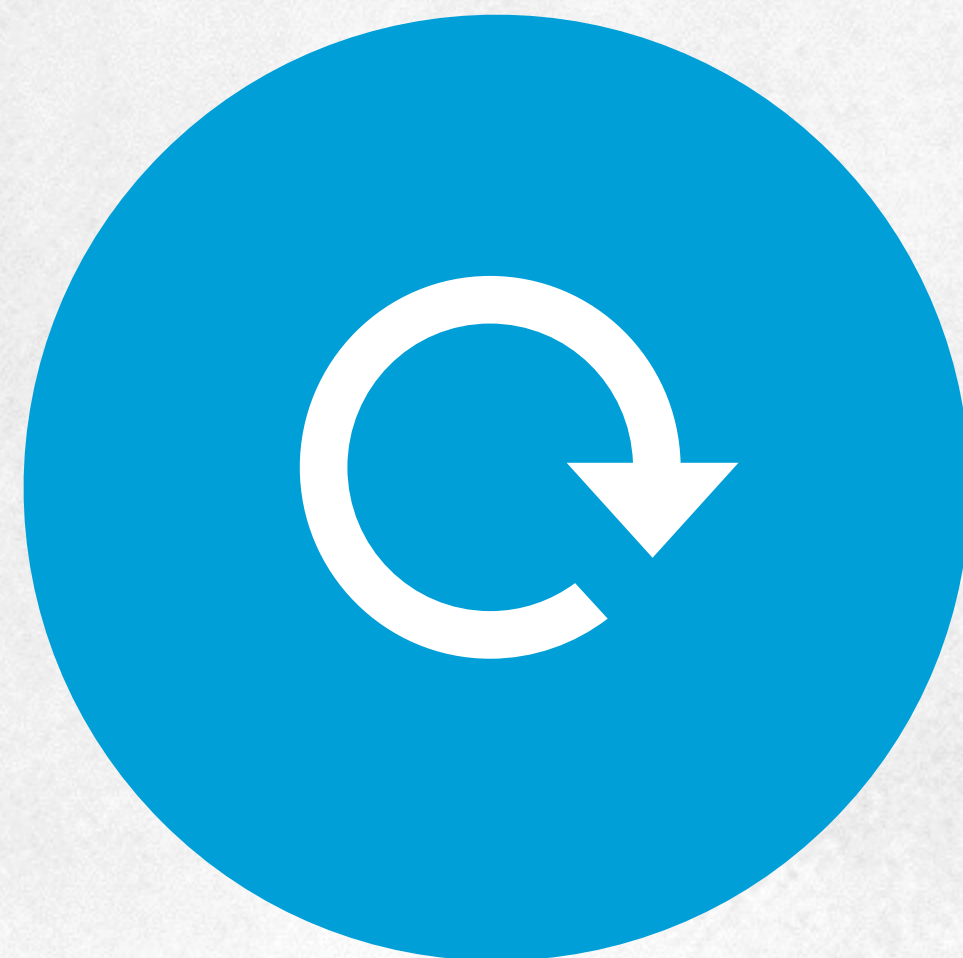
# WHY HERON?



