CC5212-1

Procesamiento Masivo de Datos

Otoño 2021

Lecture 6

Streaming: Kafka

Aidan Hogan

aidhog@gmail.com

Files

Files

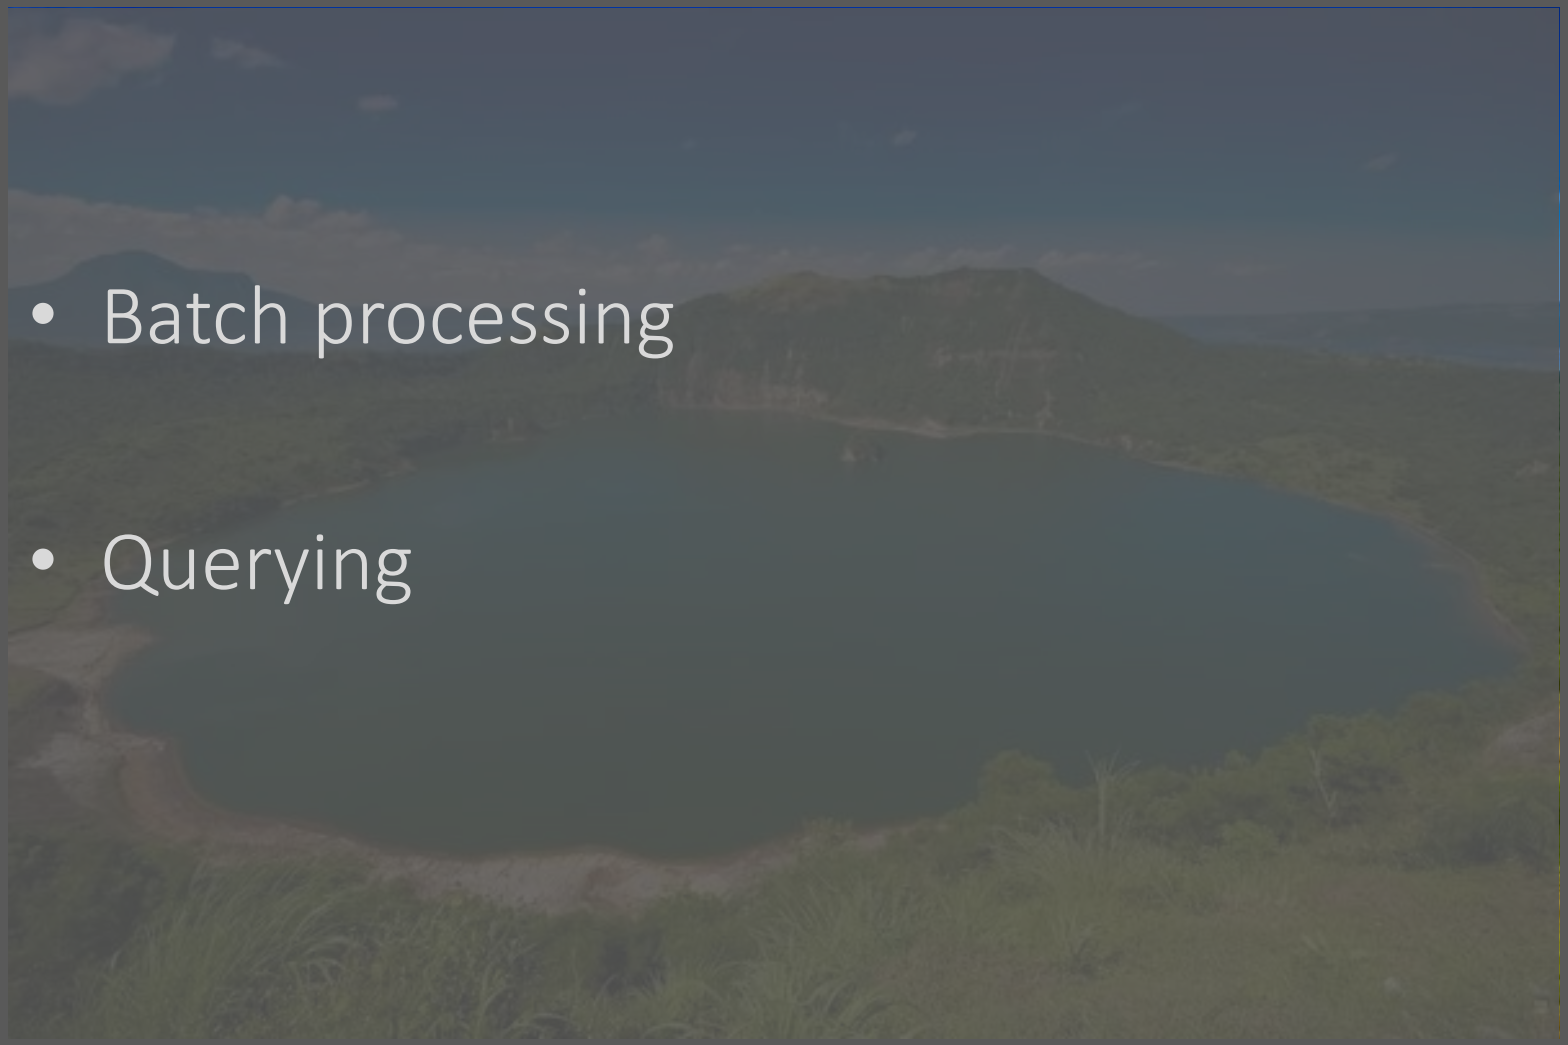

- Batch processing
- Querying

Streams

Streams

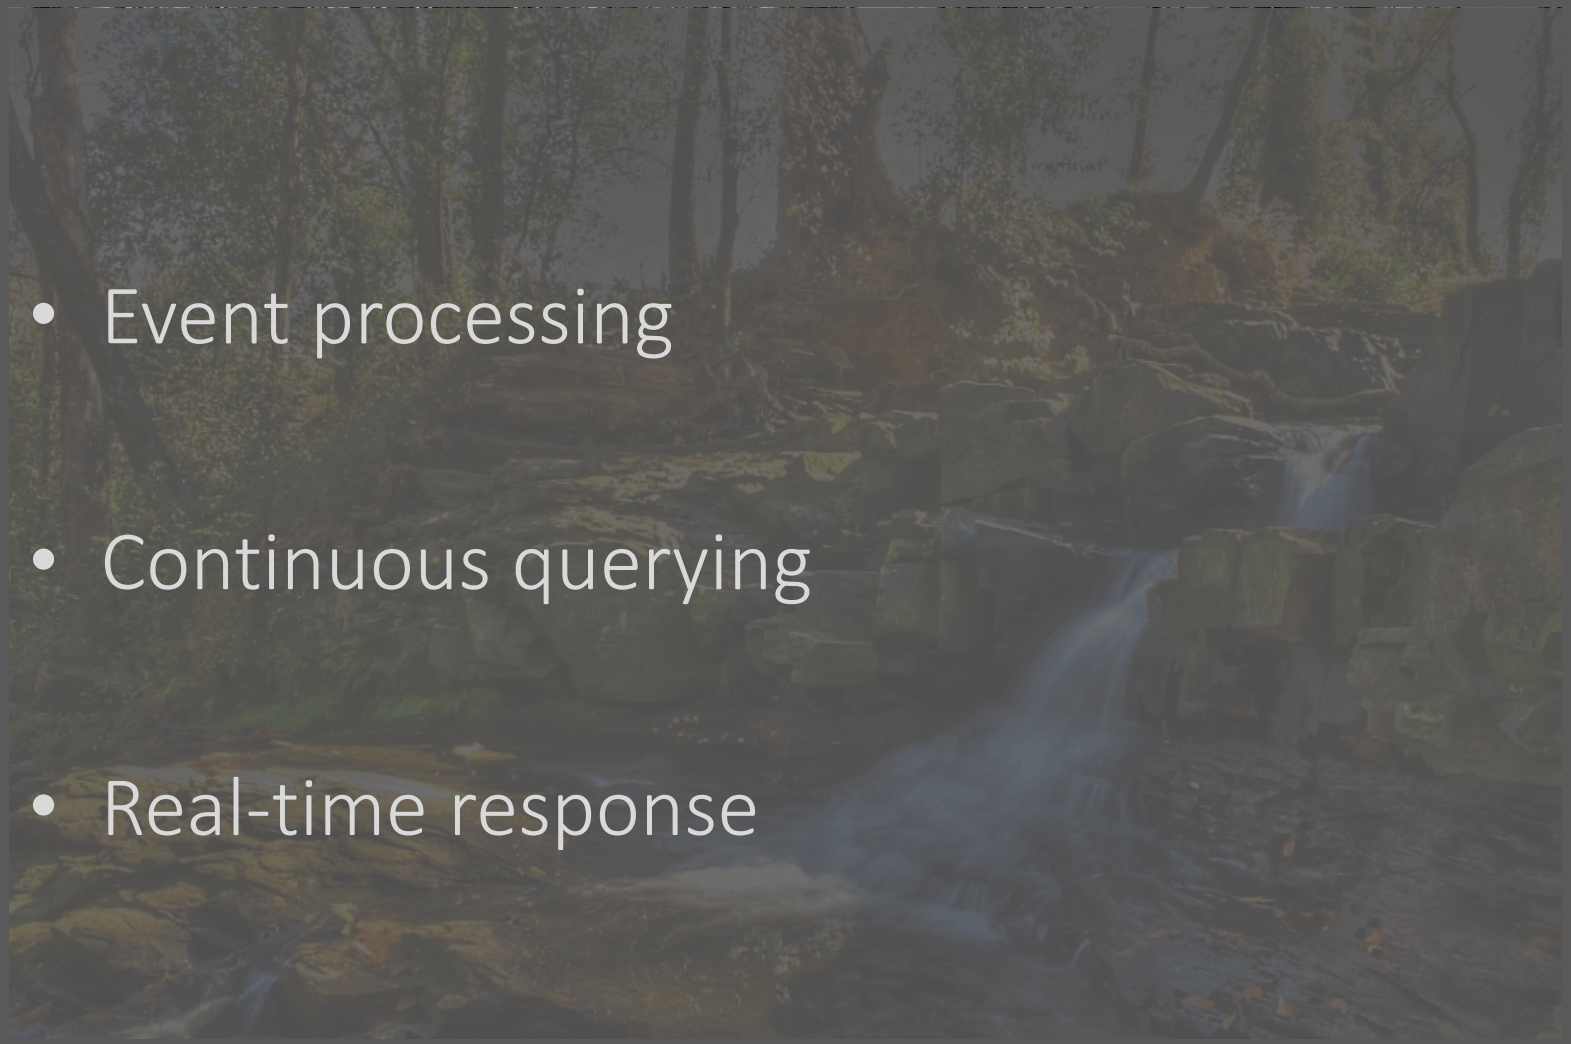

- Event processing
- Continuous querying
- Real-time response

Applications: Social Media Analytics

Applications: Social Media Analytics

- Event processing
  - Kitten video goes viral
  - Burst of tweets about earthquakes
- Continuous querying
  - Track sentiment for company's products
  - Monitor popular users tweeting about me
- Real-time response
  - Put Emergency Services on alert
  - Schedule Quality Control (QC) review

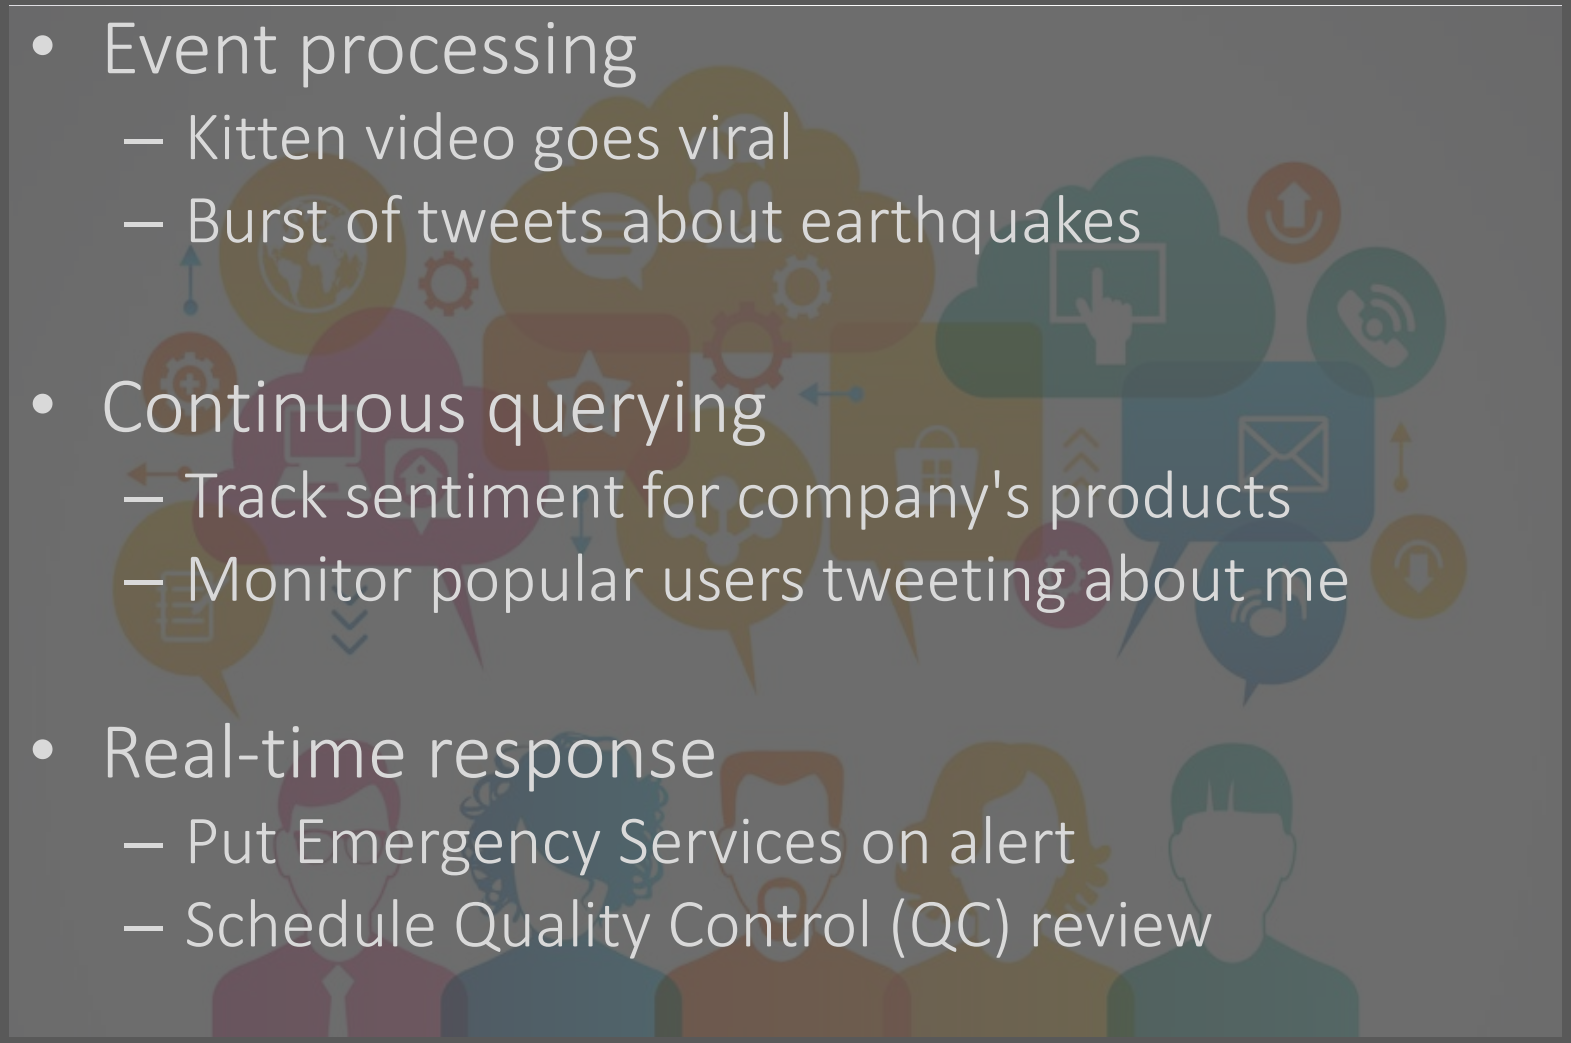

Applications: Log Monitoring

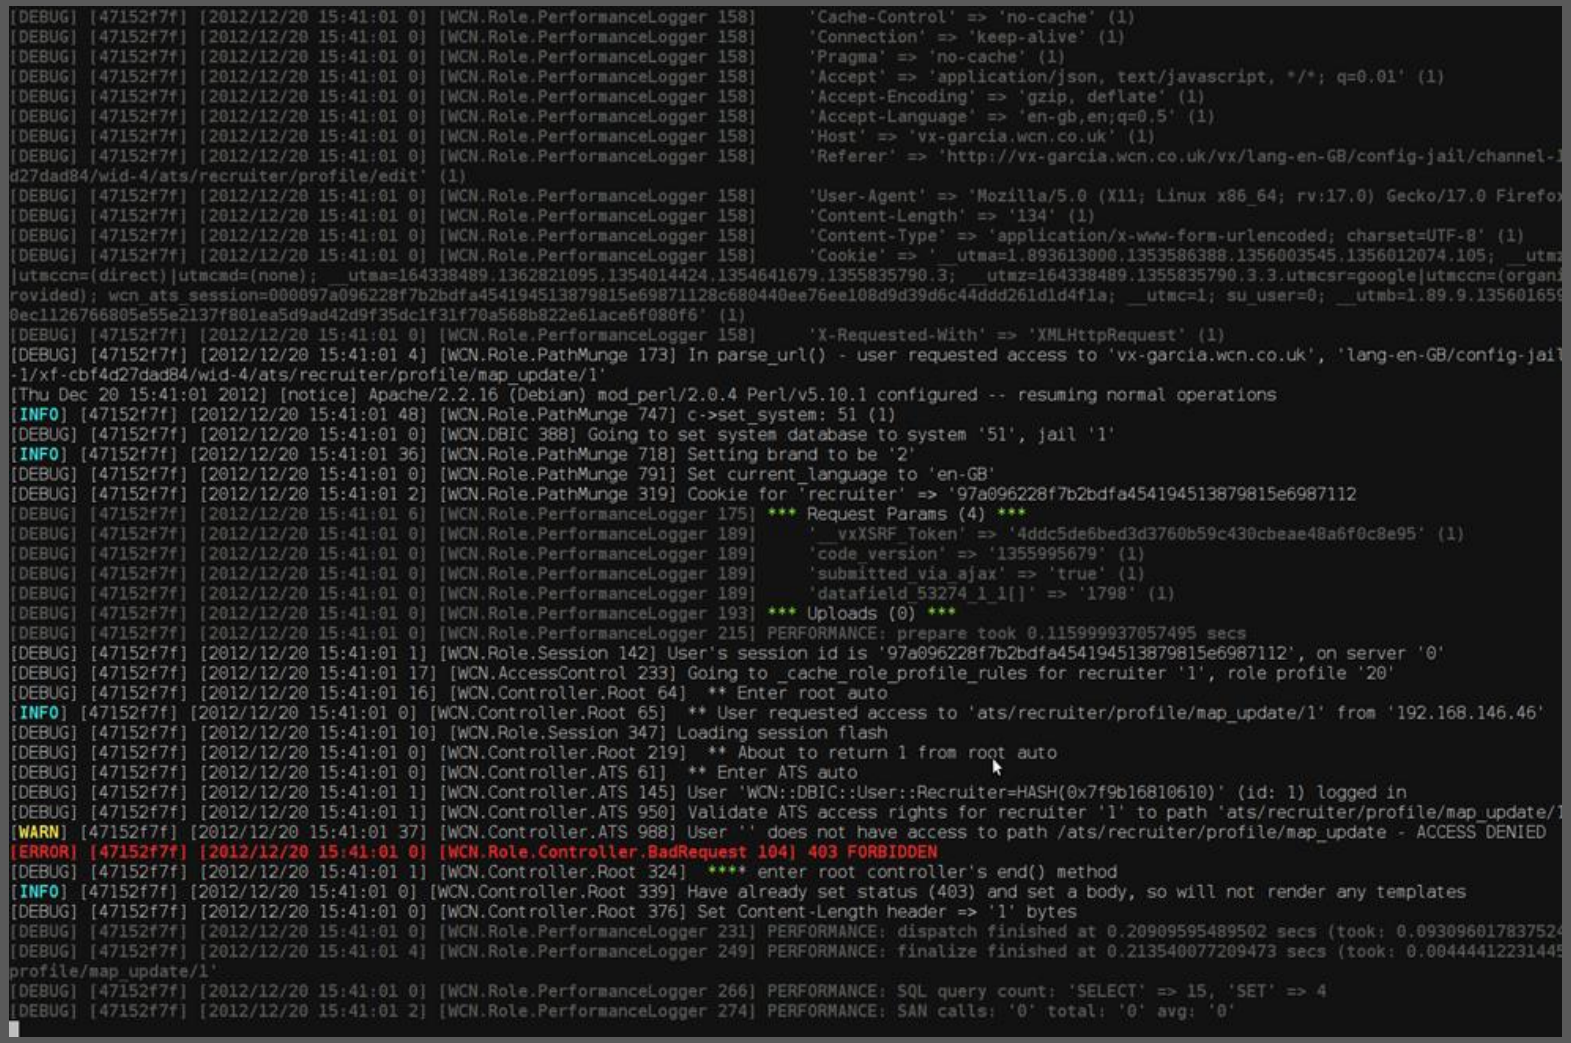

Applications: Log Monitoring

- Event processing
  - Burst of log messages
  - Critical error message
- Continuous querying
  - Track most critical fixes today
  - Monitor memory leaks in new release
- Real-time response
  - Disable unsafe feature in a web-site