PERFORMANCE PREDICTABILITY



IMPROVE DEVELOPER PRODUCTIVITY



EASE OF MANAGEABILITY





# HERON DESIGN DECISIONS



## FULLY API COMPATIBLE WITH STORM

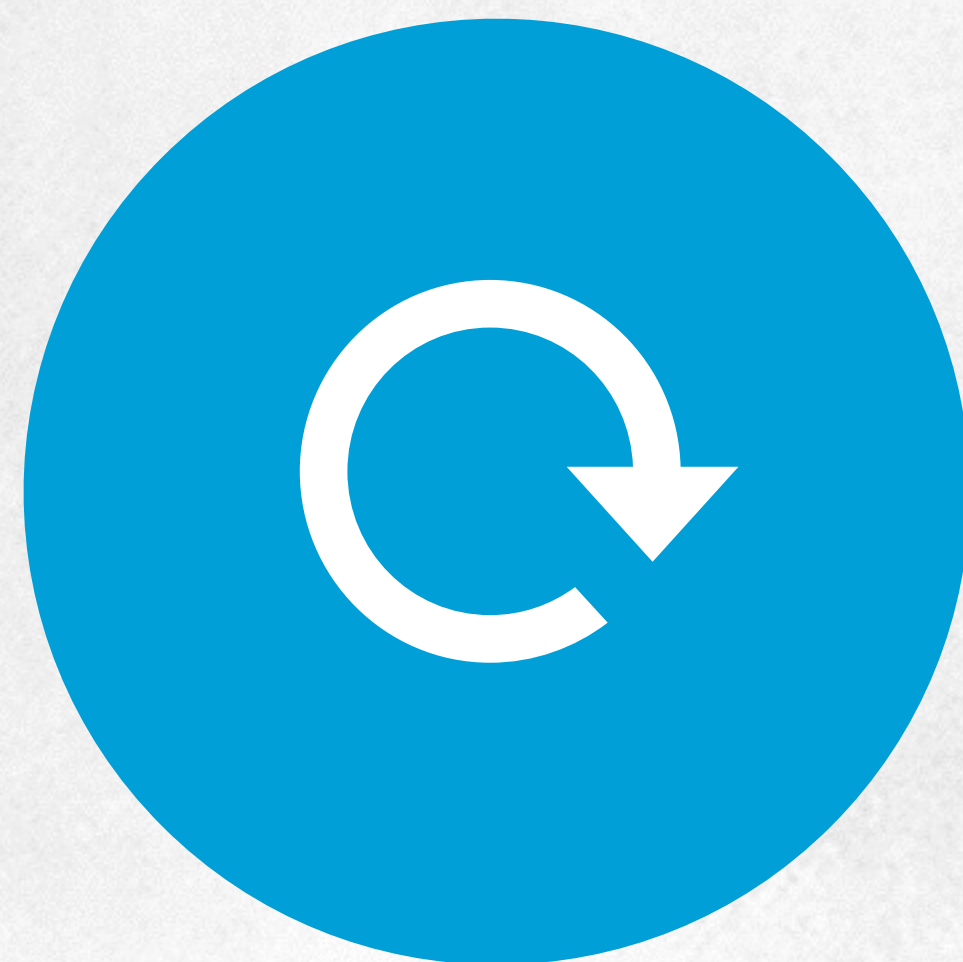
Directed acyclic graph

Topologies, spouts and bolts



## TASK ISOLATION

Ease of debug ability/resource isolation/profiling



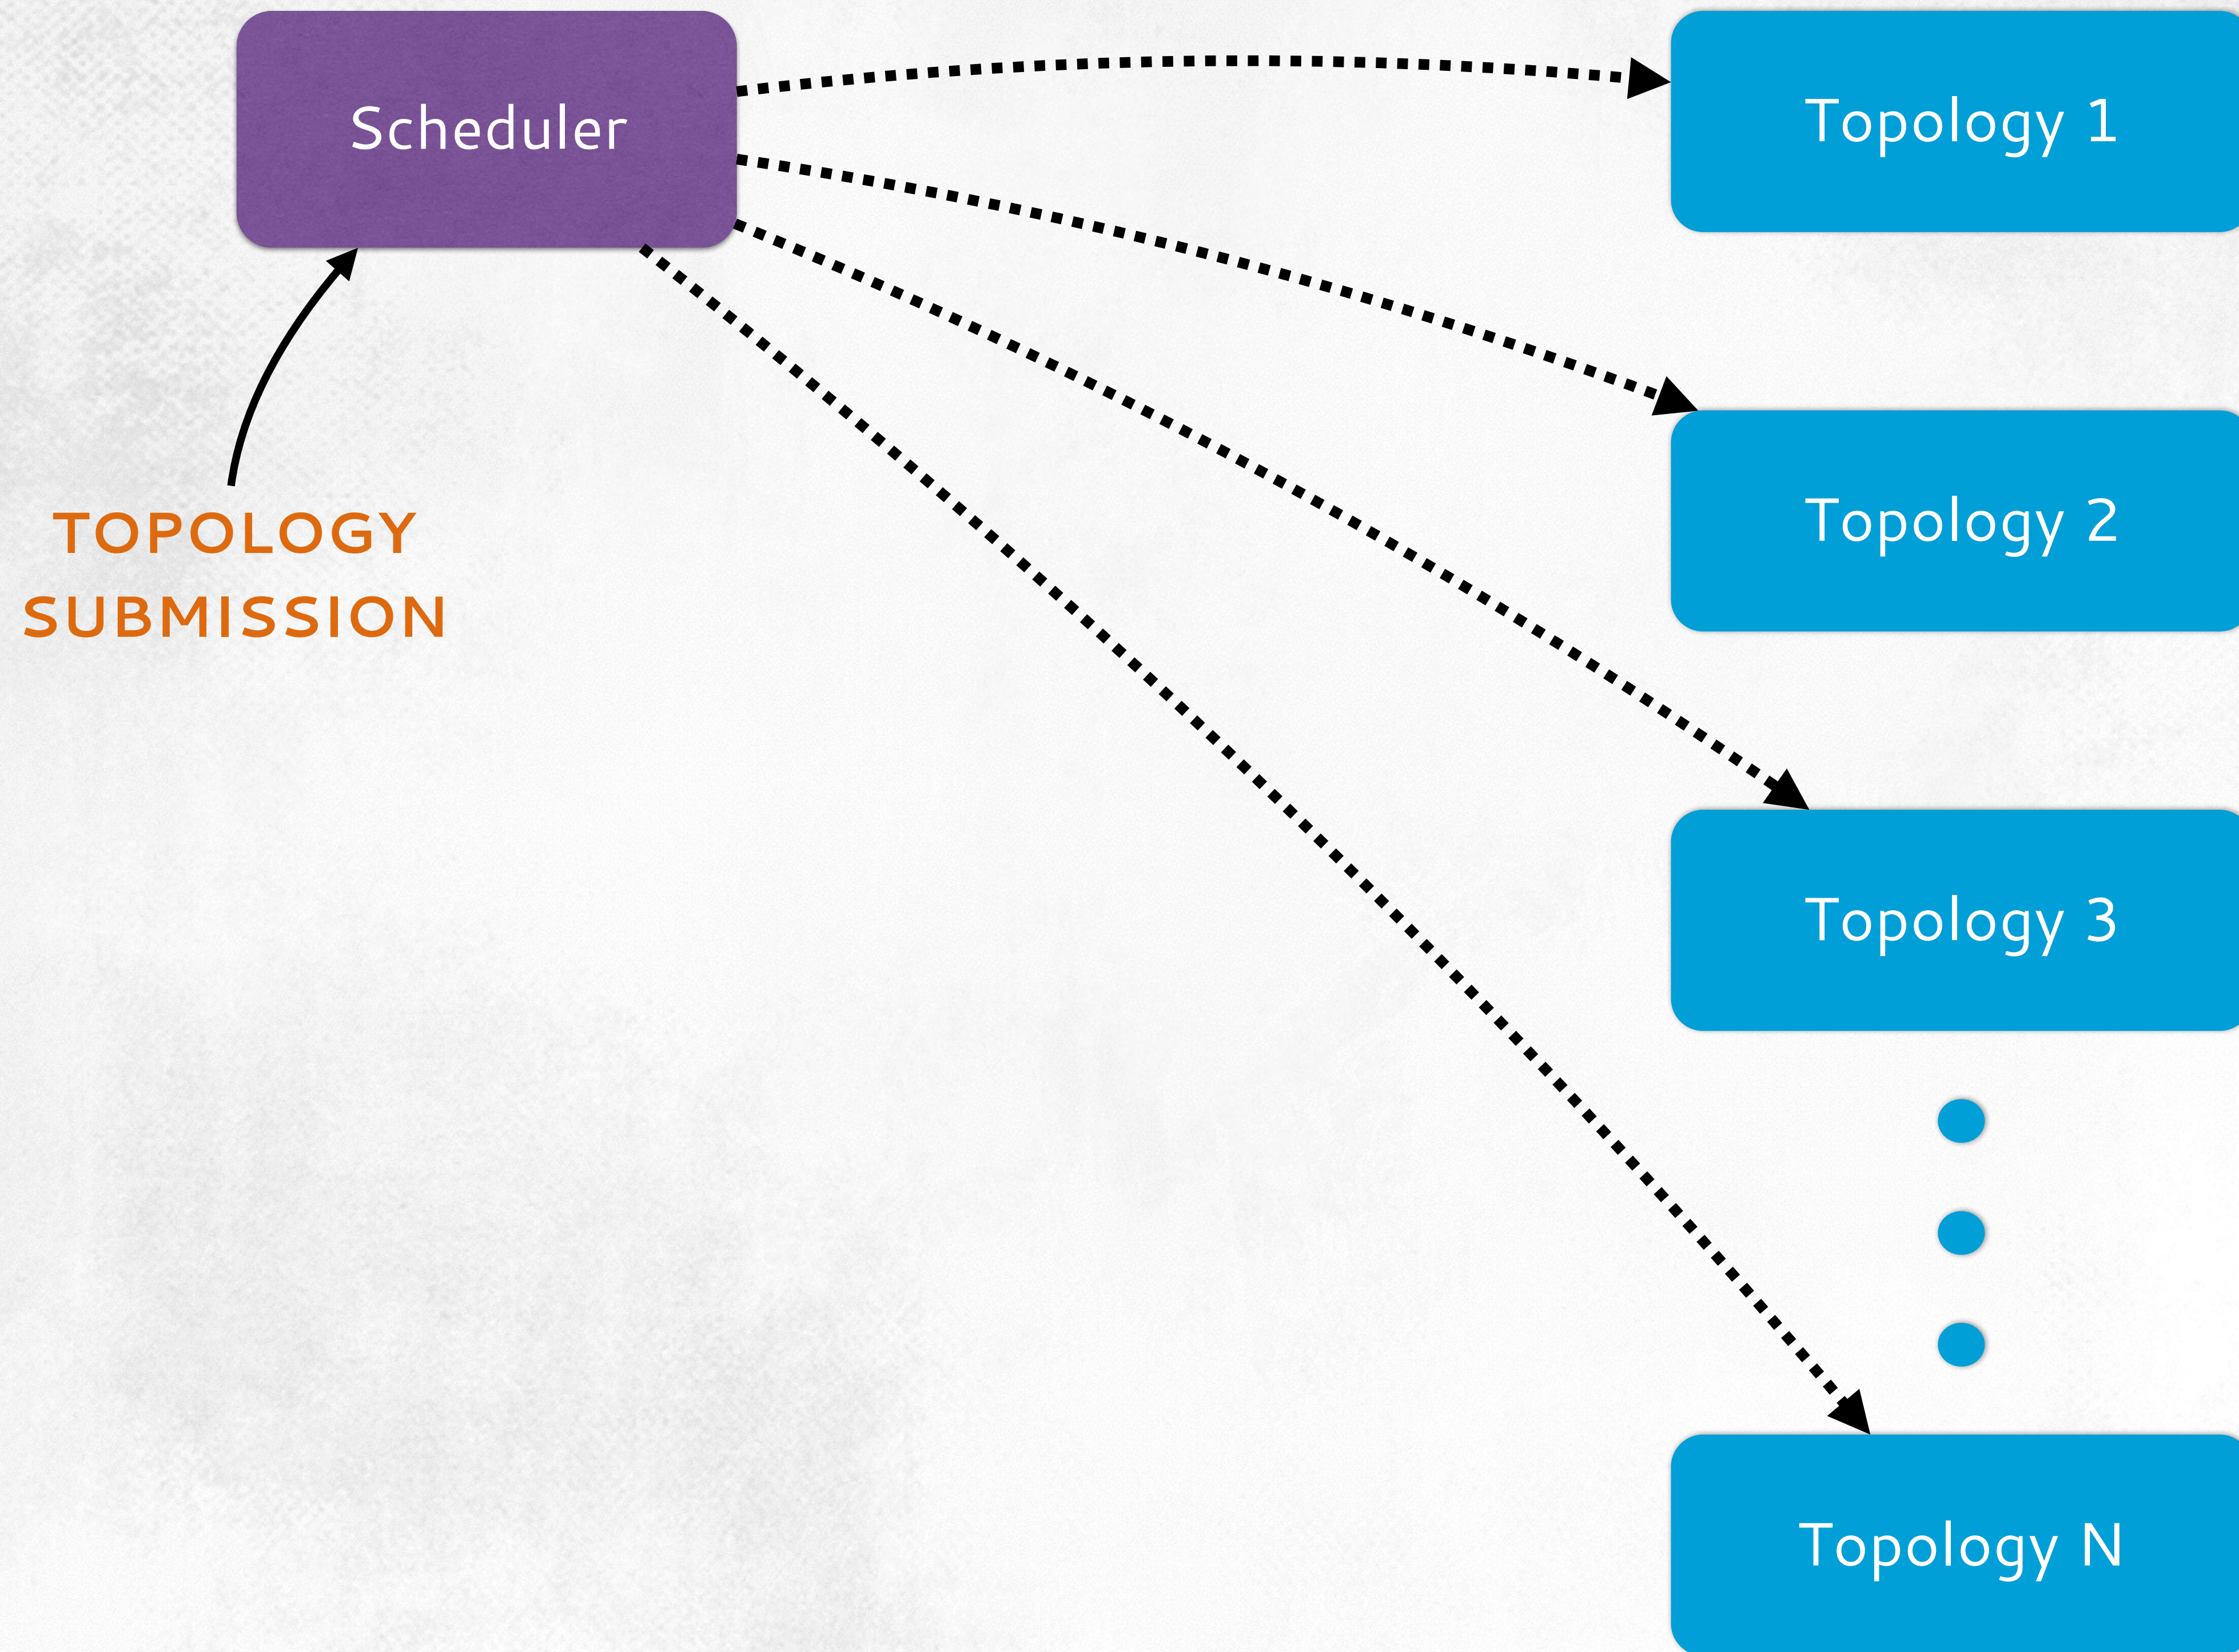
## USE OF MAIN STREAM LANGUAGES

C++/JAVA/Python



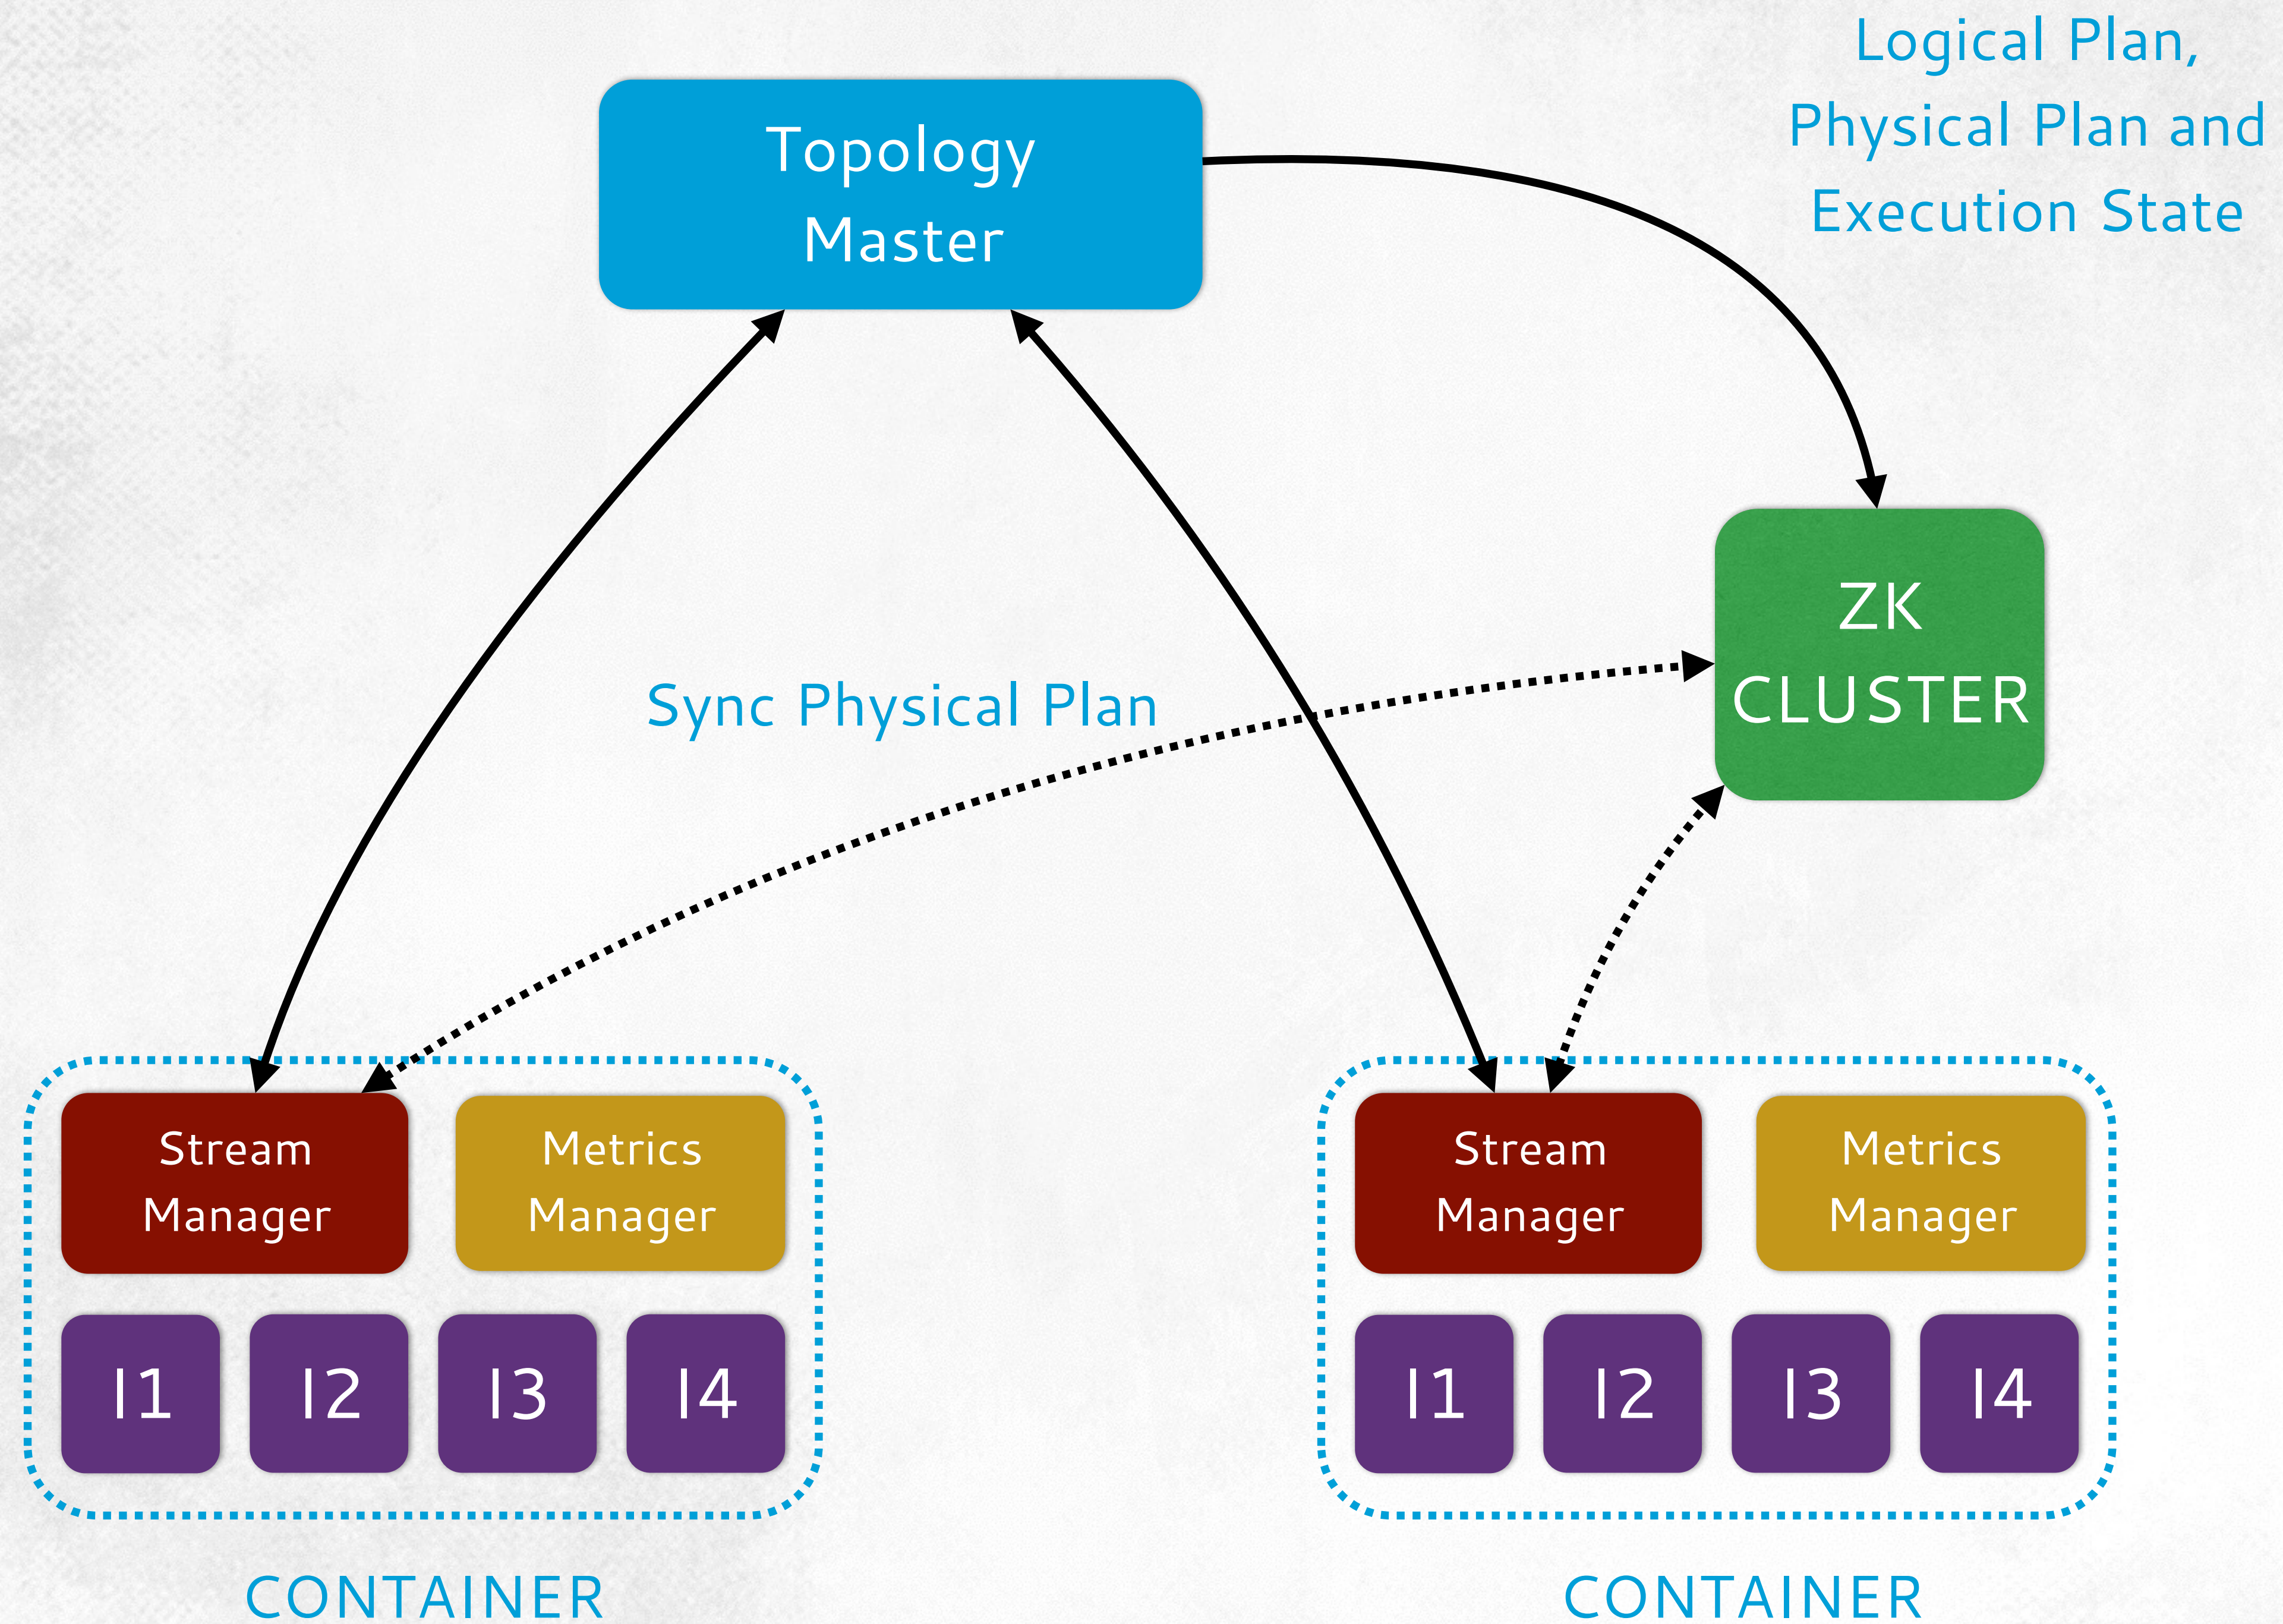


# HERON ARCHITECTURE



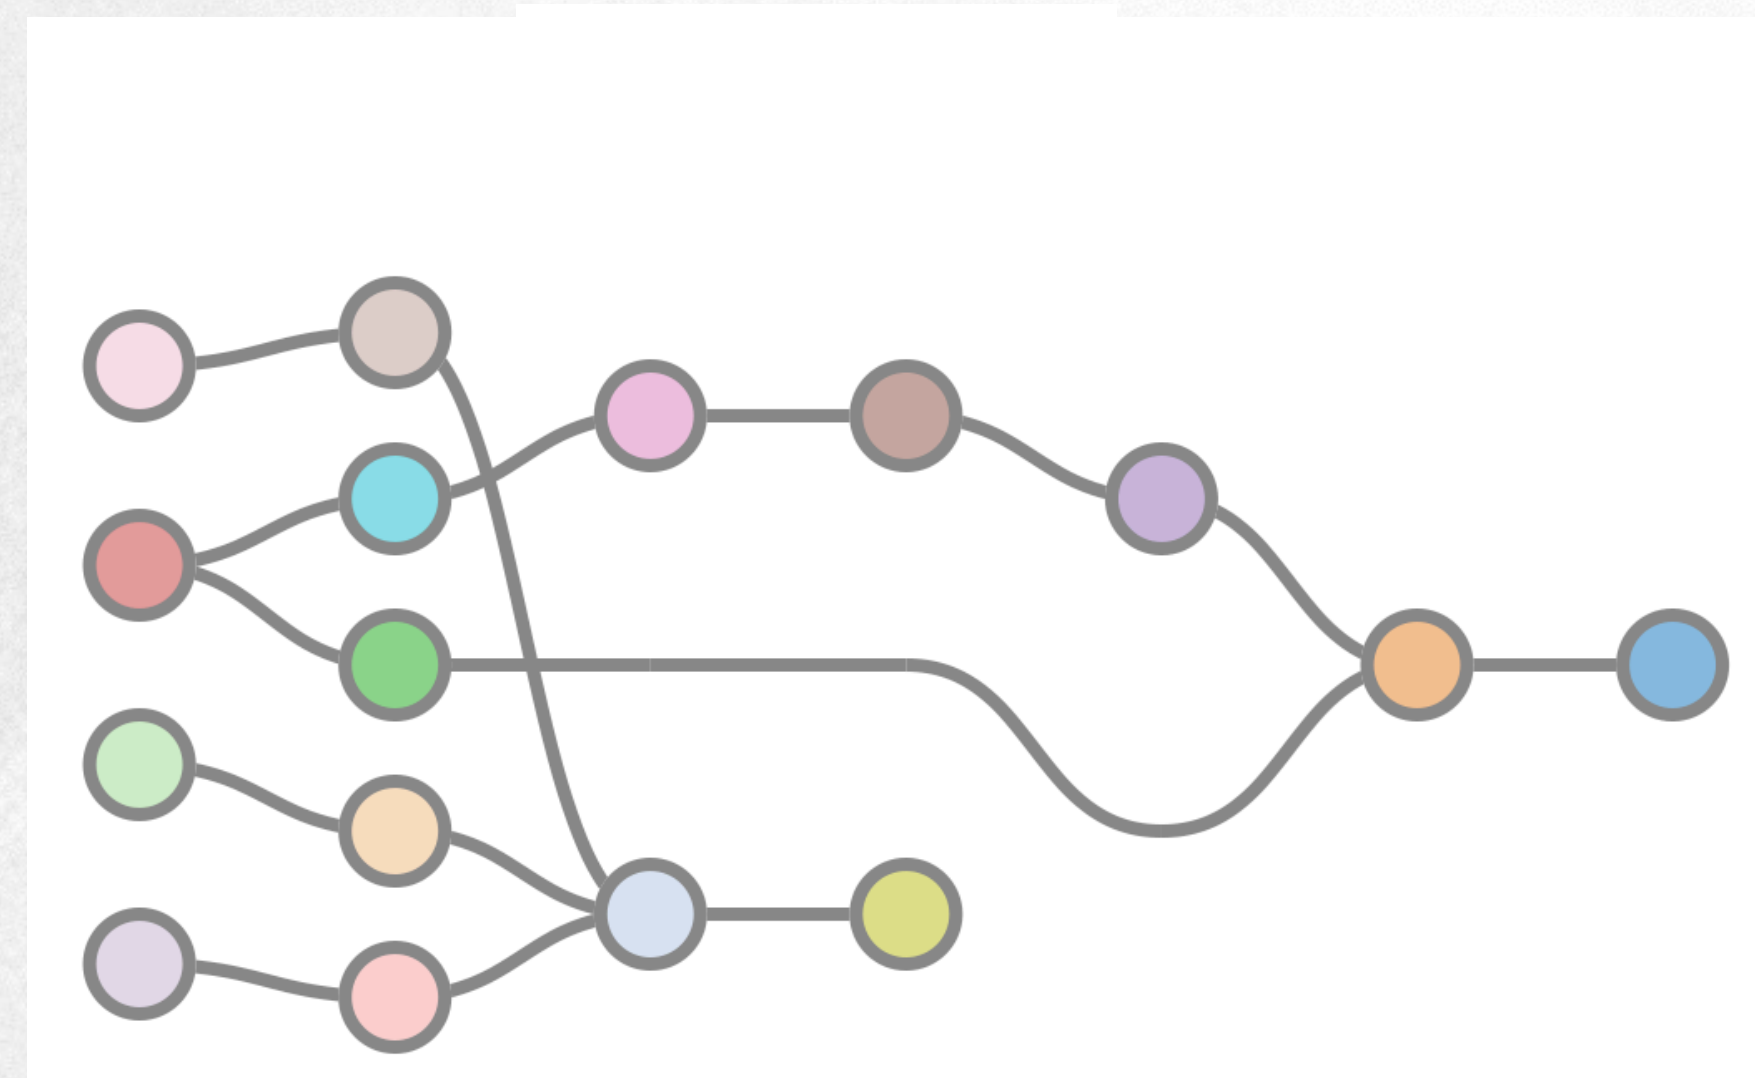
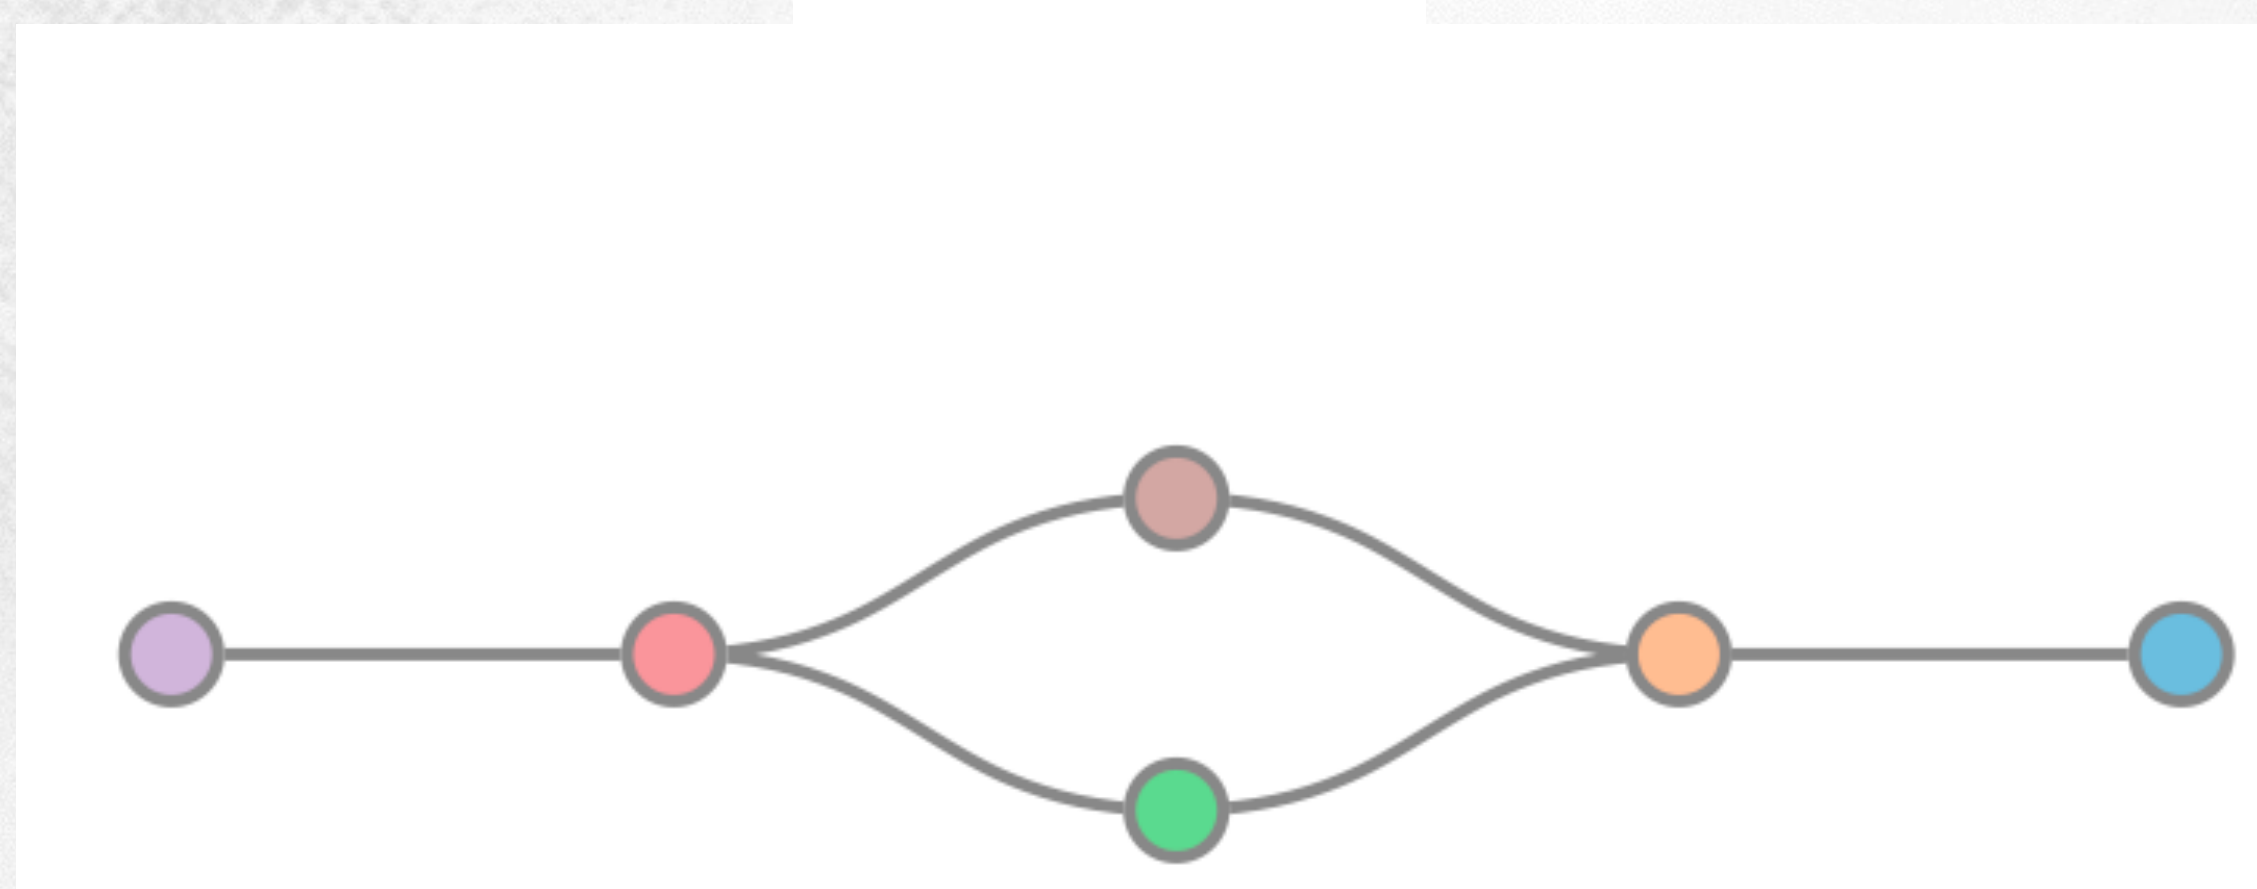
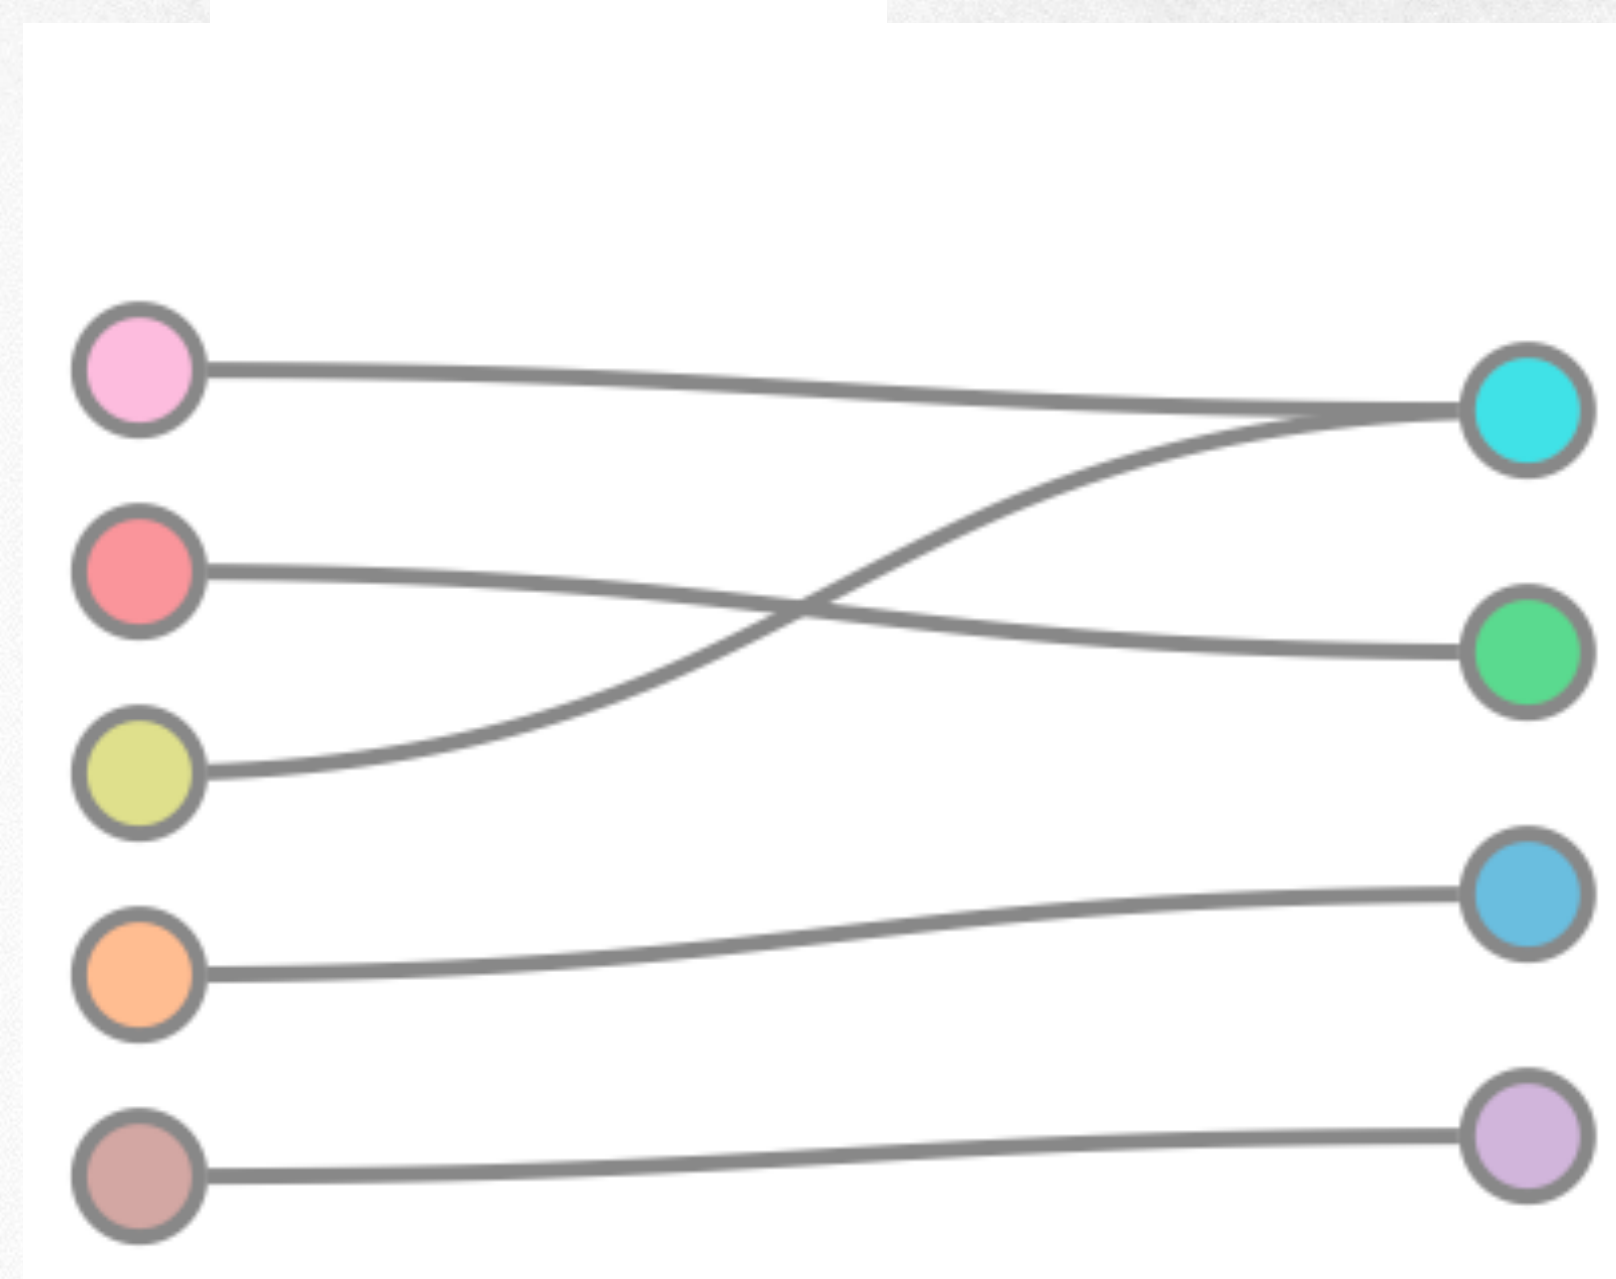
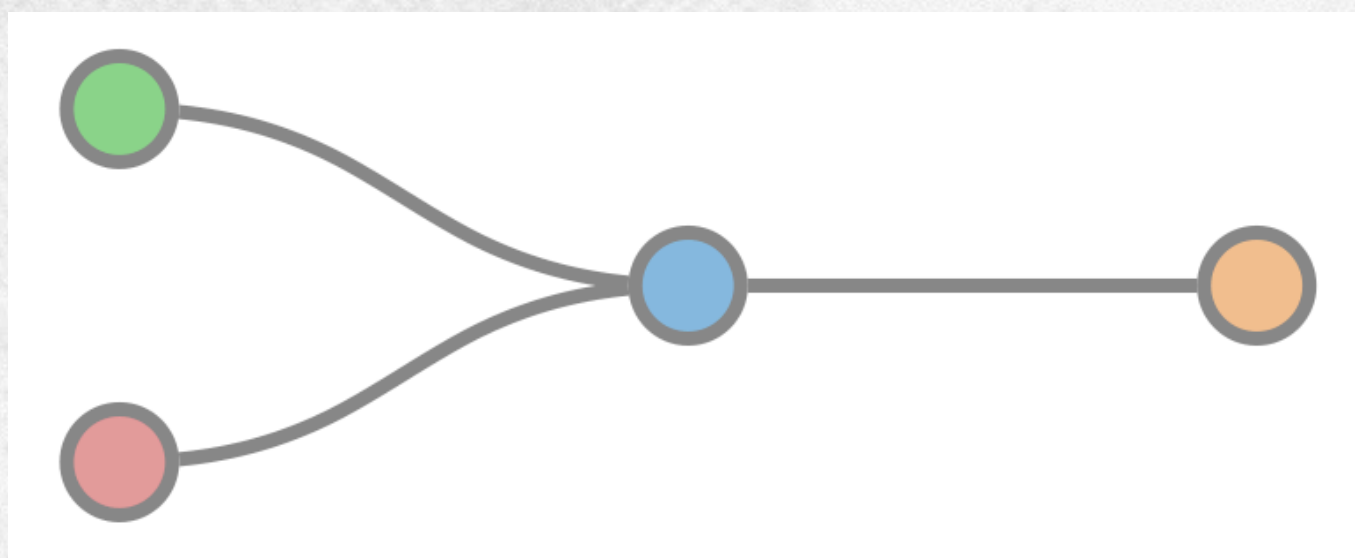