  - Automatically fire new developer

Applications: Finance

Applications: Finance

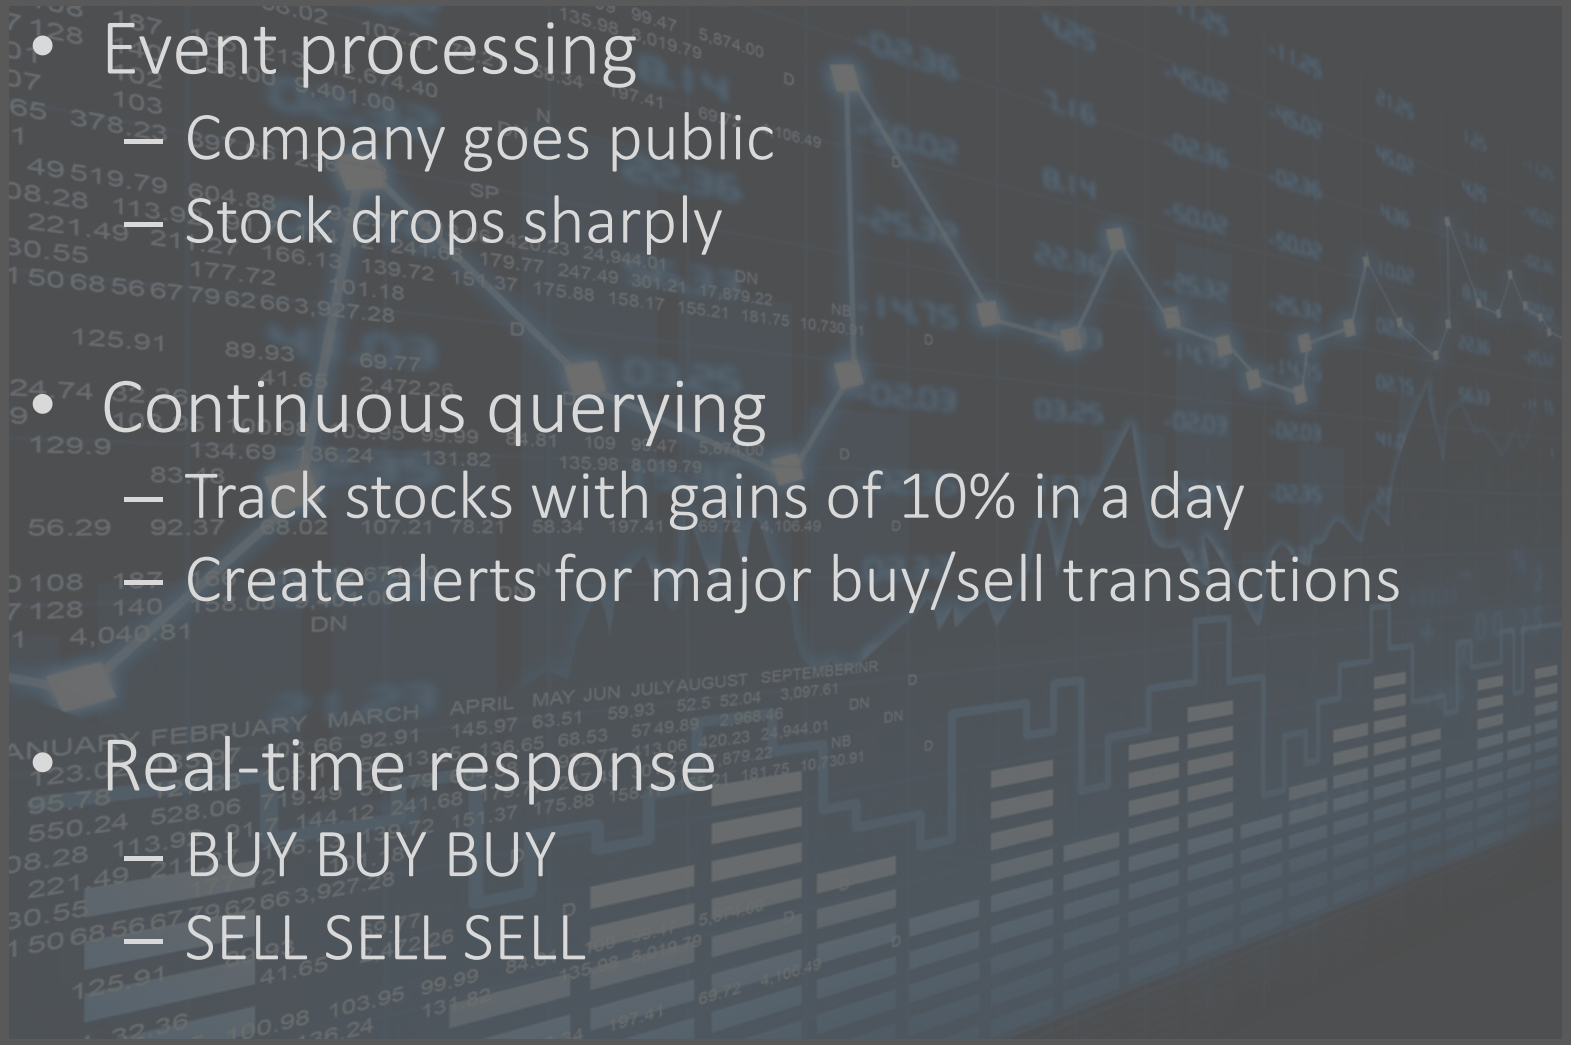

- Event processing
  - Company goes public
  - Stock drops sharply
- Continuous querying
  - Track stocks with gains of 10% in a day
  - Create alerts for major buy/sell transactions
- Real-time response
  - BUY BUY BUY
  - SELL SELL SELL