# TOPOLOGY ARCHITECTURE





# HERON SAMPLE TOPOLOGIES





# HERON @TWITTER

Heron has been in production for 2 years

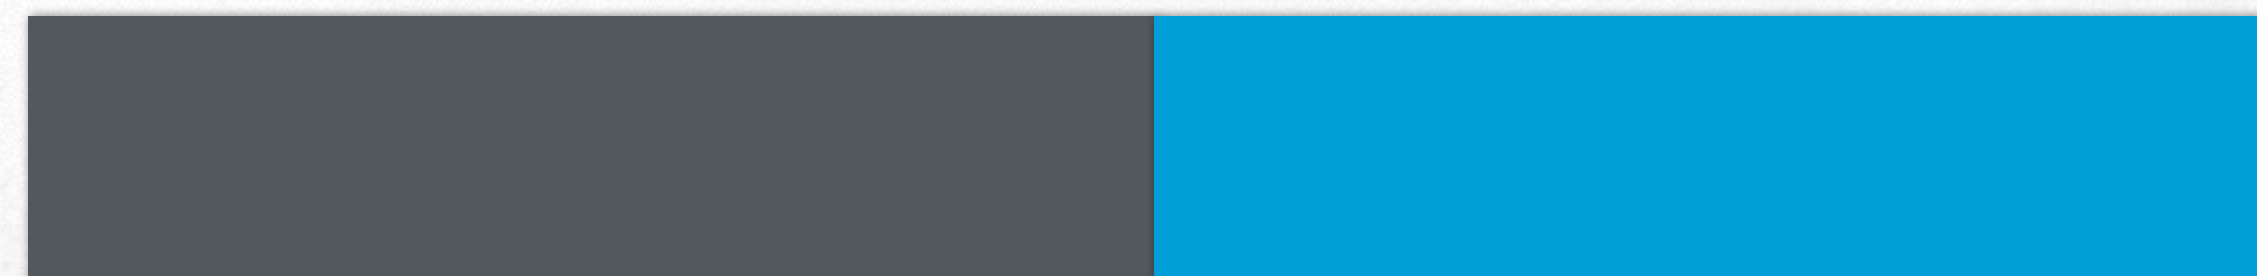
Large amount of data  
produced every day

Large cluster

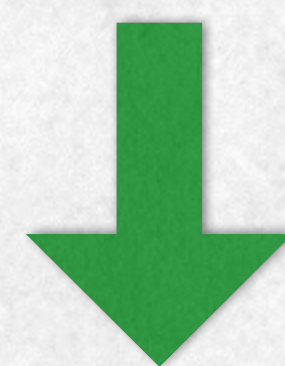
Several hundred  
topologies deployed

Several billion  
messages every day

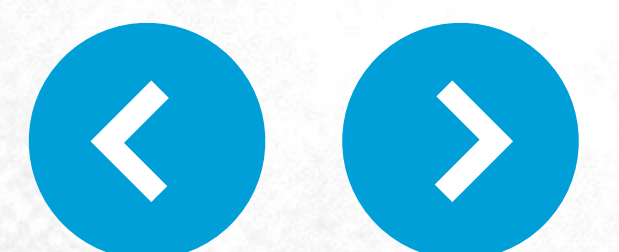
1 stage



10 stages



3x reduction in cores and memory





# HERON USE CASES

REALTIME  
ETL

REAL TIME  
BI

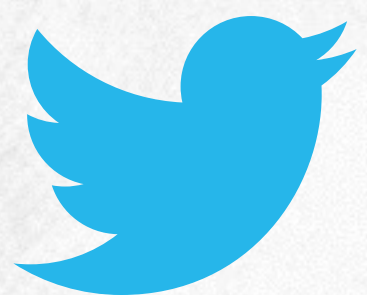
SPAM  
DETECTION

REAL TIME  
TRENDS

REALTIME  
ML

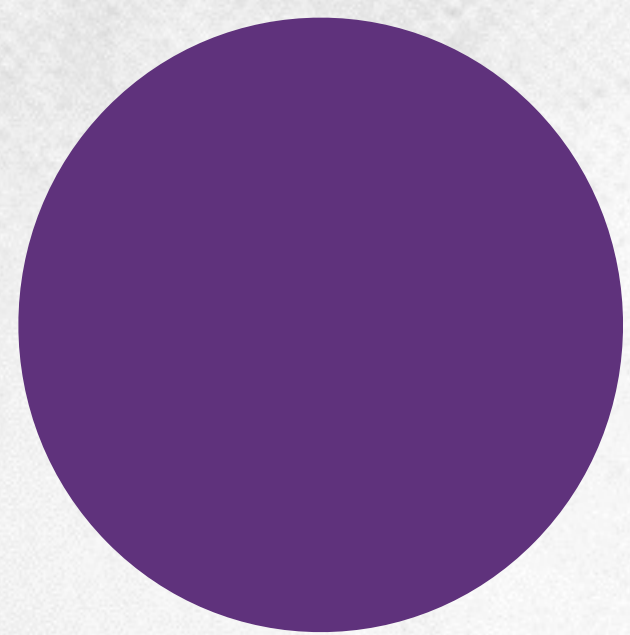
REAL TIME  
MEDIA

REAL TIME  
OPS

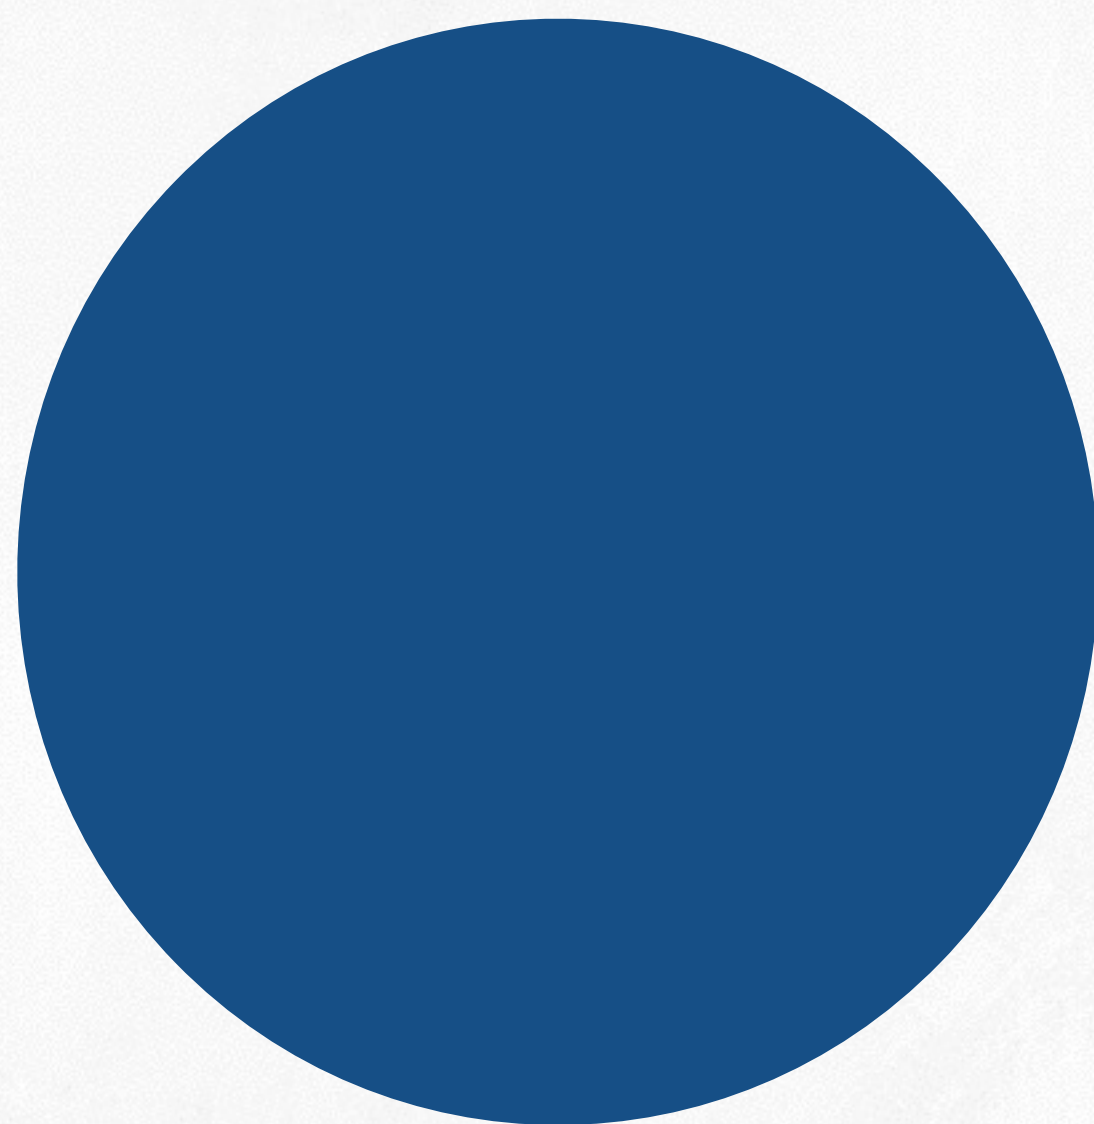




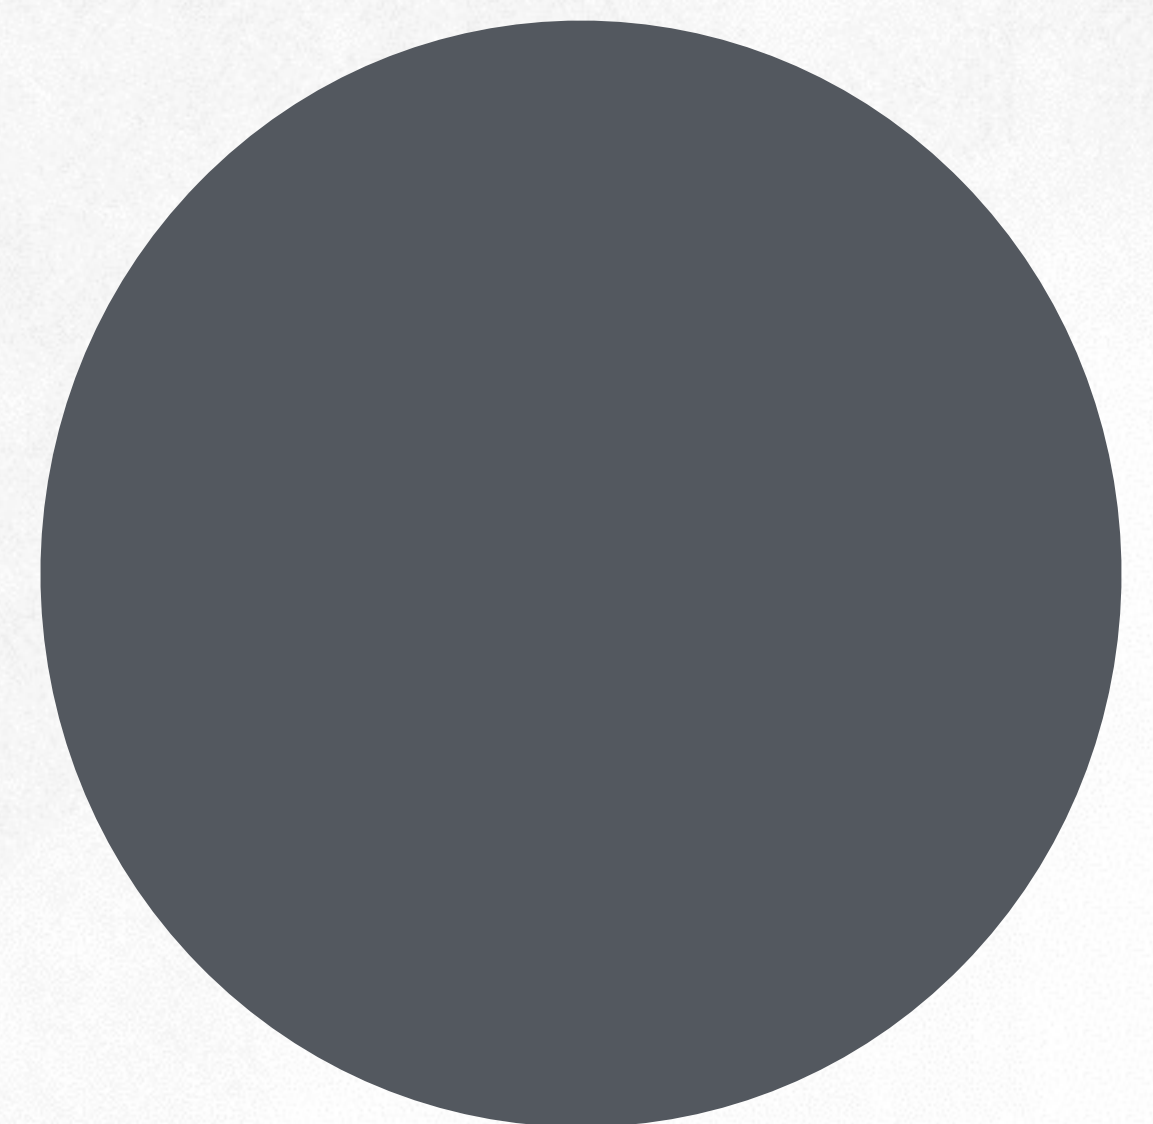
# HERON ENVIRONMENT



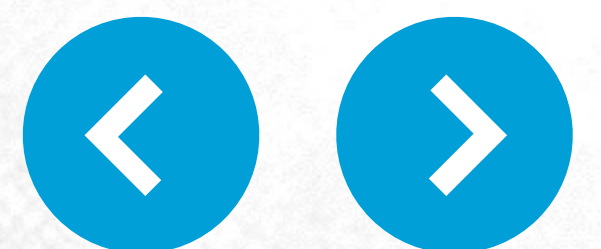
Laptop/Server



Cluster/Aurora



Cluster/Mesos







# HERON RESOURCE USAGE





# HERON PERFORMANCE

## Settings

COMPONENTS	EXPT #1	EXPT #2	EXPT #3	EXPT #4
Spout	25	100	200	300
Bolt	25	100	200	300
# Heron containers	25	100	200	300
# Storm workers	25	100	200	300

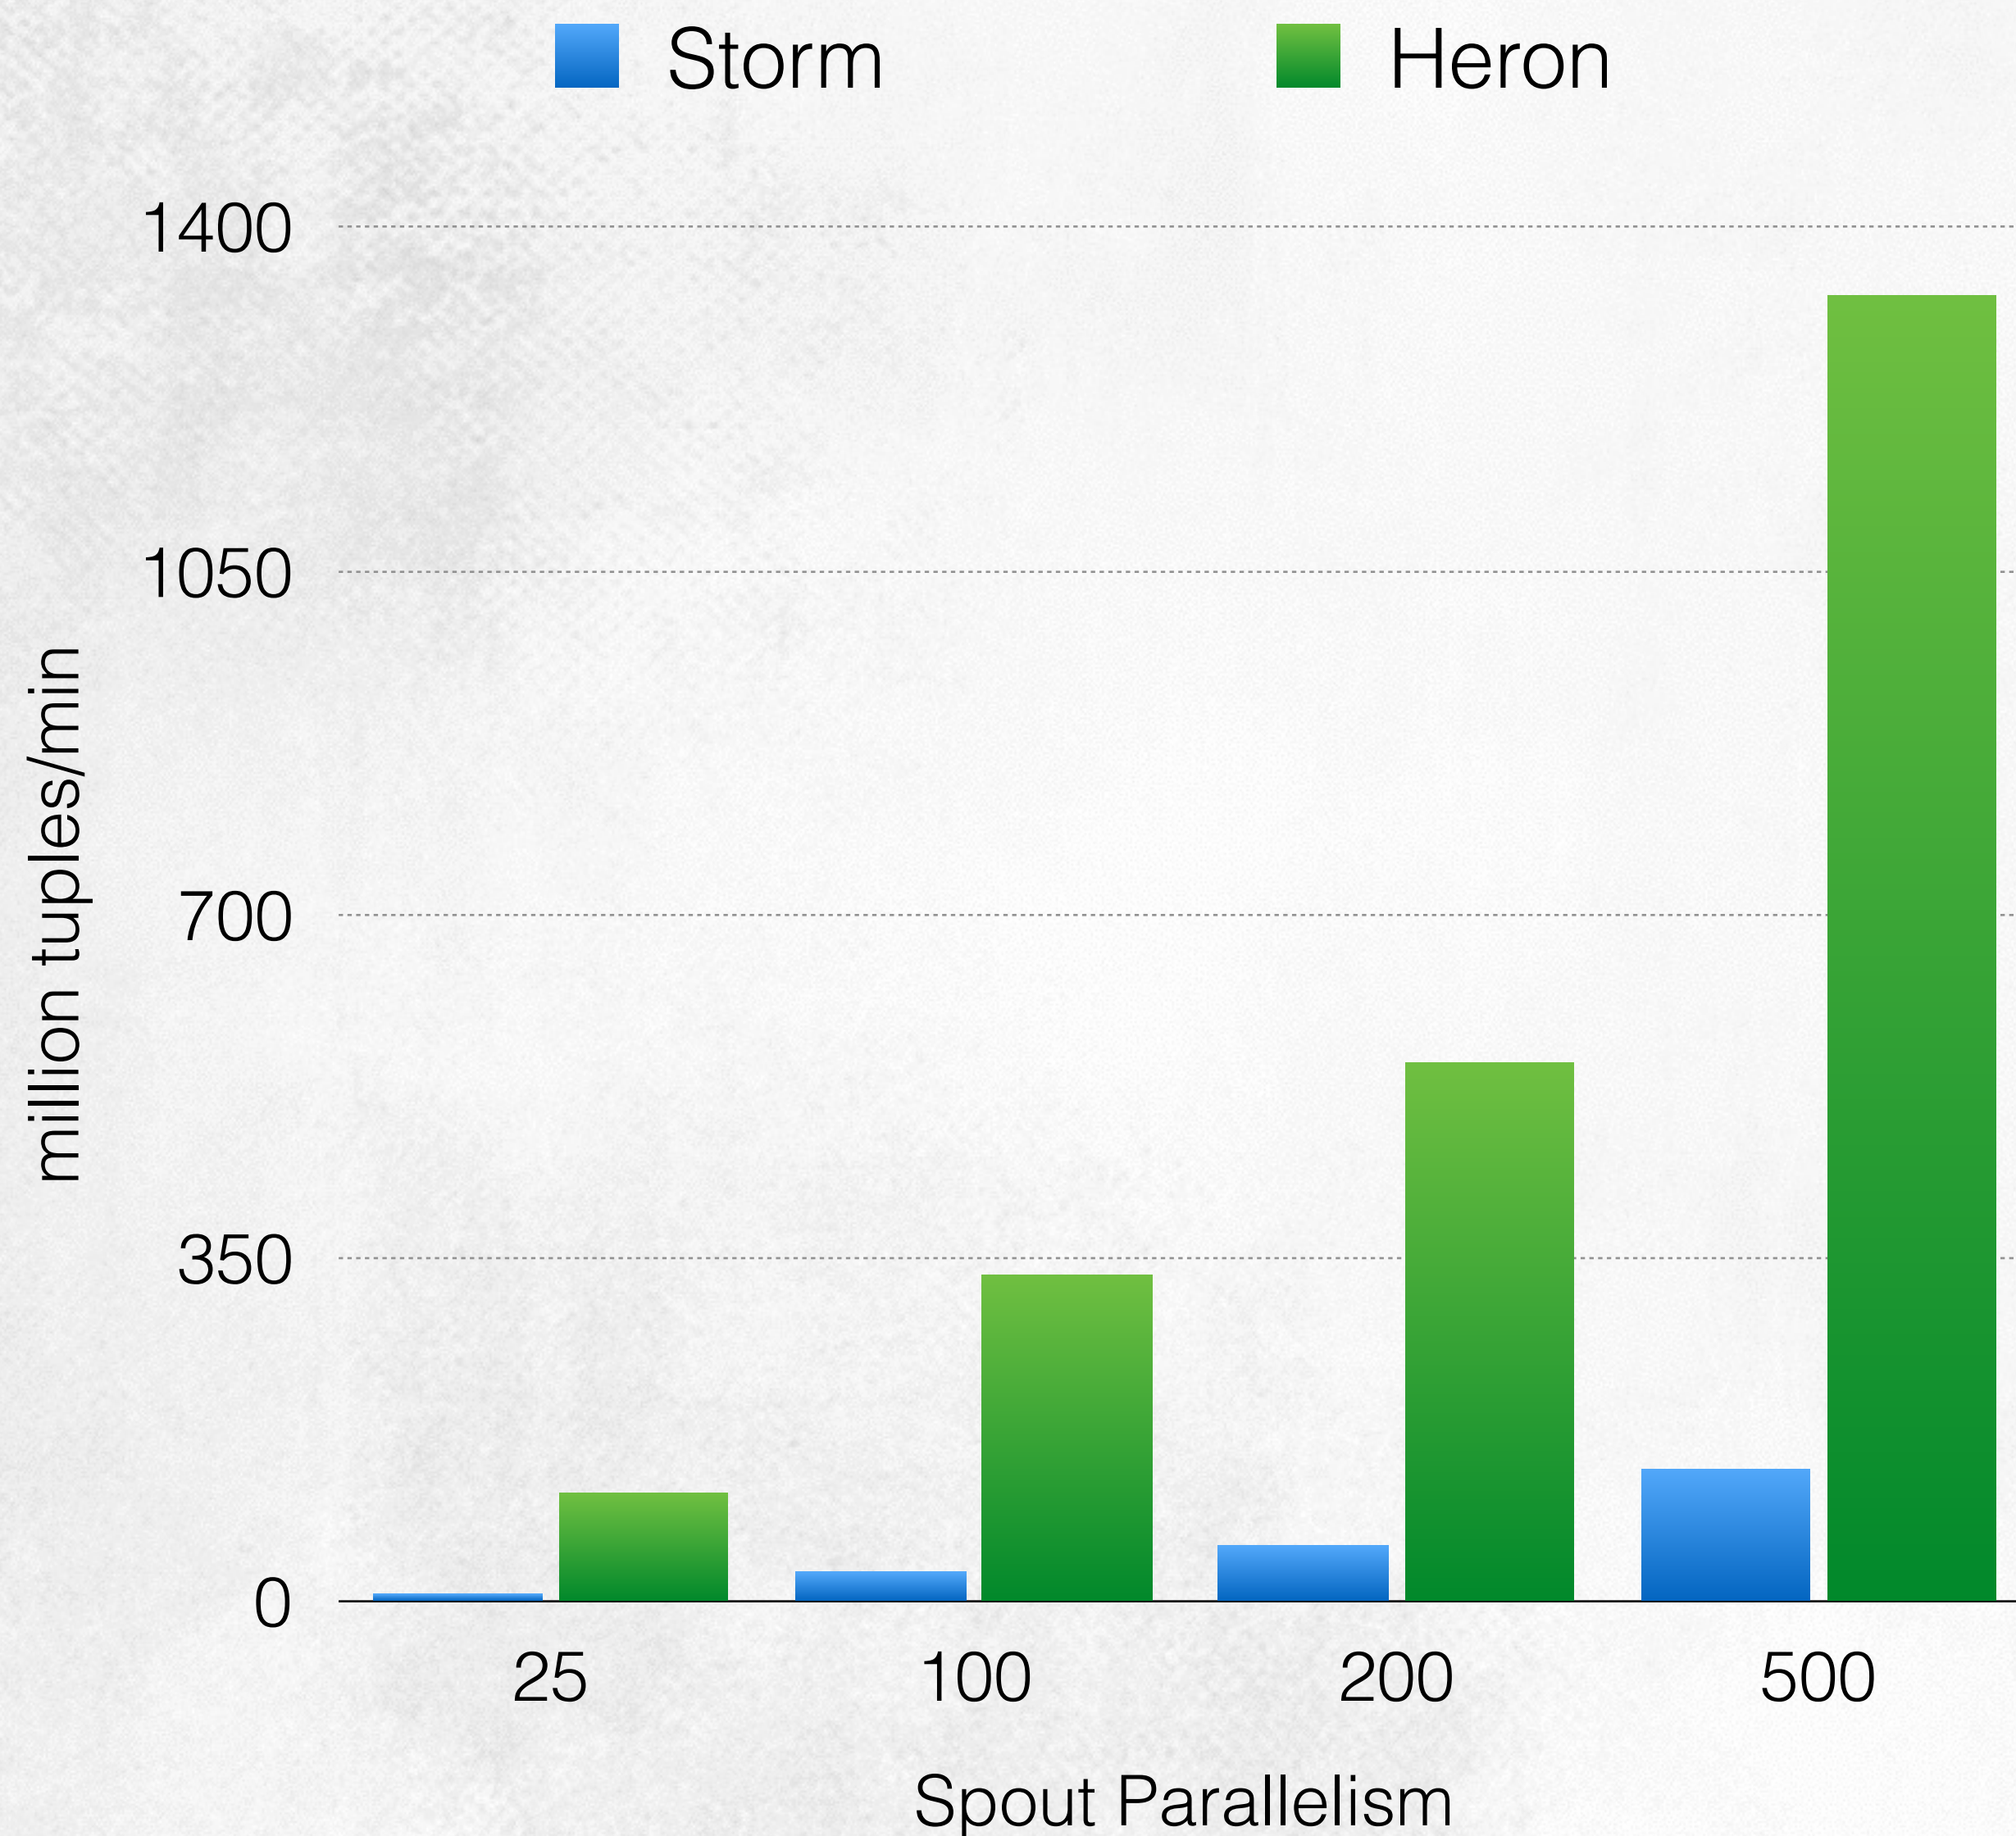




# HERON PERFORMANCE

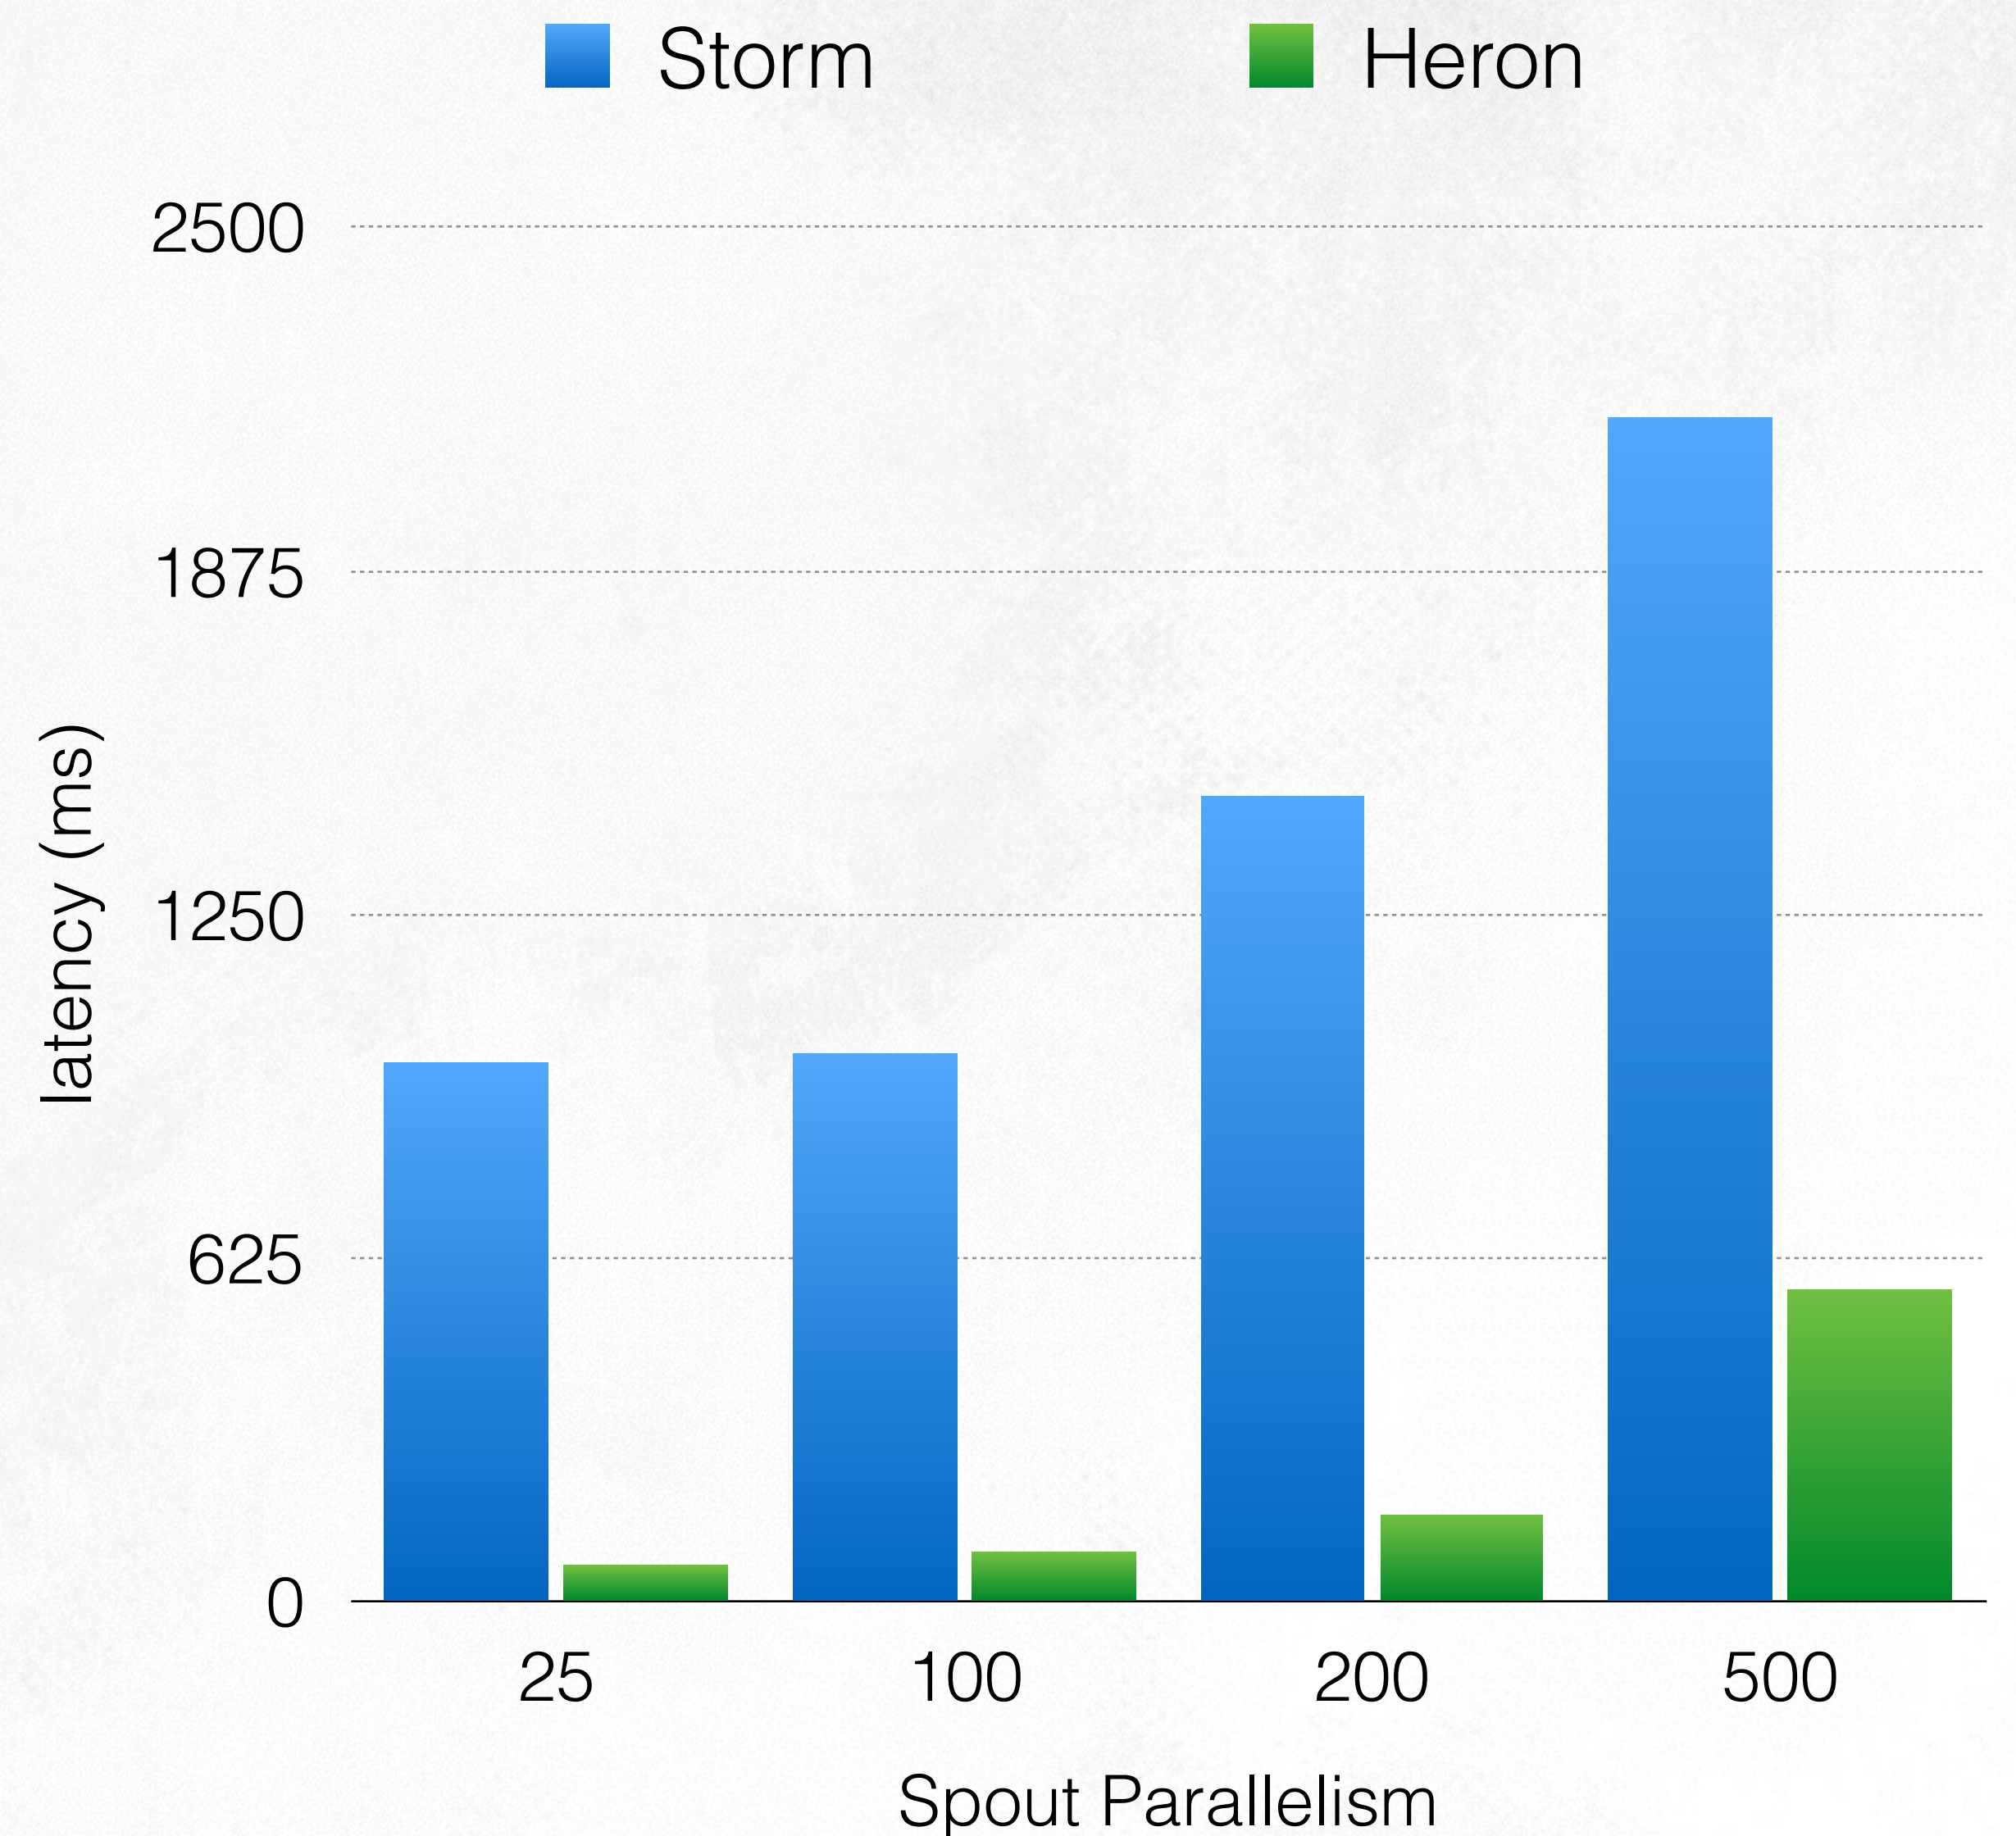
Word count topology – Acknowledgements enabled