Applications: Astronomy

Applications: Astronomy

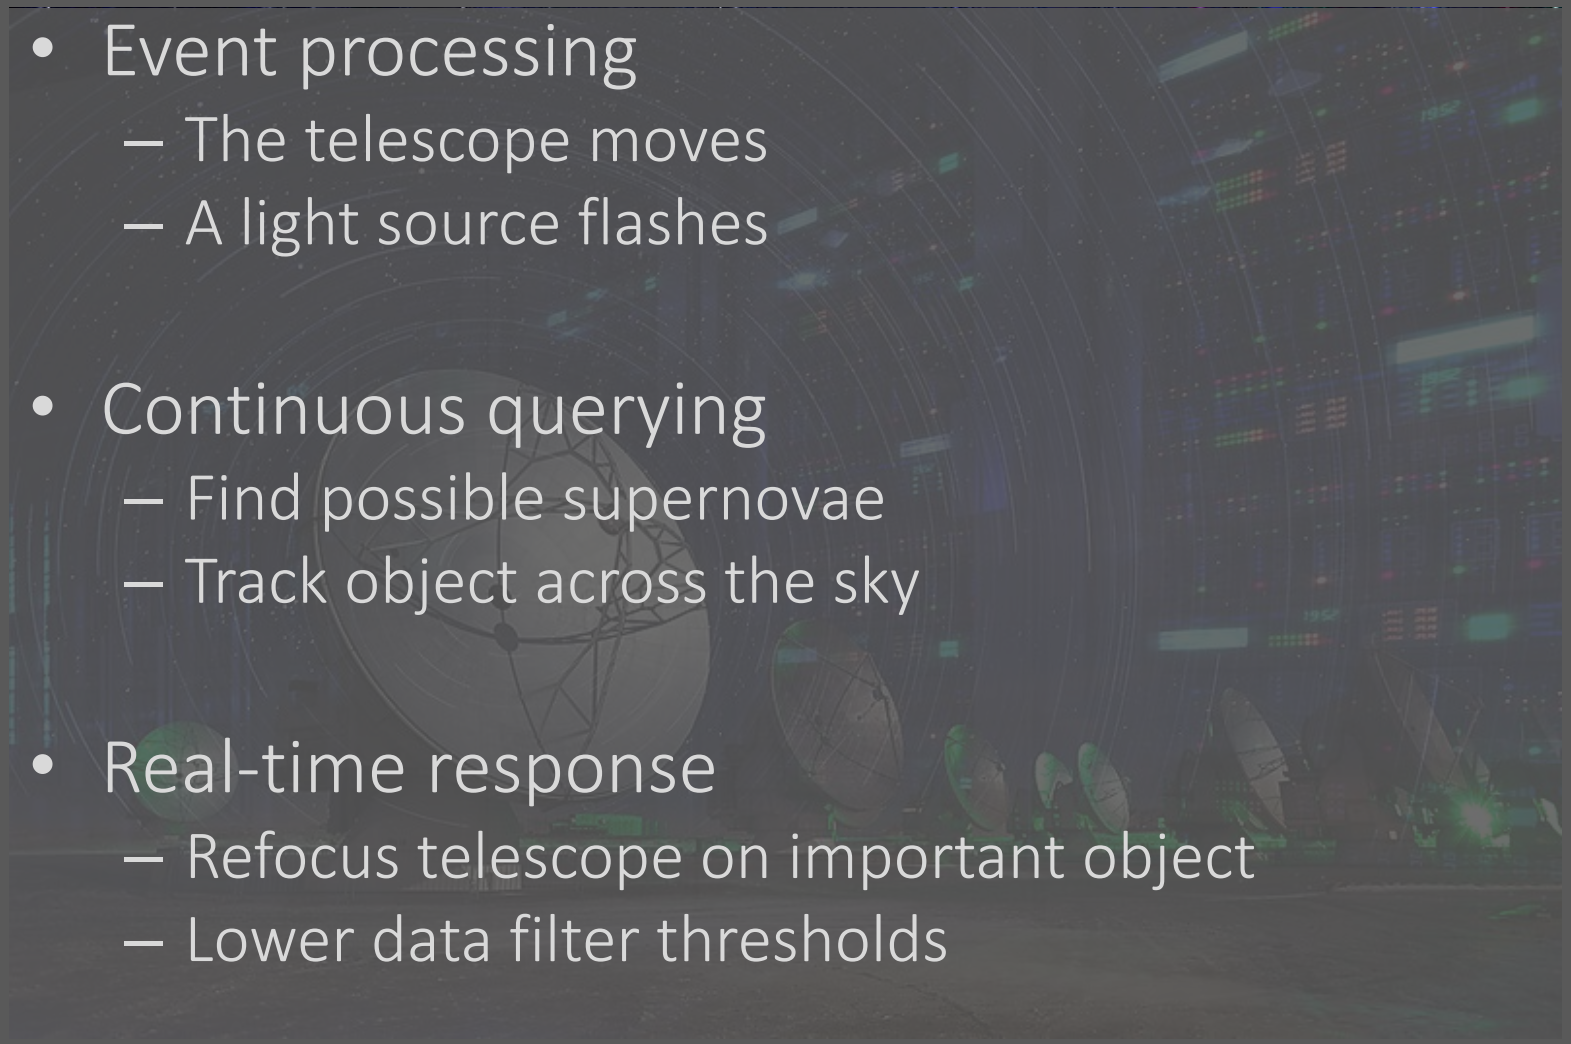

- Event processing
  - The telescope moves
  - A light source flashes
- Continuous querying
  - Find possible supernovae
  - Track object across the sky
- Real-time response
  - Refocus telescope on important object
  - Lower data filter thresholds