## Throughput

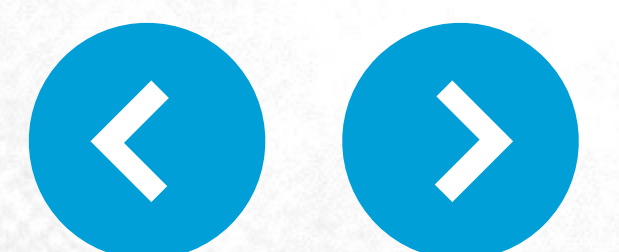


10-14x

## Latency

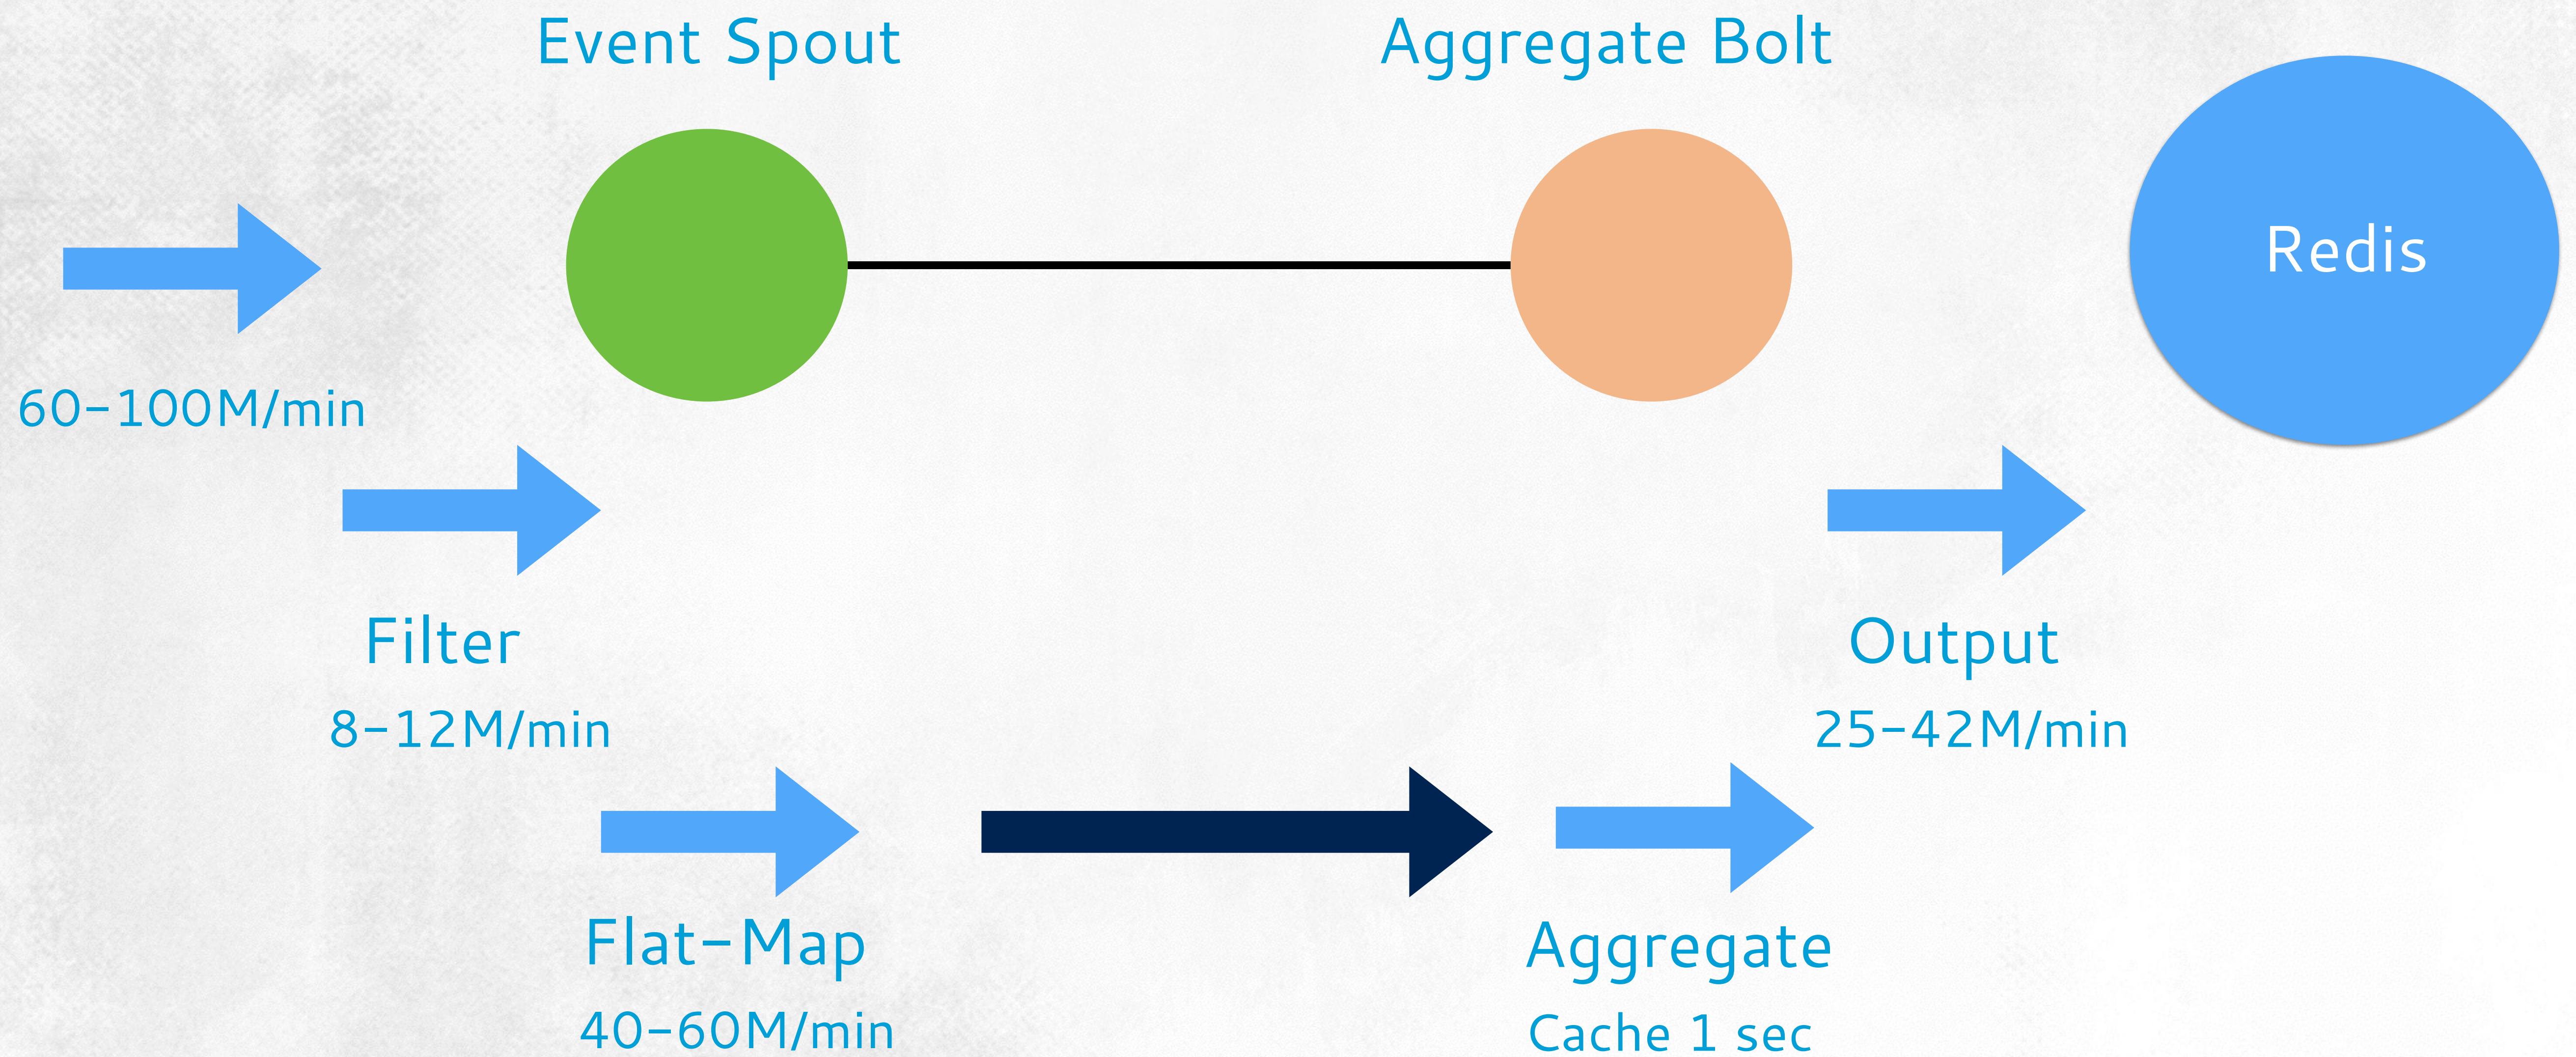


5-15x





# HERON RESOURCE USAGE





# RESOURCE CONSUMPTION

	Cores Requested	Cores Used	Memory Requested (GB)	Memory Used
Redis	24	2-4	48	N/A
Heron	120	30-50	200	180



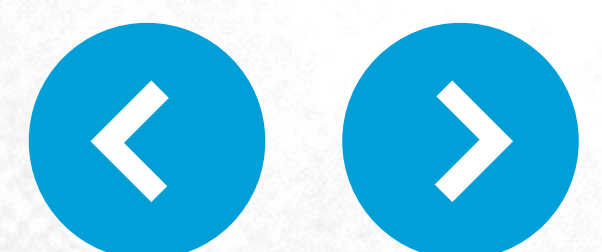
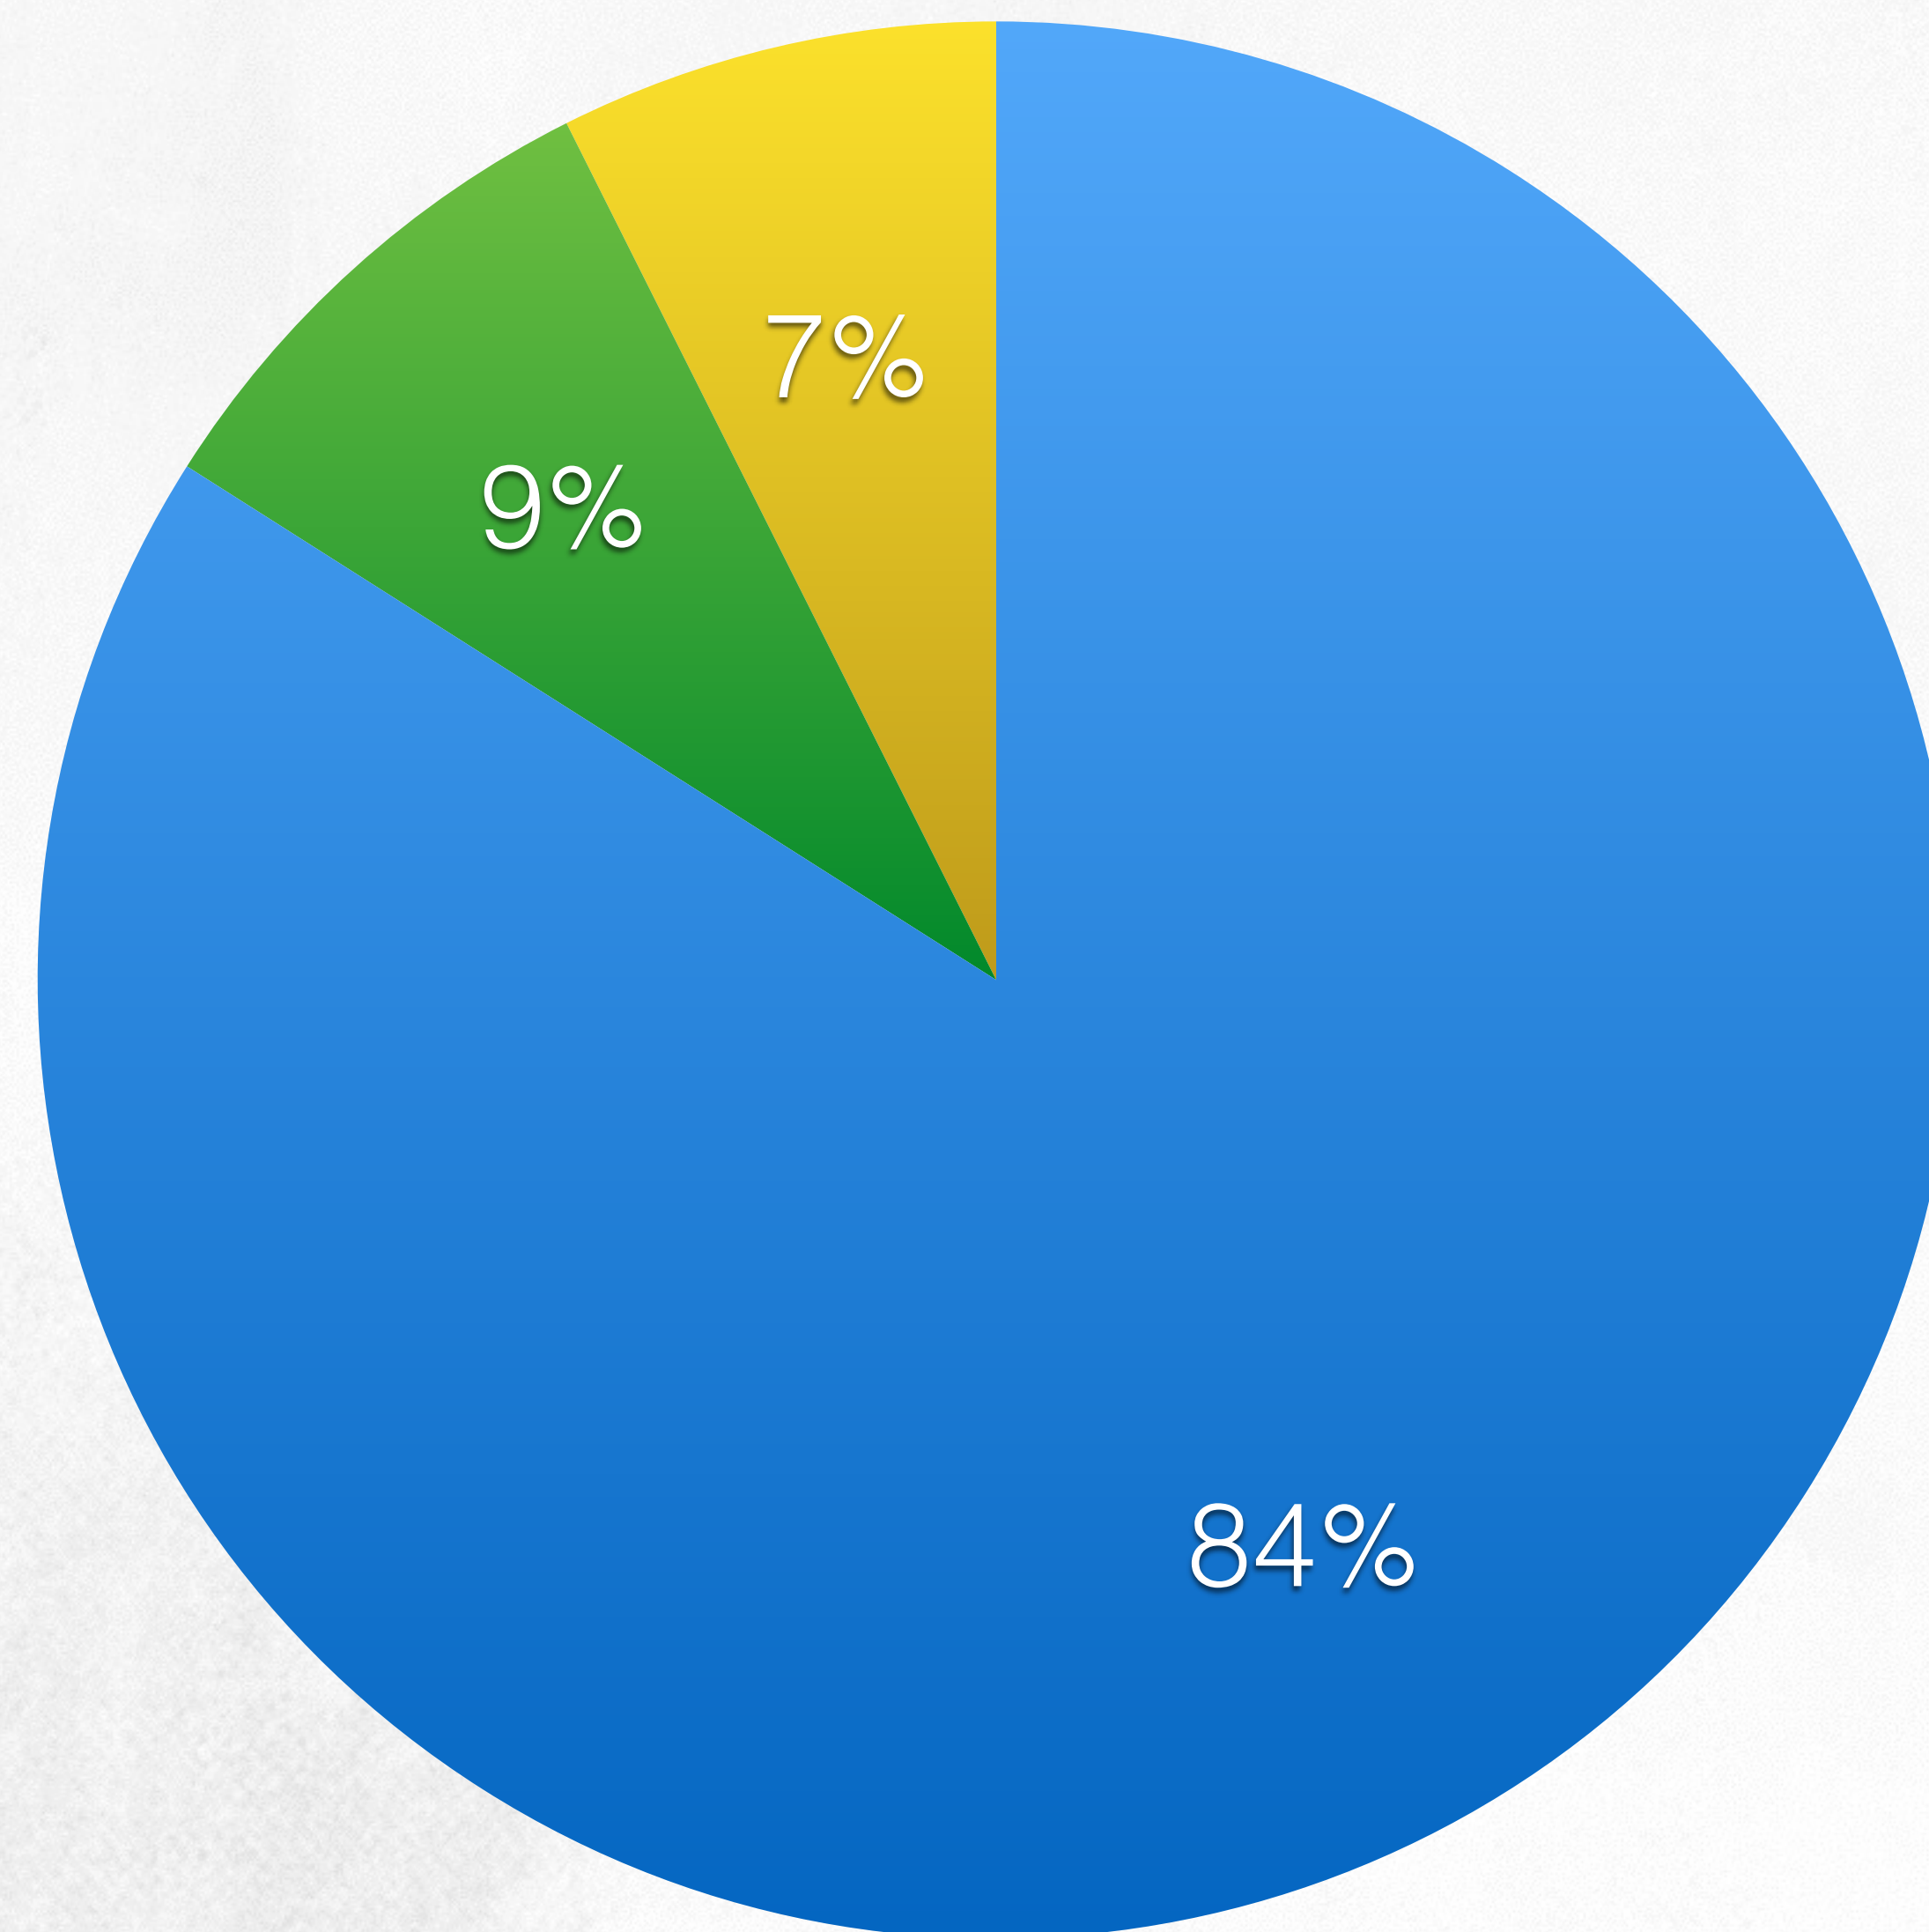