Applications: Astronomy

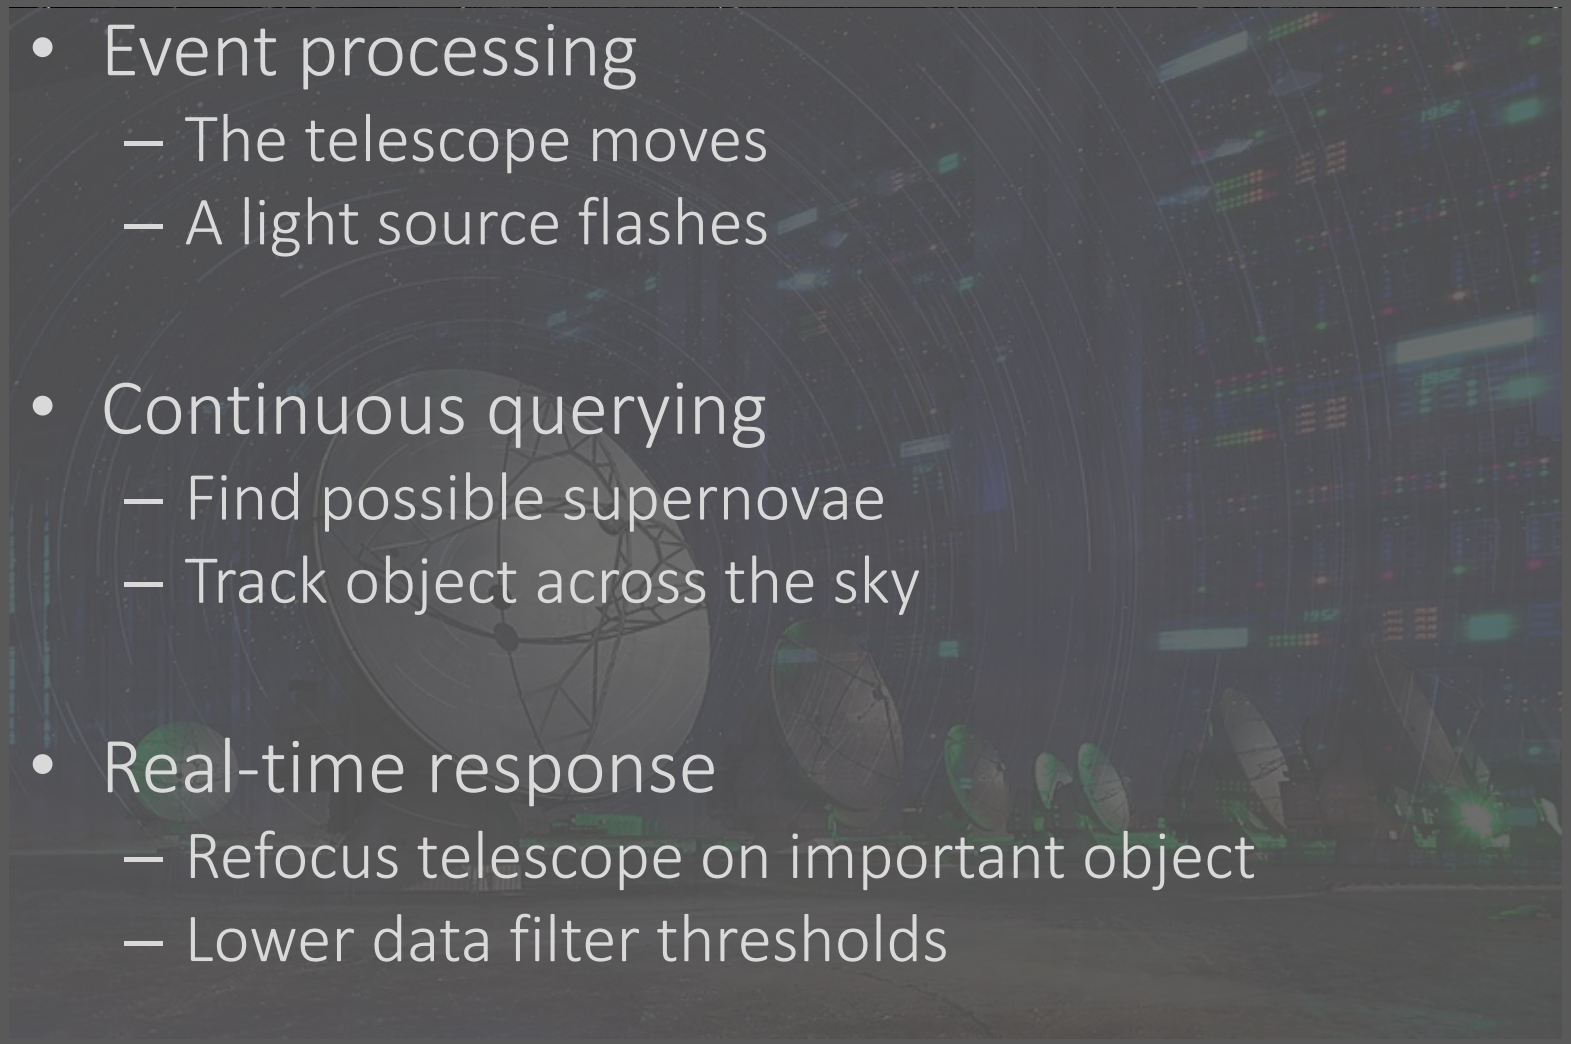

- Event processing
  - The telescope moves
  - A light source flashes
- Continuous querying
  - Find possible supernovae
  - Track object across the sky
- Real-time response
  - Refocus telescope on important object
  - Lower data filter thresholds

Streams: Internet of Things

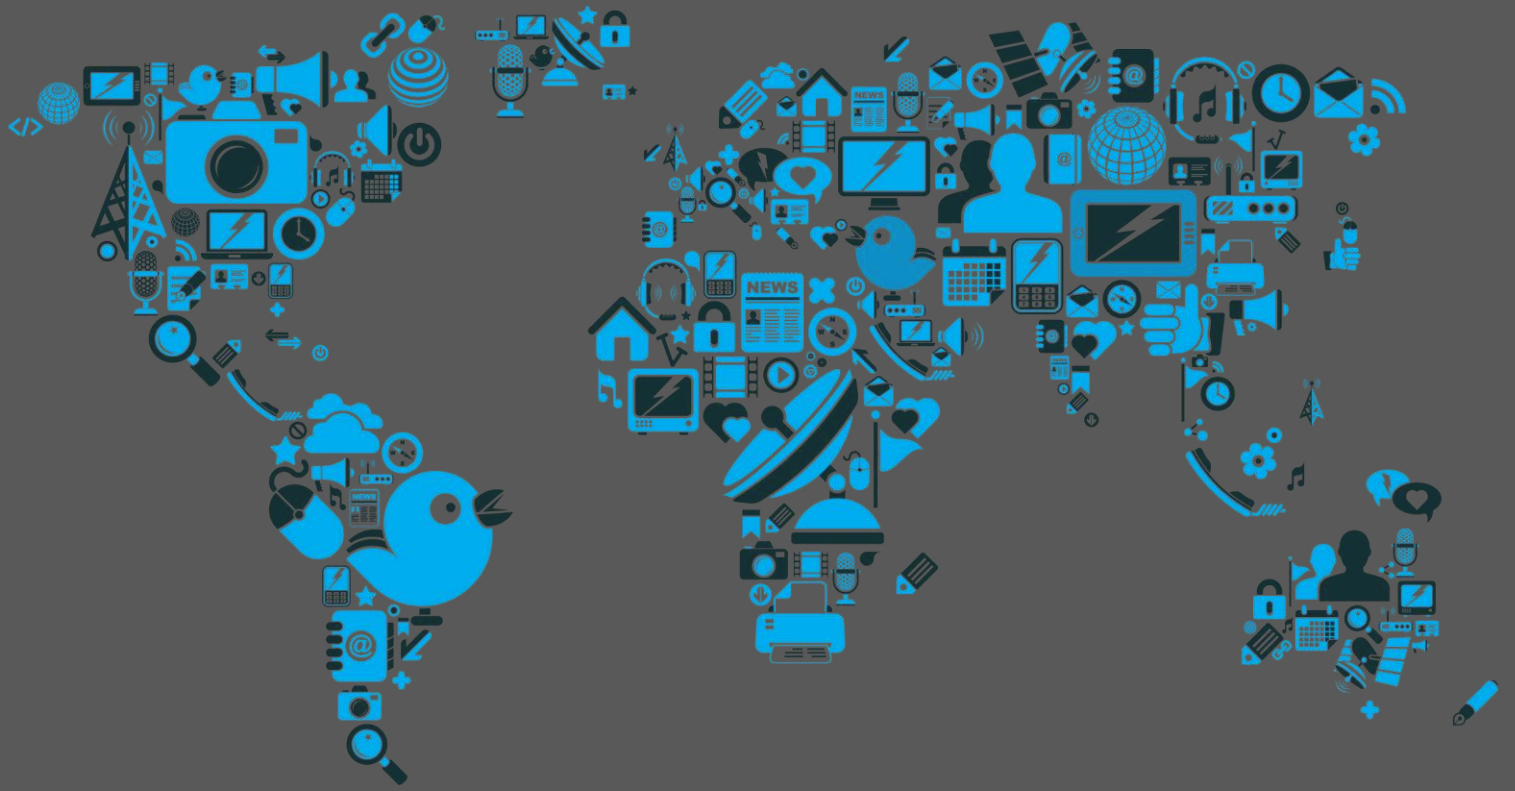

Streams: Internet of Things

- Event processing
  - A light turns on
  - It starts to rain
- Continuous querying
  - Tell me when temperature reaches 30°
  - Update position of vehicle
- Real-time response
  - Turn off air conditioning
  - Take another route

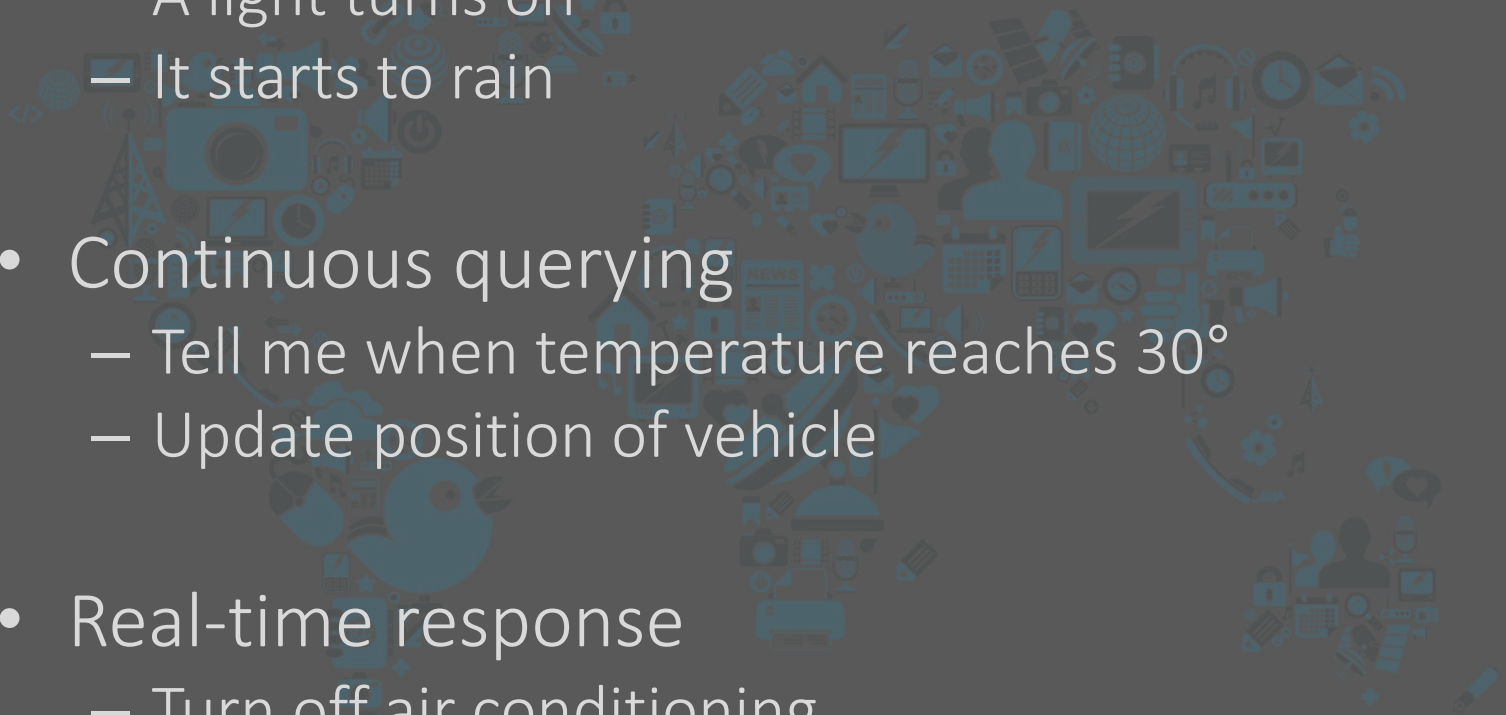

Distributed Streaming Platform

Available Frameworks

Application: Emergency Response