# RESOURCE CONSUMPTION

● Spout Instances

● Bolt Instances

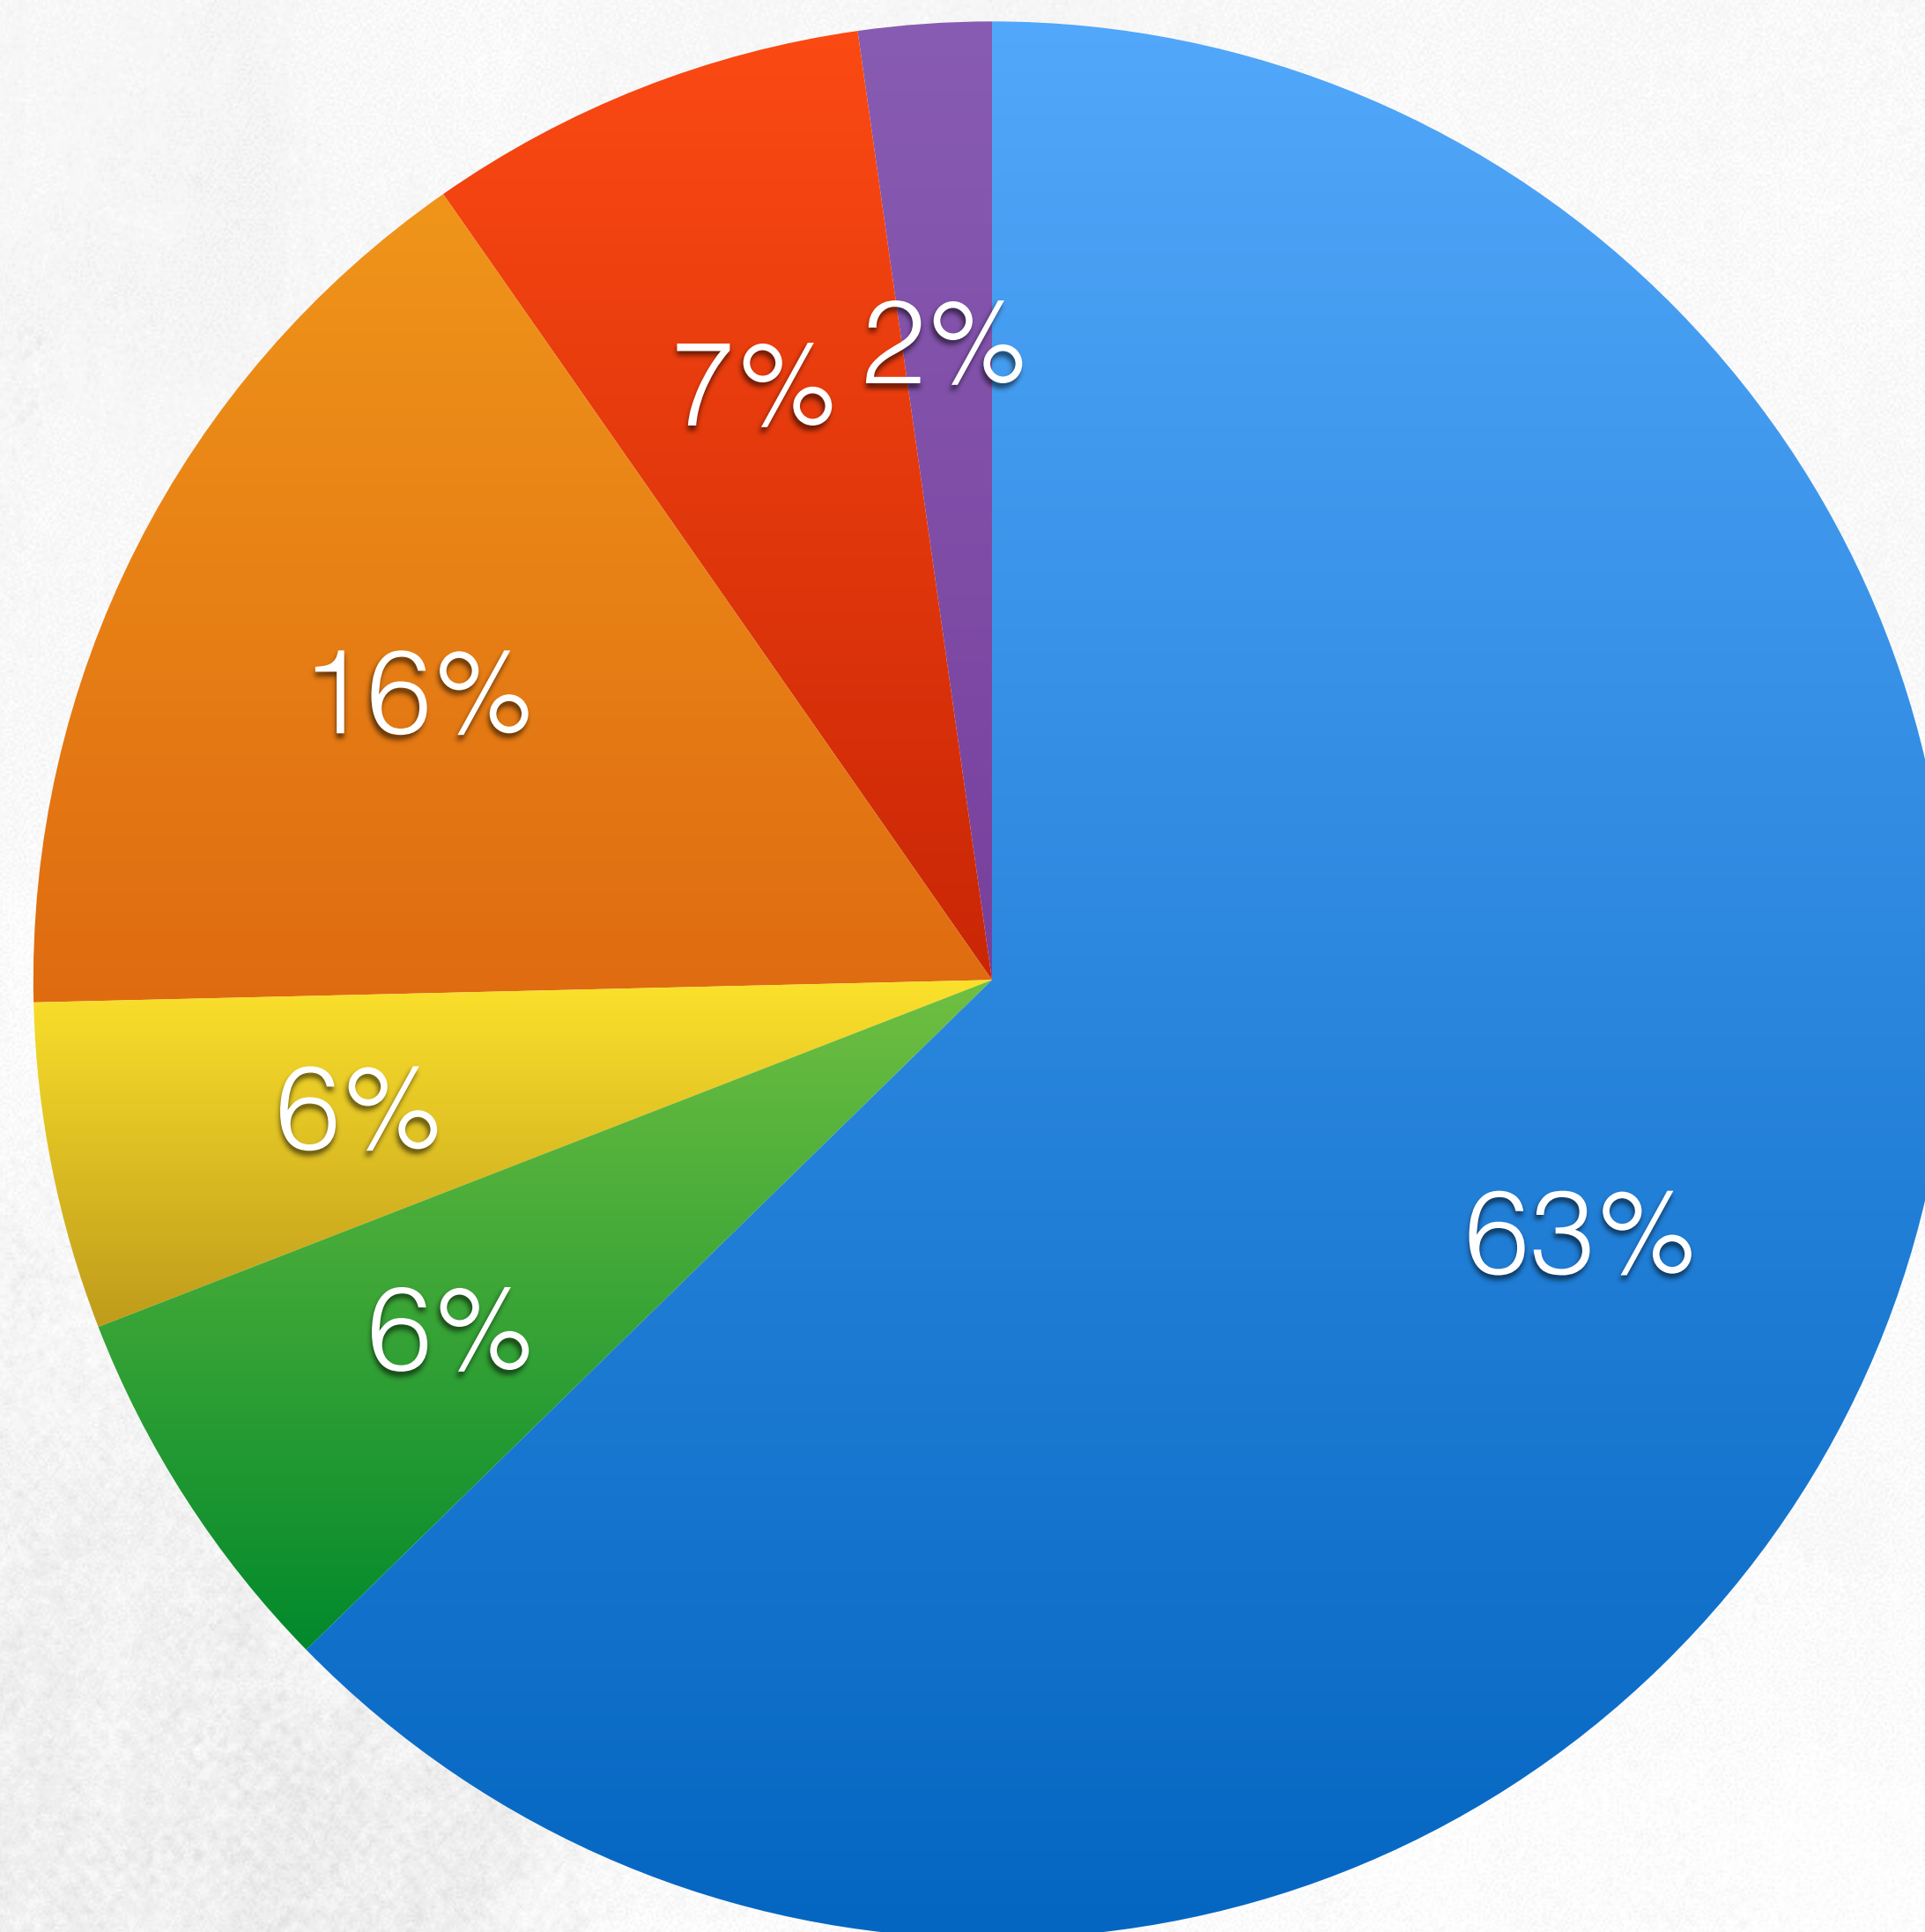
● Heron Overhead





# PROFILING SPOUTS

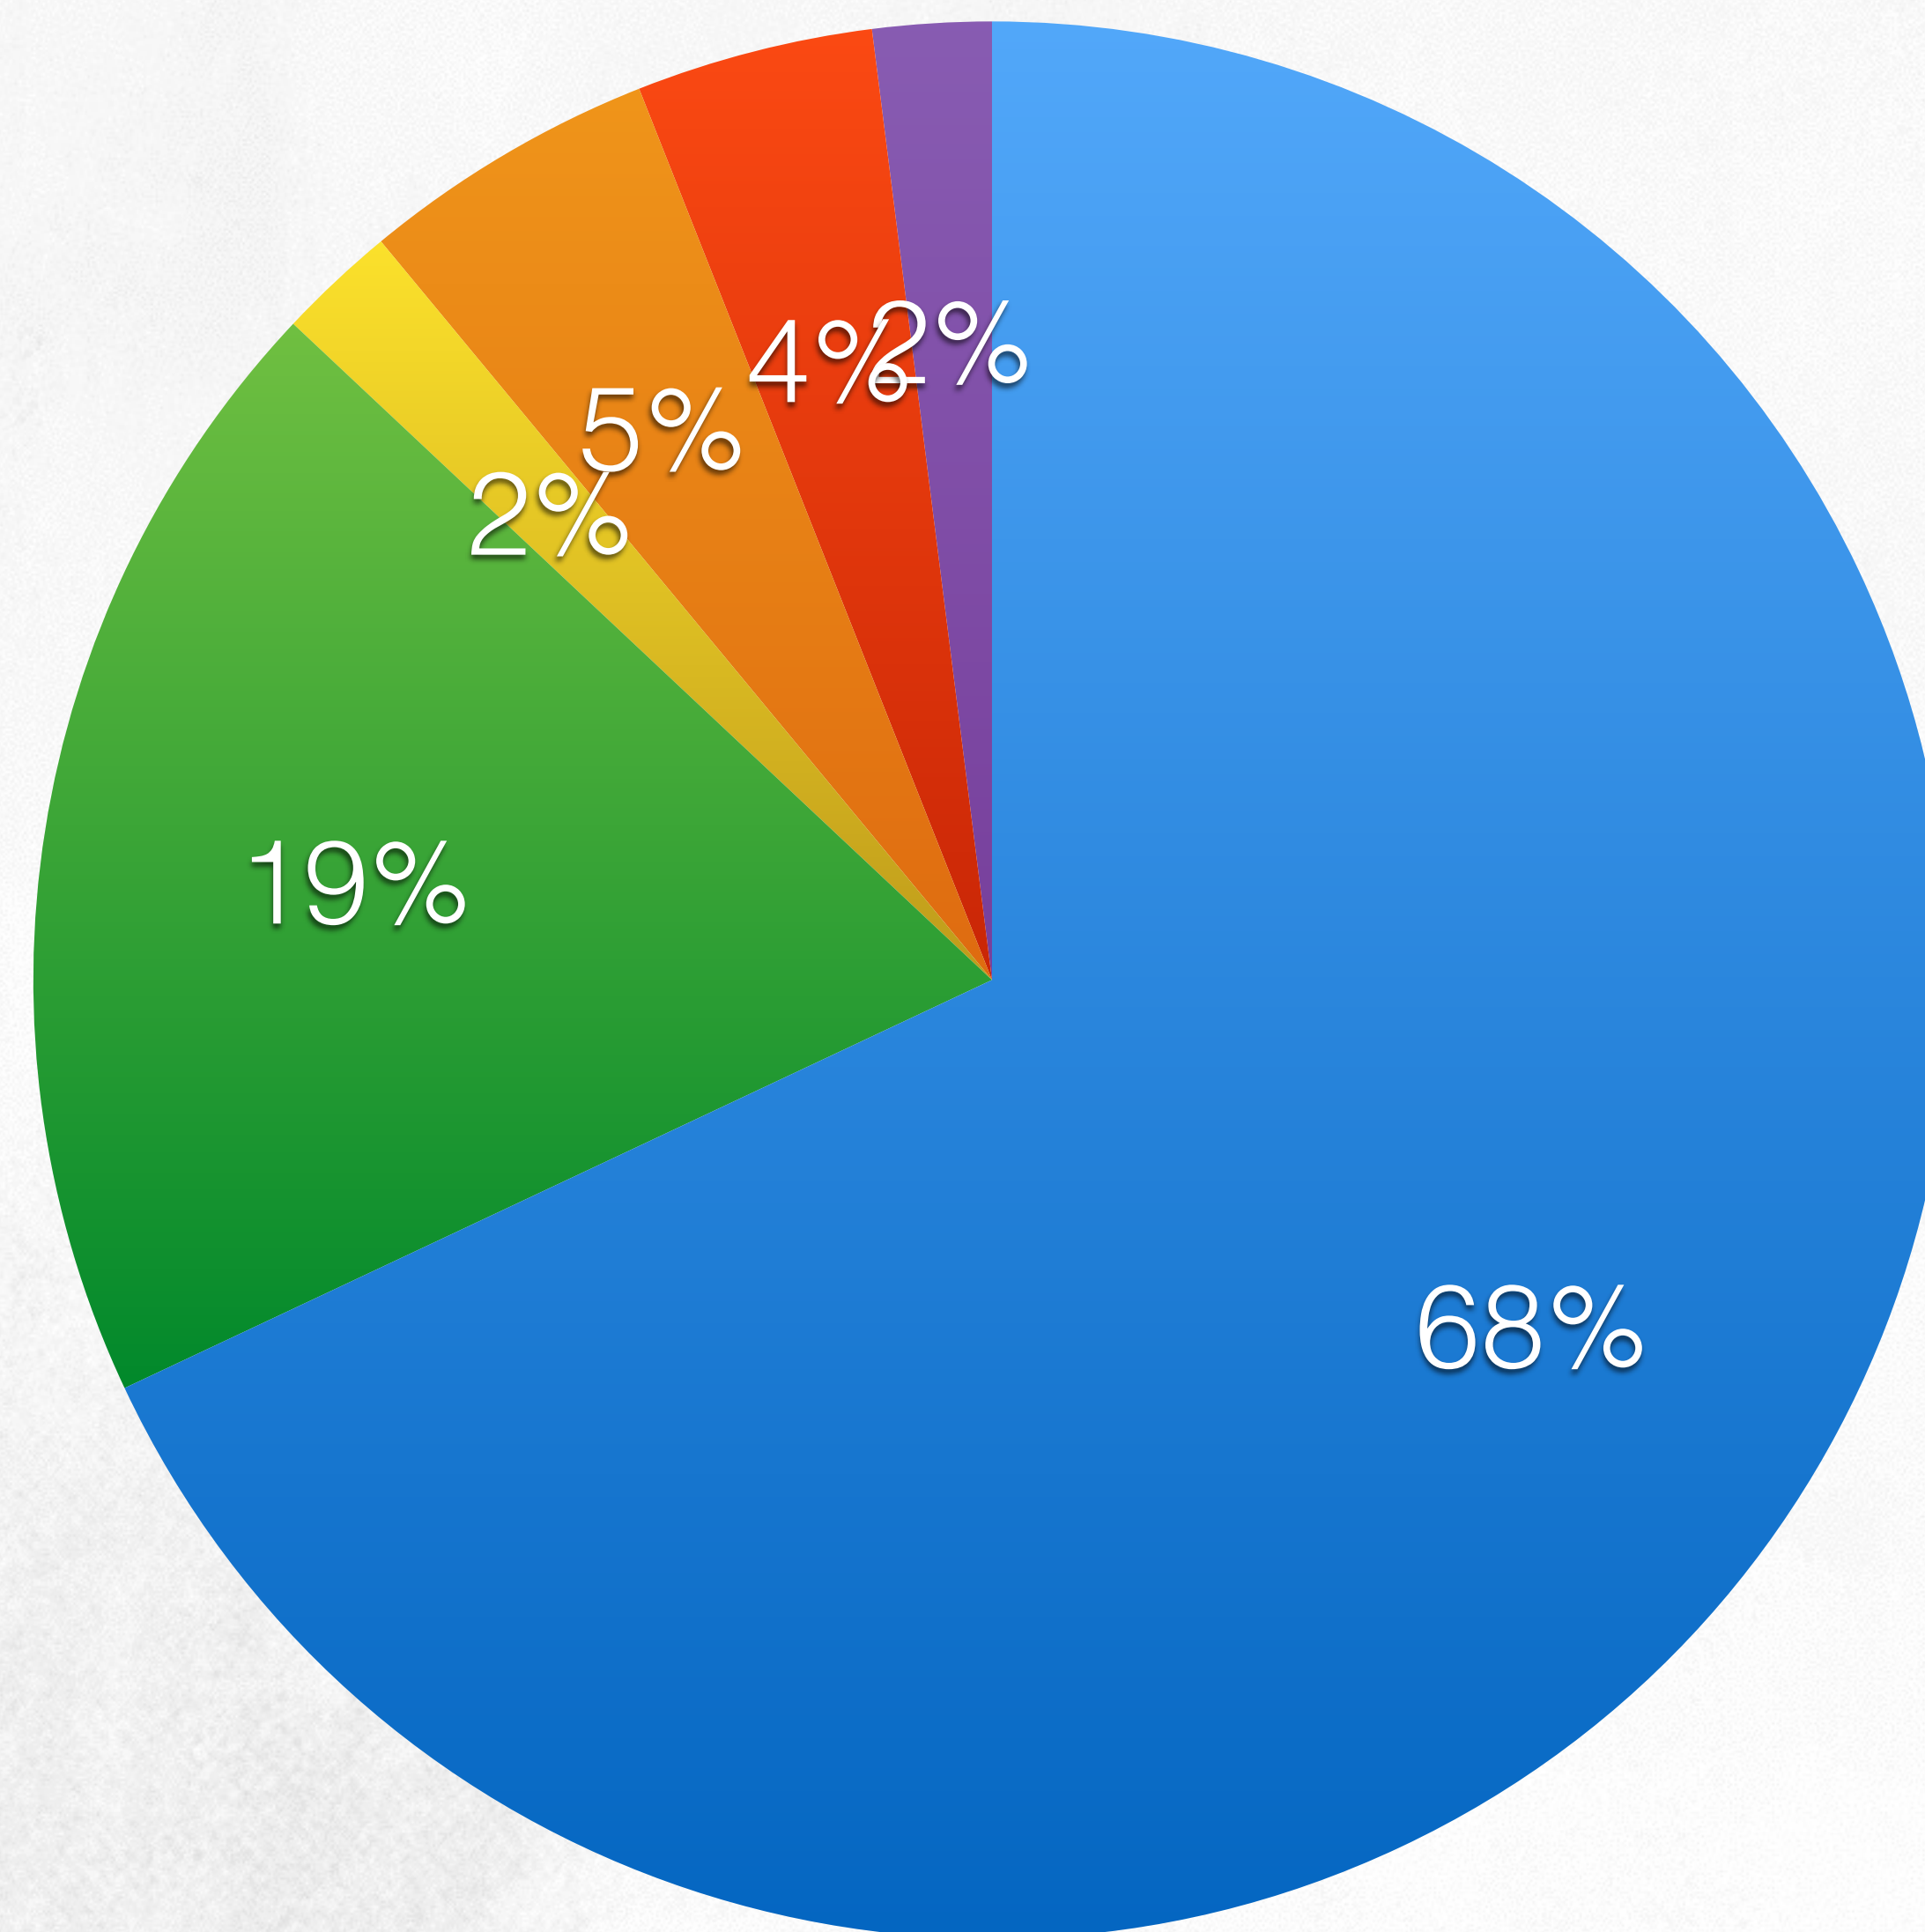
● Deserialize ● Parse/Filter ● Mapping ● Kafka Iterator ● Kafka Fetch ● Rest





# PROFILING BOLTS

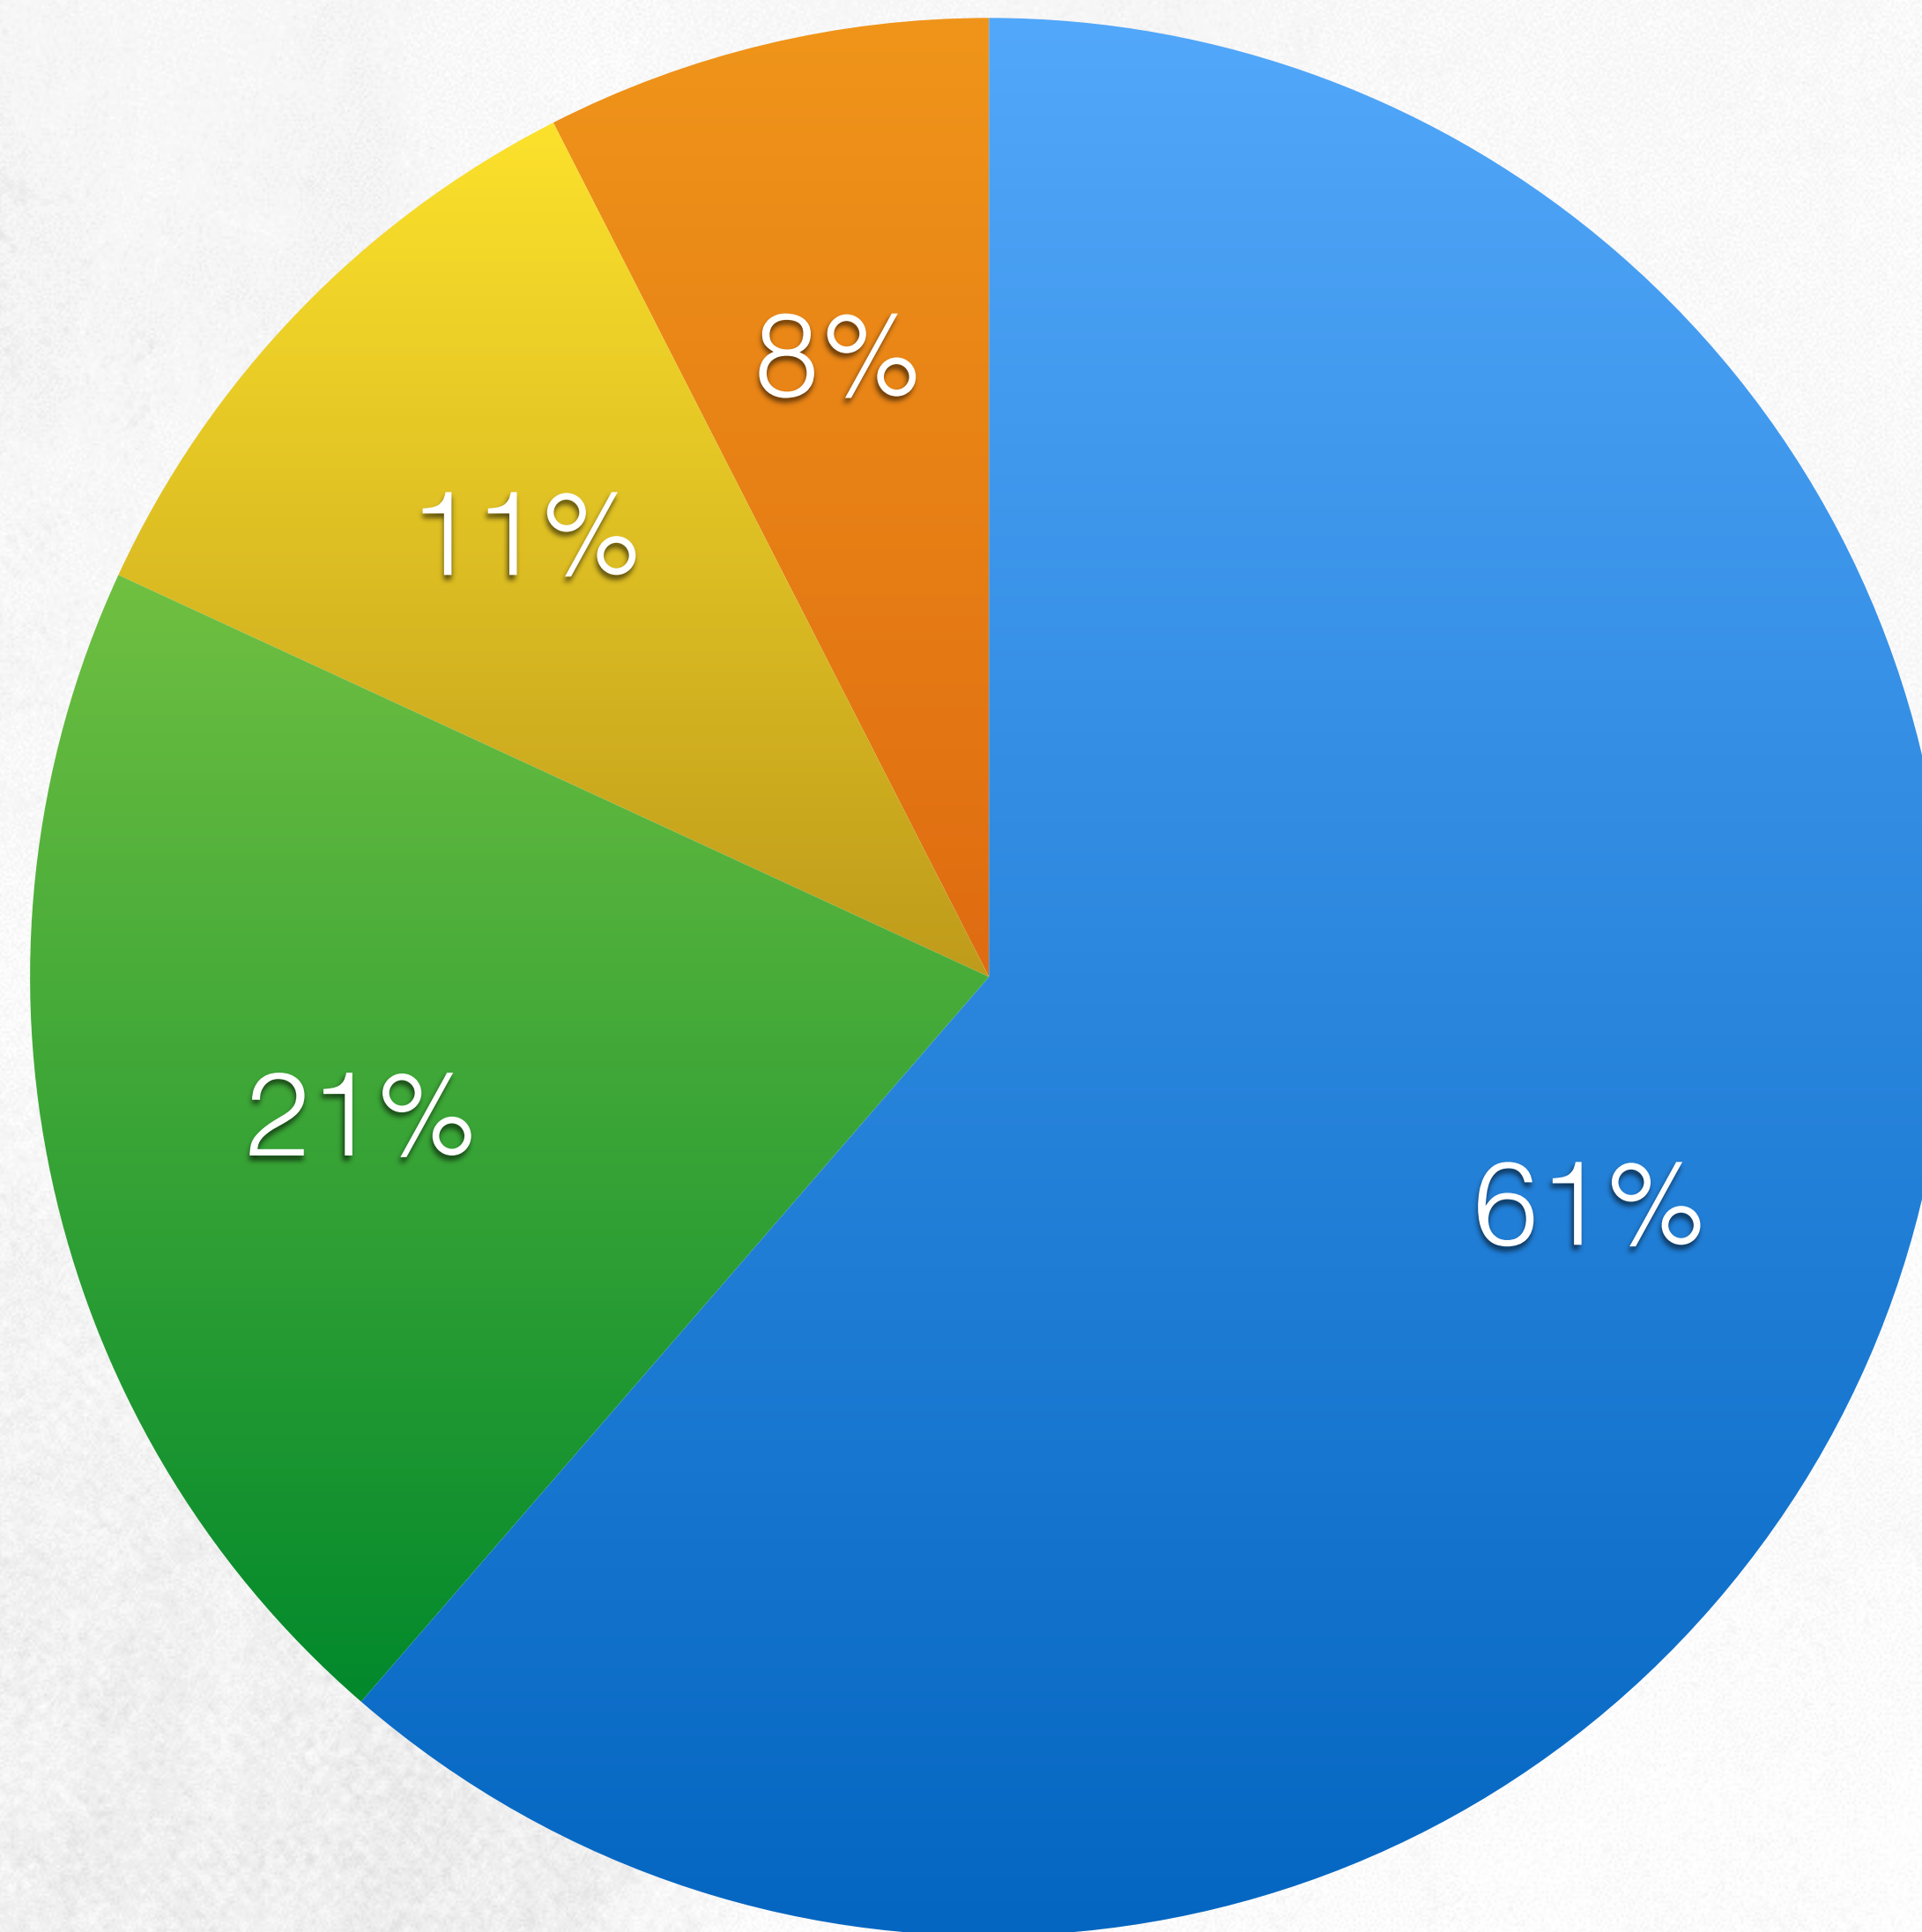
● Write Data ● Serialize ● Deserialize ● Aggregation ● Data Transport ● Rest