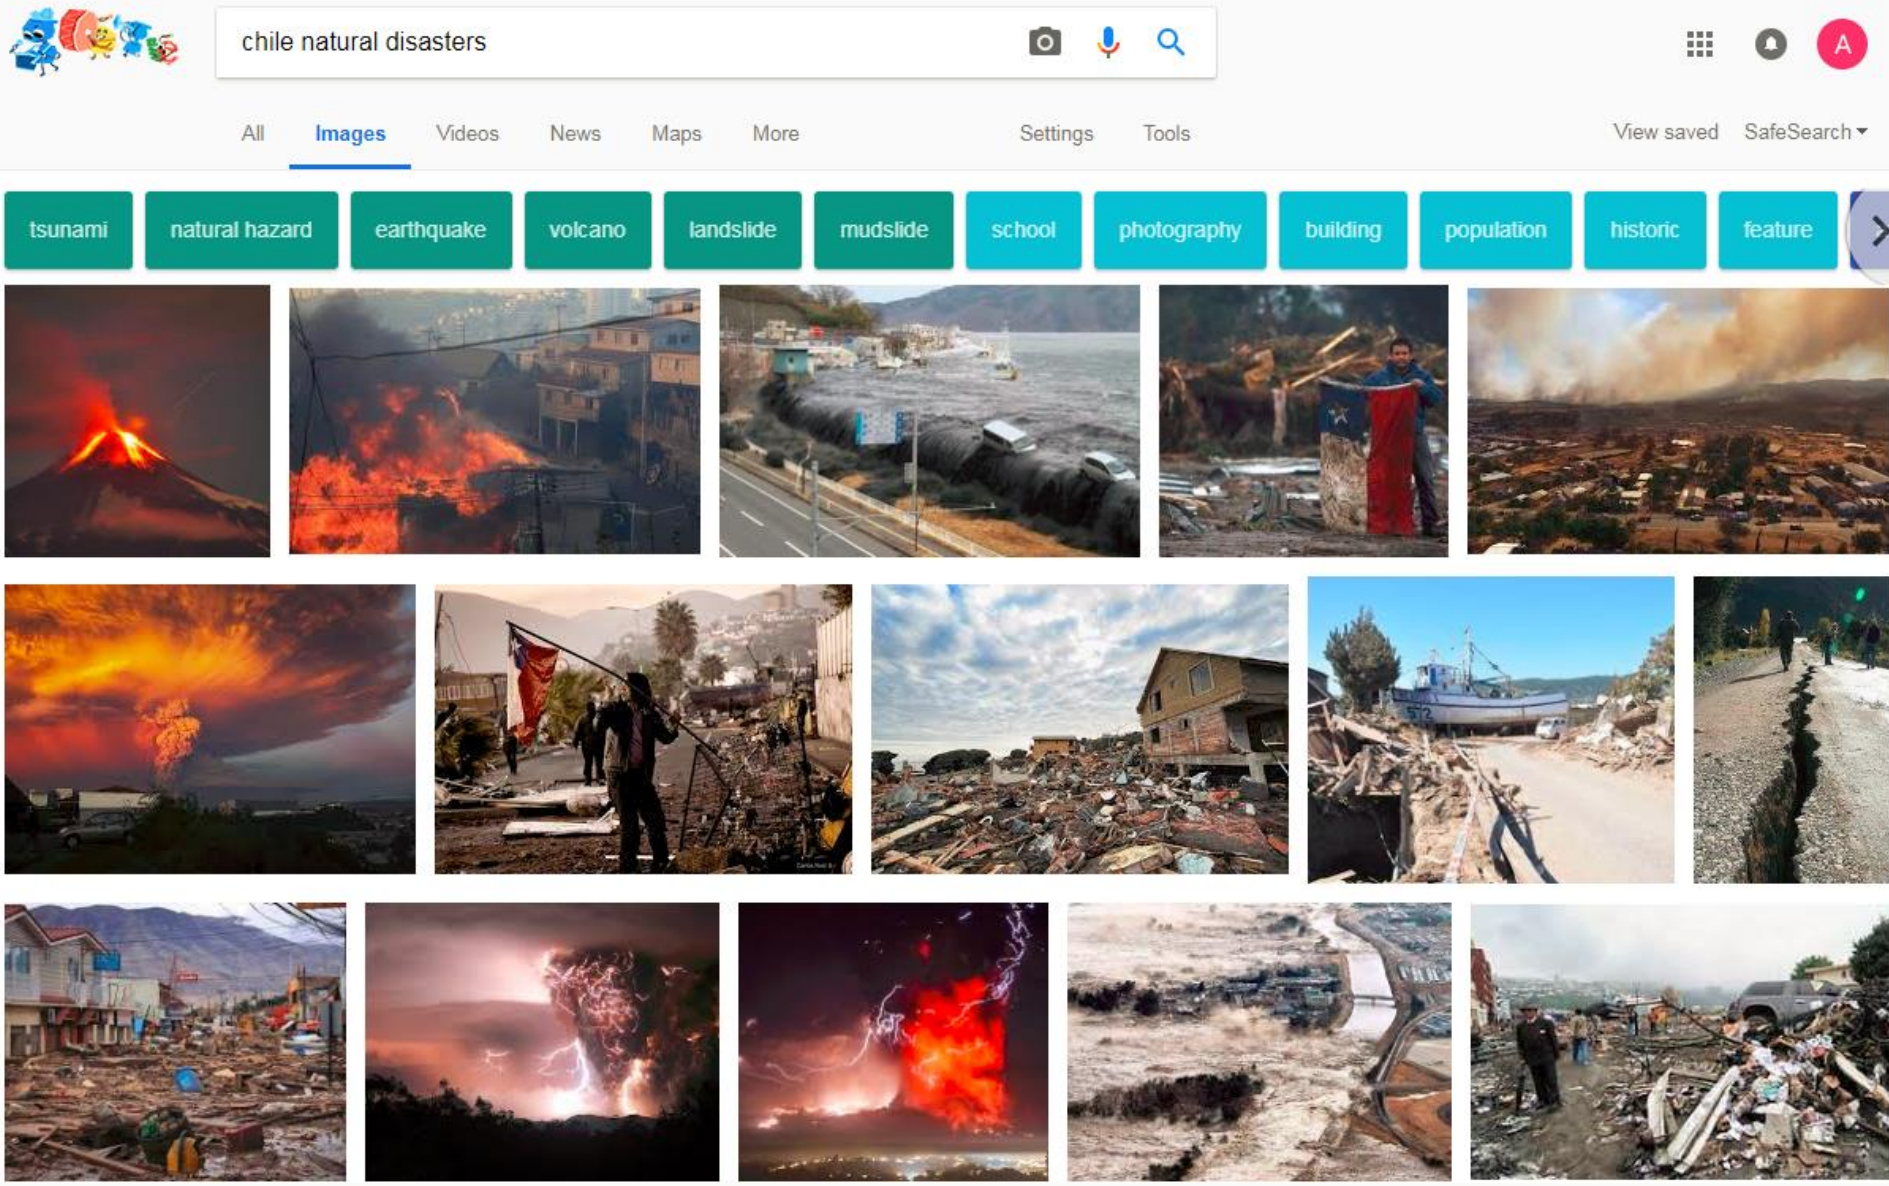

Real-Time Emergency Response

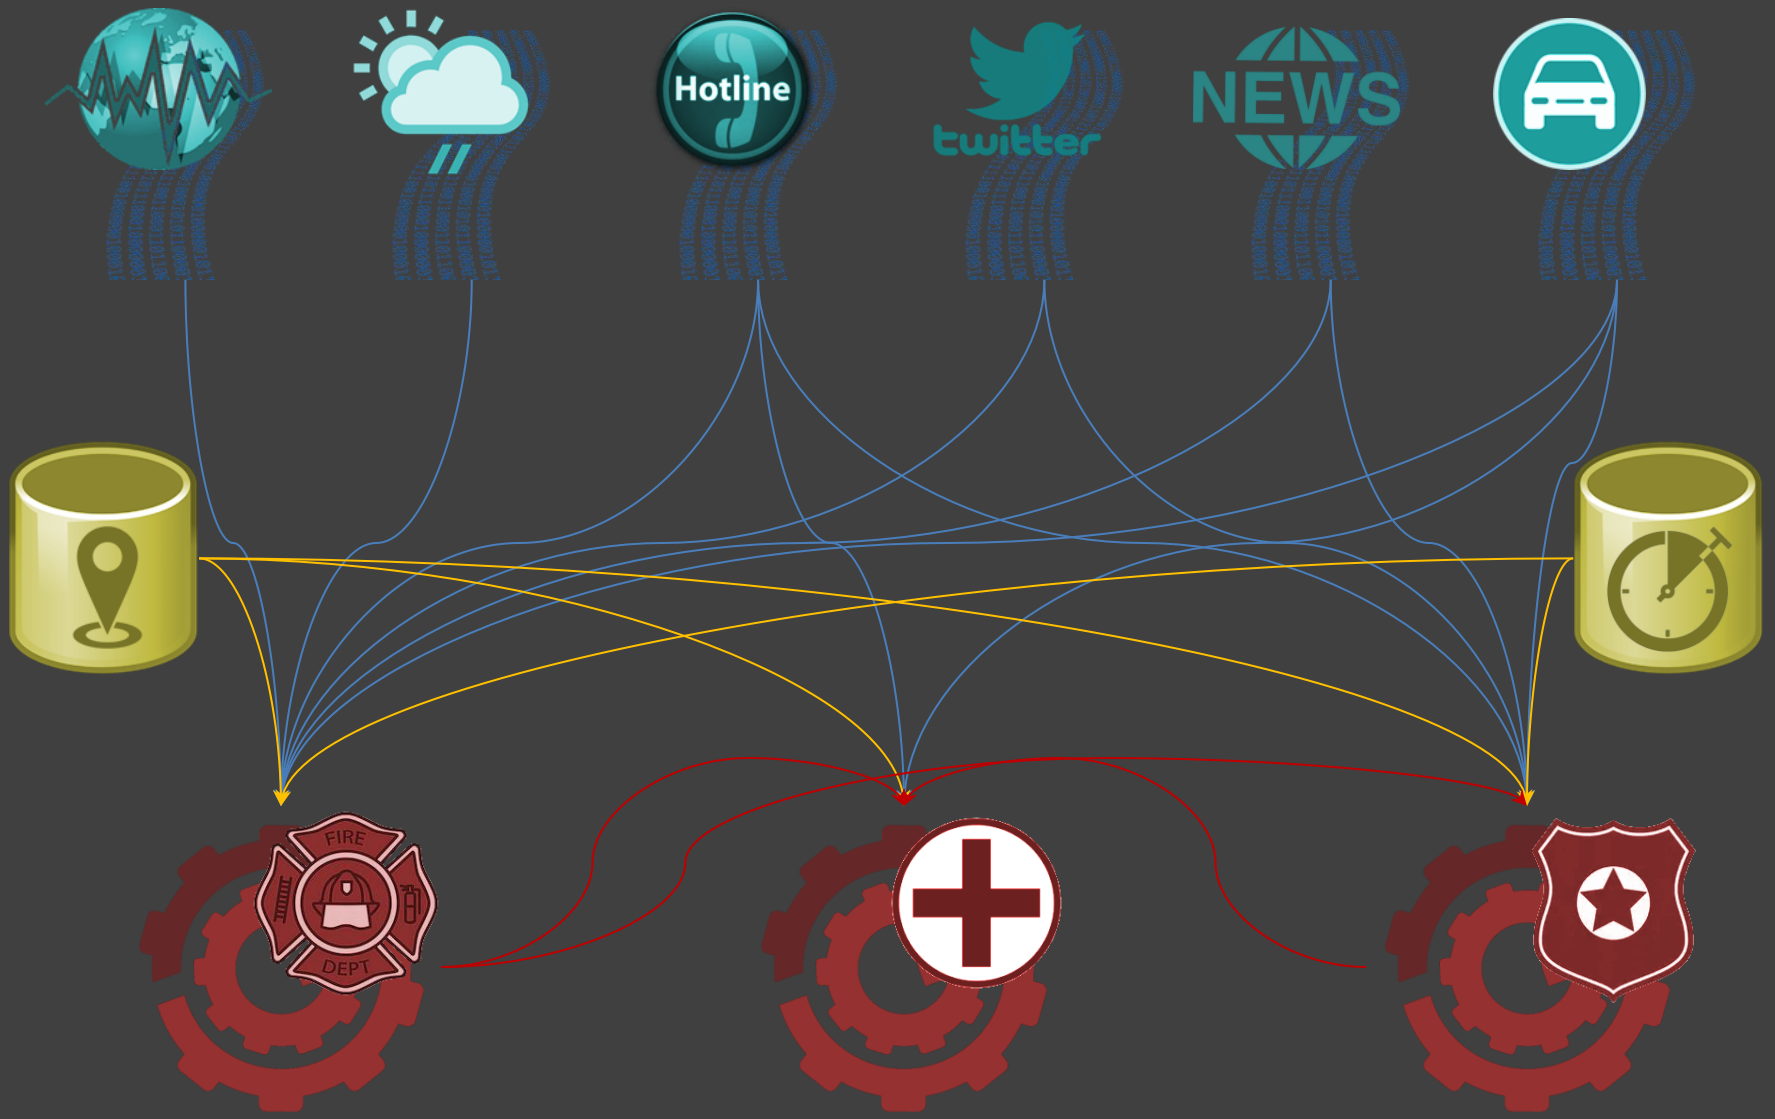

Real-Time Emergency Response