# RESOURCE CONSUMPTION – BREAKDOWN

● Fetching Data    ● User Logic    ● Heron Usage    ● Writing Data







# HERON BACK PRESSURE



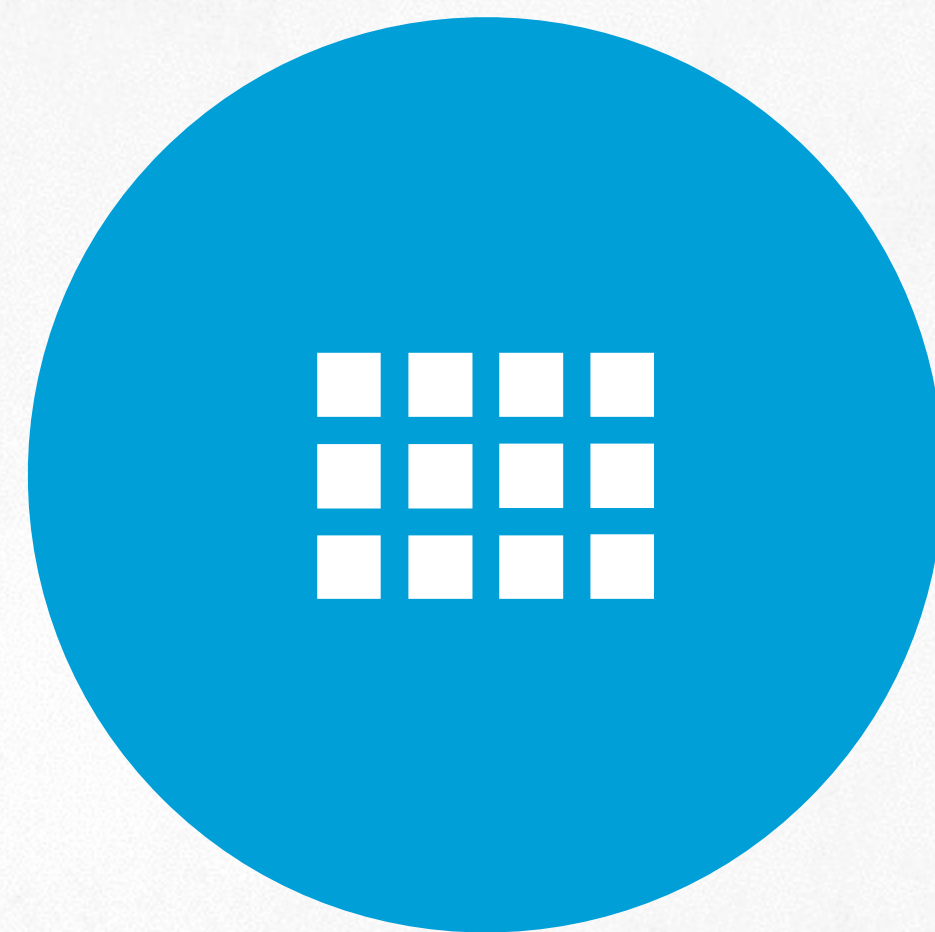


# BACK PRESSURE AND STRAGGLERS

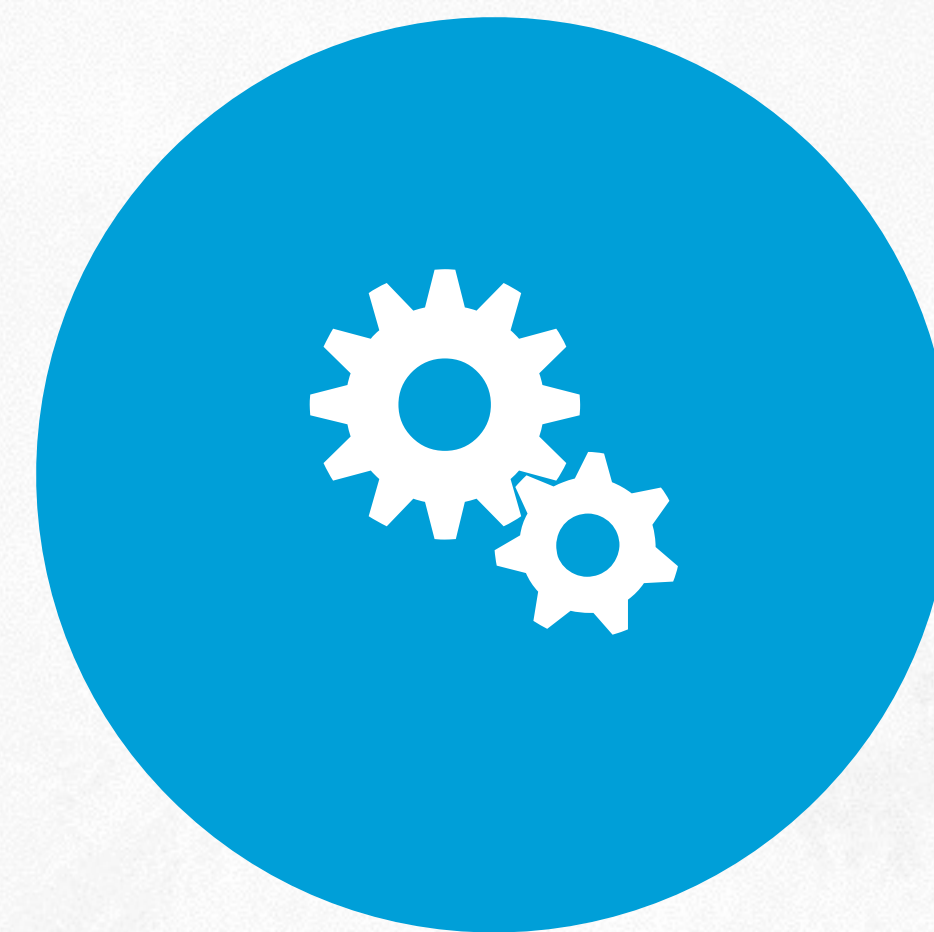
**Stragglers** are the norm in a multi-tenant distributed systems  
Bad machine, inadequate provisioning and hot keys



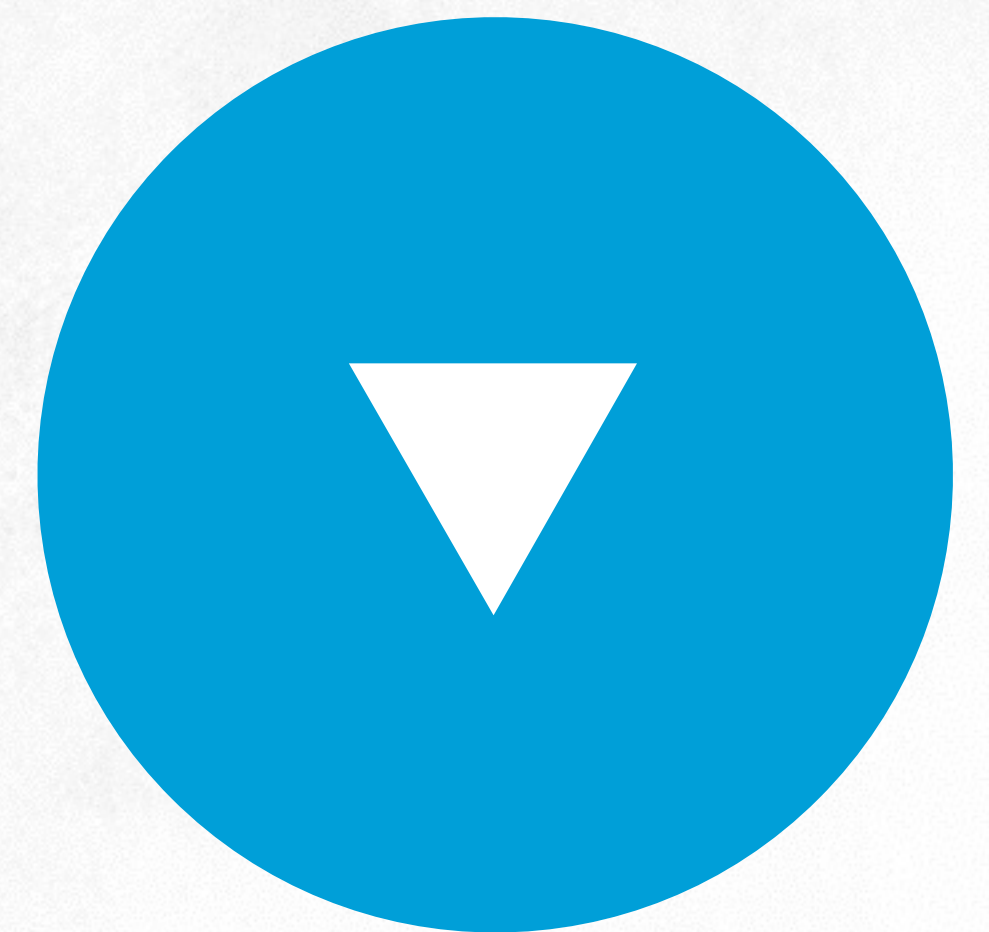
PROVIDES  
PREDICTABILITY



PROCESSES  
DATA AT  
MAXIMUM  
RATE



REDUCE  
RECOVERY  
TIMES



HANDLES  
TEMPORARY  
SPIKES



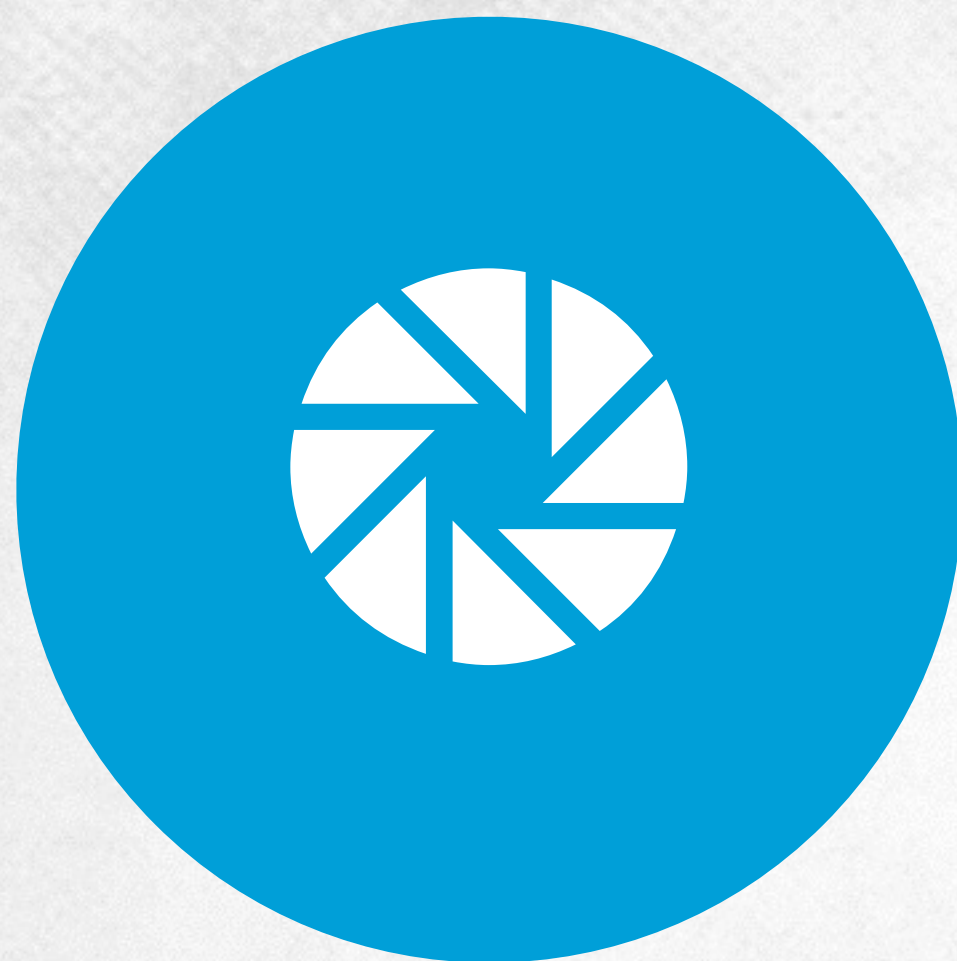


# BACK PRESSURE **AND STRAGGLERS**



## MOST SCENARIOS BACK PRESSURE RECOVERS

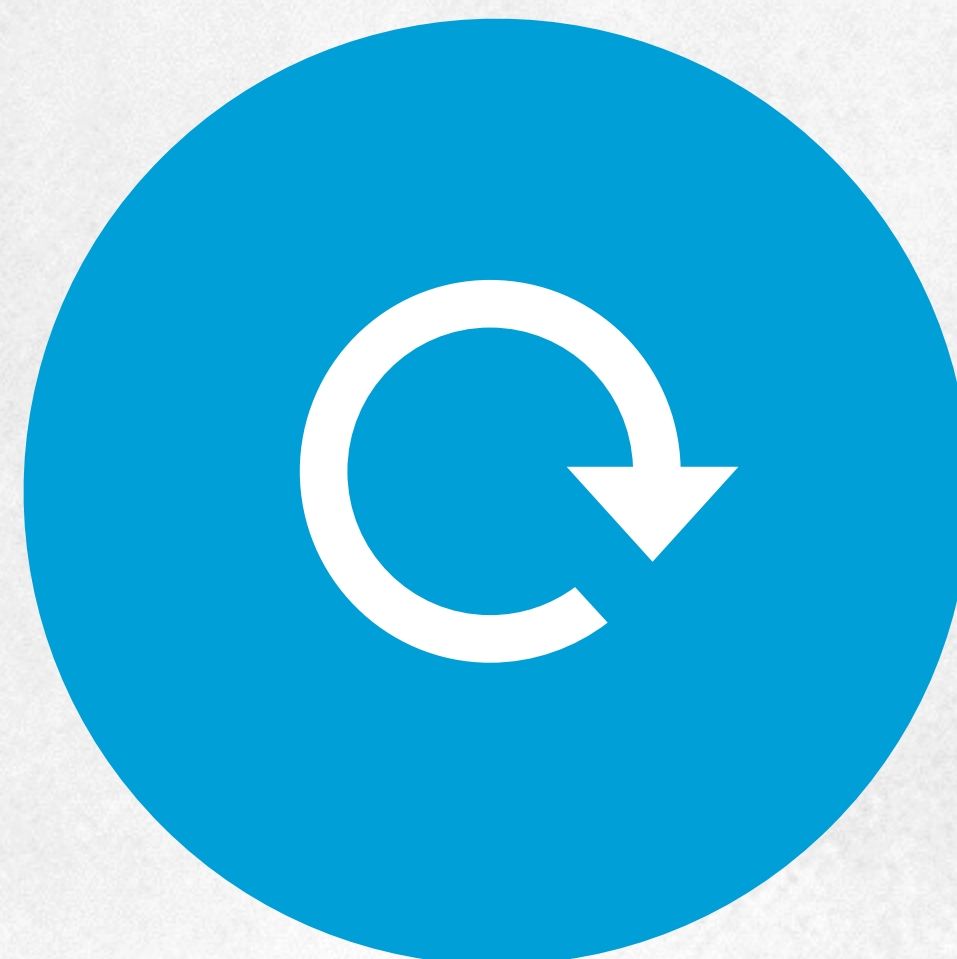
Without any manual intervention



## SUSTAINED BACK PRESSURE

Irrecoverable GC cycles

Bad or faulty host



## SOMETIMES USER PREFER DROPPING OF DATA

Care about only latest data





# LOAD SHEDDING

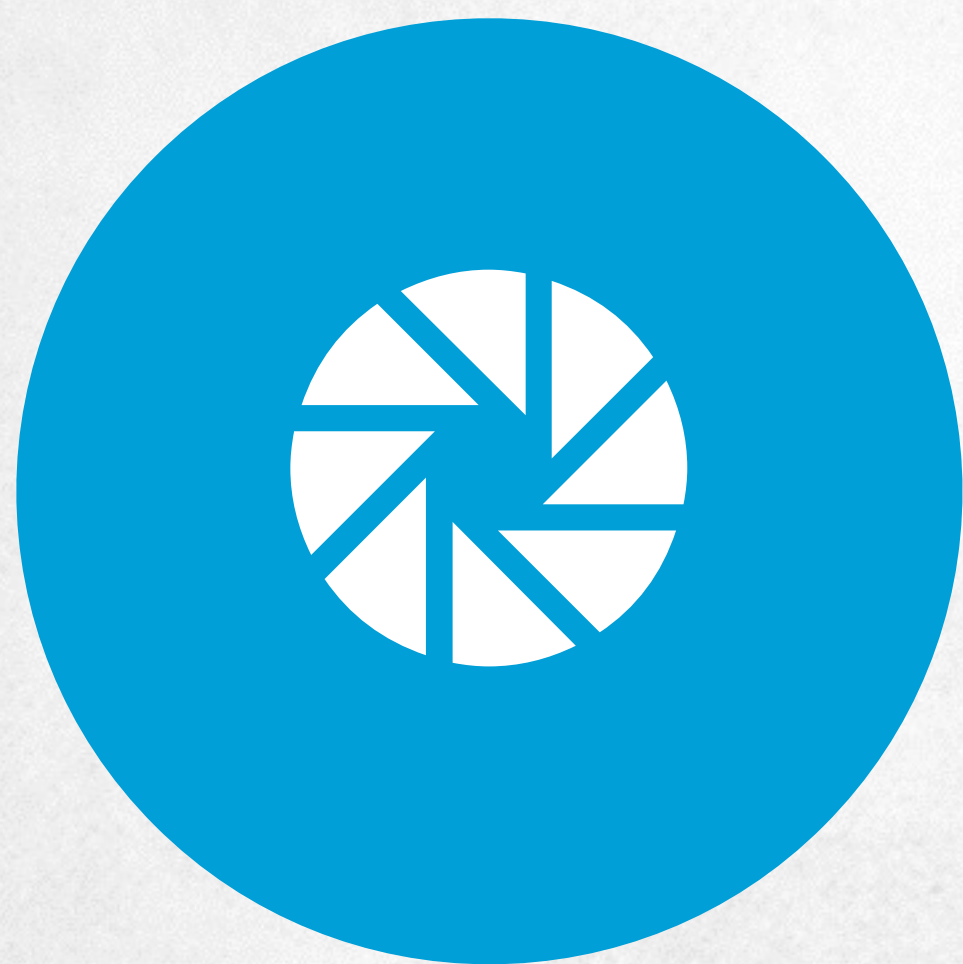


## SAMPLING BASED APPROACHES

Down sample the incoming stream and scale up the results

Easy to reason if the sampling is uniform

Hard to achieve uniformity across distributed spouts



## DROP BASED APPROACHES

Simply drop older data

Spouts takes a lag threshold and a lag adjustment value

Works well in practice





# CURIOUS TO **LEARN MORE...**

## **Streaming@Twitter**

Maosong Fu, Sailesh Mittal, Vikas Kedigehalli, Karthik Ramasamy, Michael Barry, Andrew Jorgensen, Christopher Kellogg, Neng Lu, Bill Graham, Jingwei Wu

Twitter, Inc.

## **Twitter Heron: Stream Processing at Scale**

Sanjeev Kulkarni, Nikunj Bhagat, Maosong Fu, Vikas Kedigehalli, Christopher Kellogg,

Sailesh Mittal, Jignesh M. Patel<sup>\*1</sup>, Karthik Ramasamy, Siddarth Taneja

@sanjeevrk, @challenger\_nik, @Louis\_Fumaosong, @vikkyrk, @cckellogg,  
@saileshmittal, @pateljm, @karthikz, @staneja

Twitter, Inc., \*University of Wisconsin – Madison

## **Storm @Twitter**

Ankit Toshniwal, Siddarth Taneja, Amit Shukla, Karthik Ramasamy, Jignesh M. Patel\*, Sanjeev Kulkarni, Jason Jackson, Krishna Gade, Maosong Fu, Jake Donham, Nikunj Bhagat, Sailesh Mittal, Dmitriy Ryaboy

@ankitoshniwal, @staneja, @amits, @karthikz, @pateljm, @sanjeevrk,  
@jason\_j, @krishnagade, @Louis\_Fumaosong, @jakedonham, @challenger\_nik, @saileshmittal, @squarecog  
Twitter, Inc., \*University of Wisconsin – Madison





**#ThankYou**

FOR LISTENING







# QUESTIONS AND ANSWERS



Go ahead. Ask away.





# HERON LOAD SHEDDING