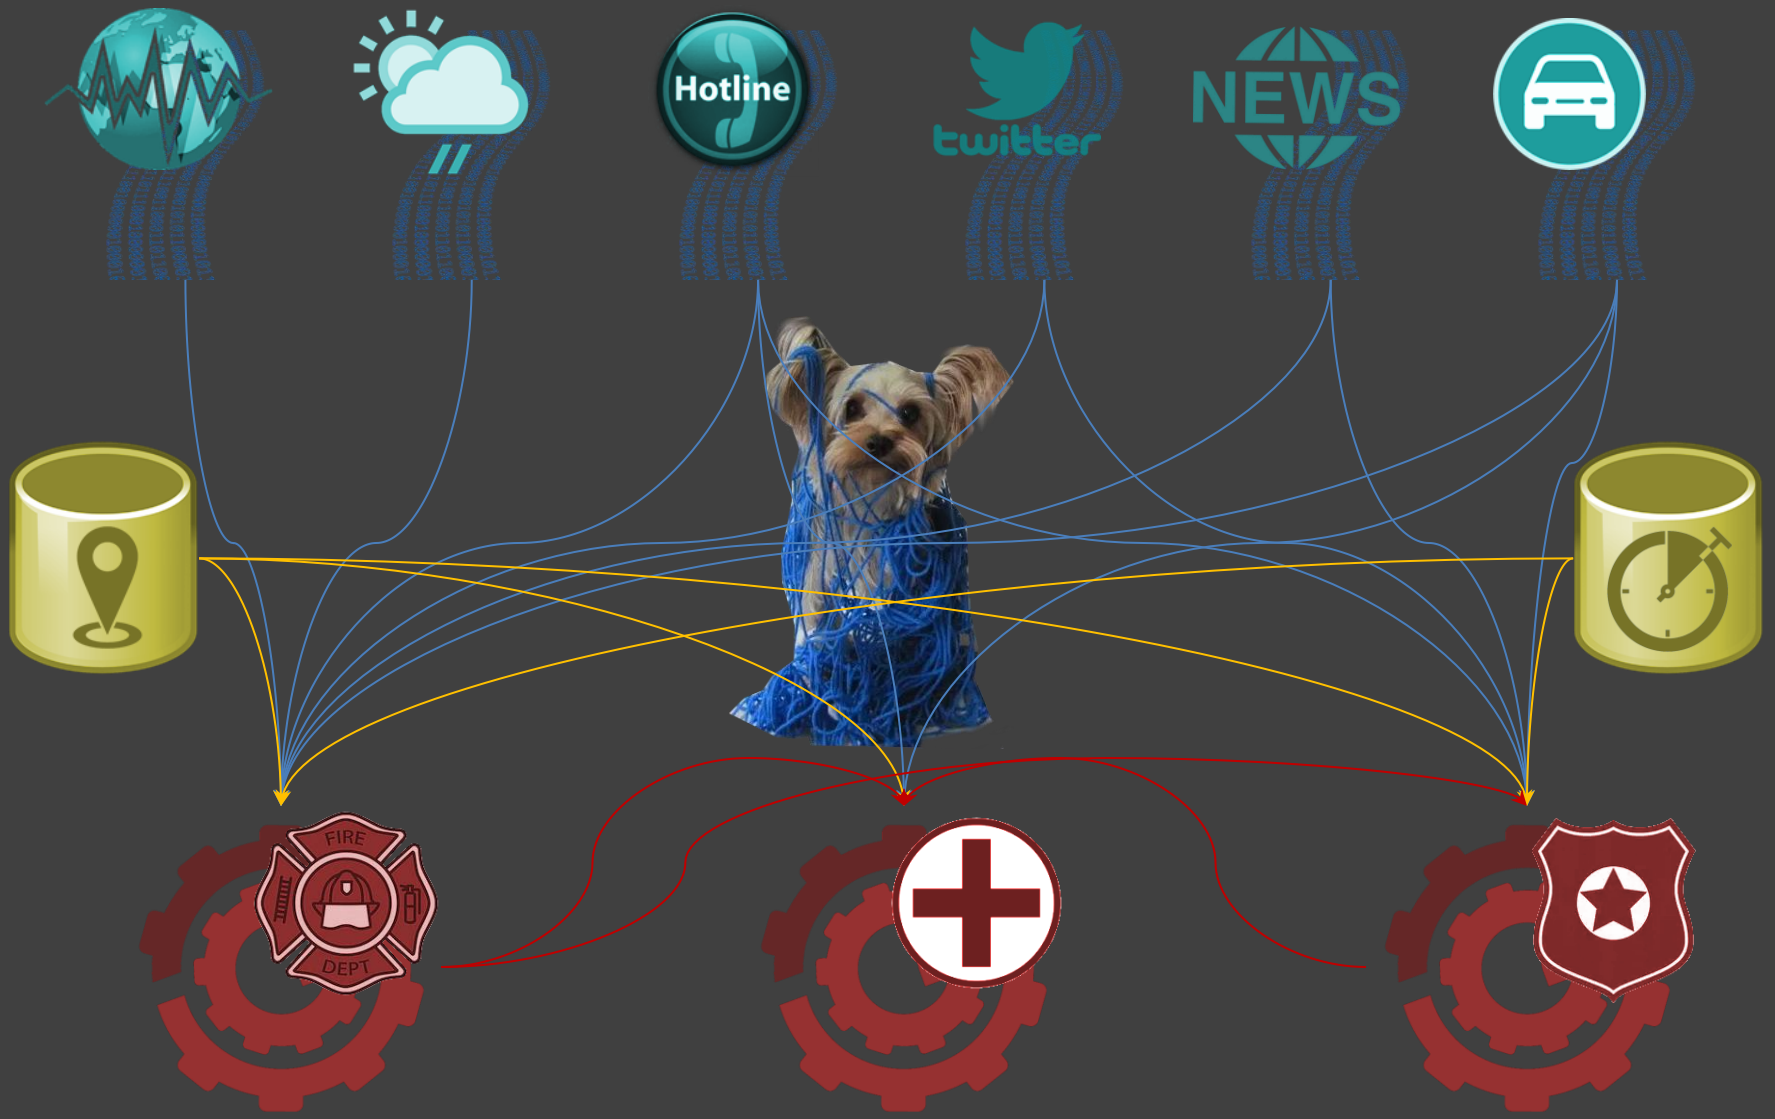

Real-Time Emergency Response

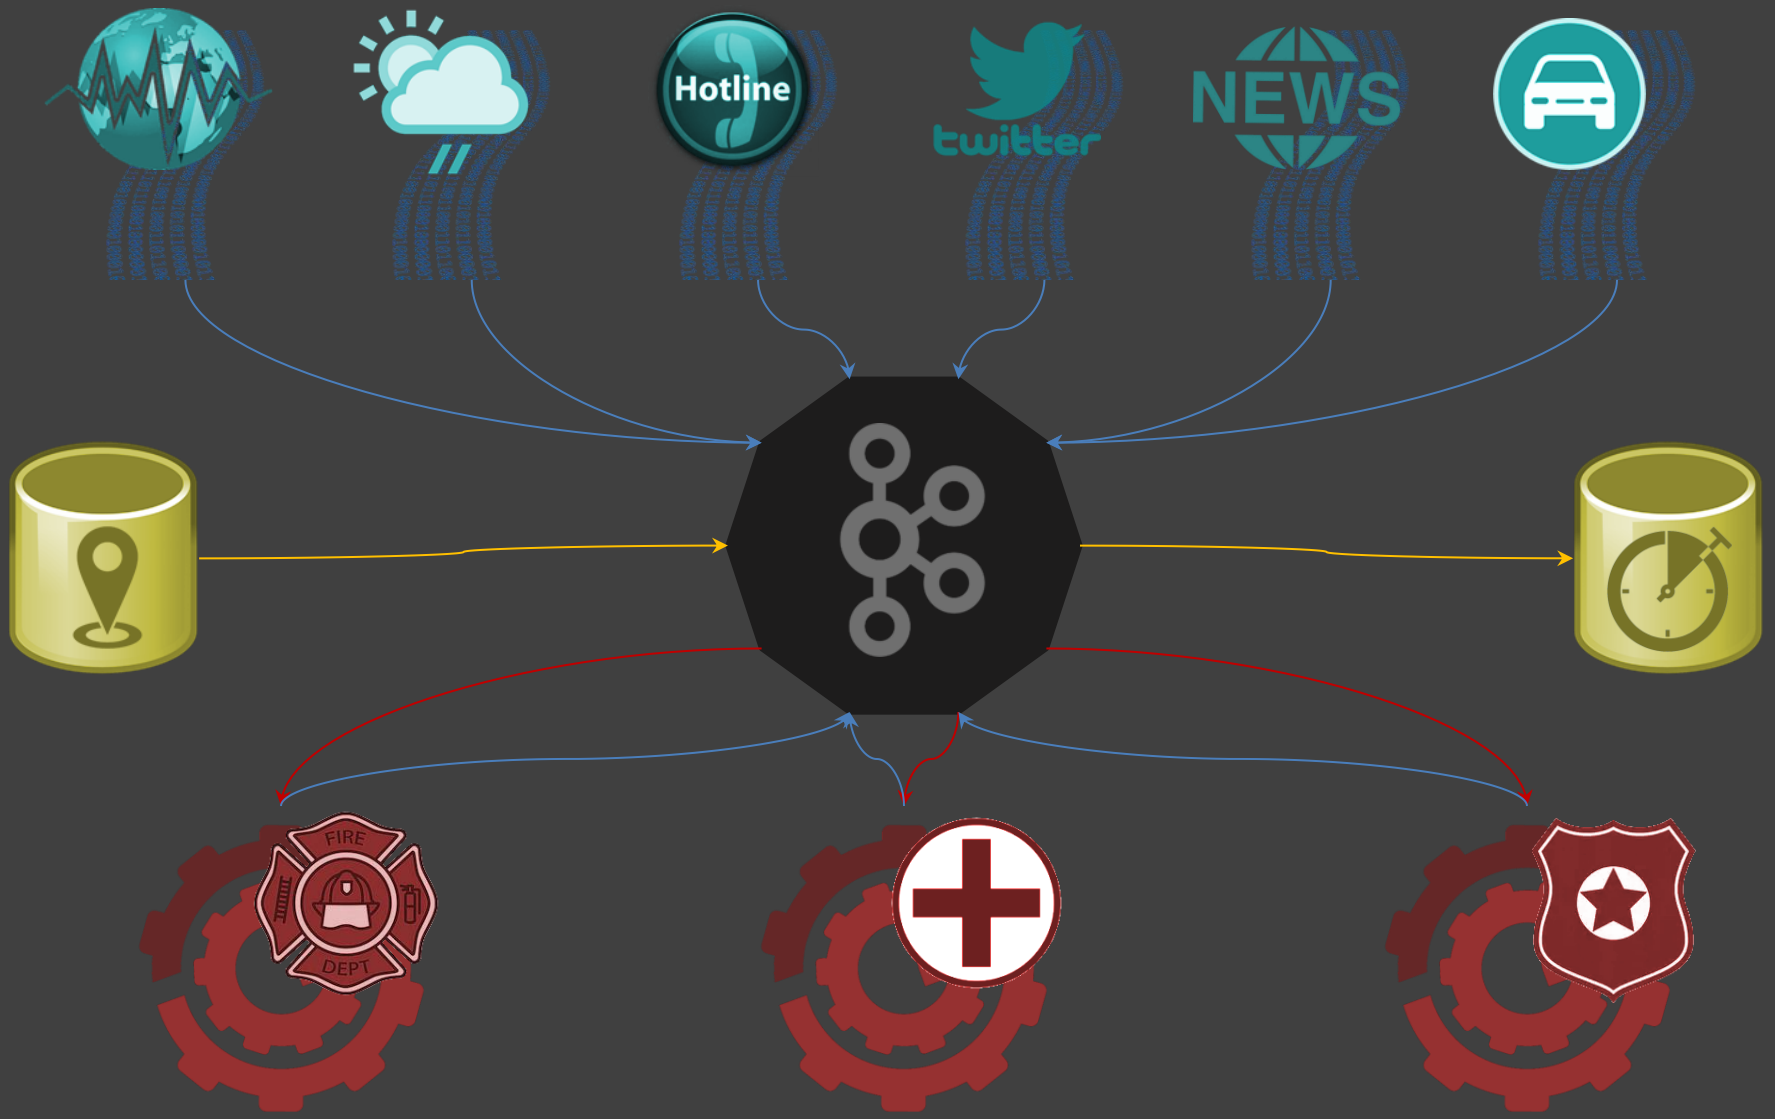

Apache Kafka

Apache Kafka vs. Franz Kafka

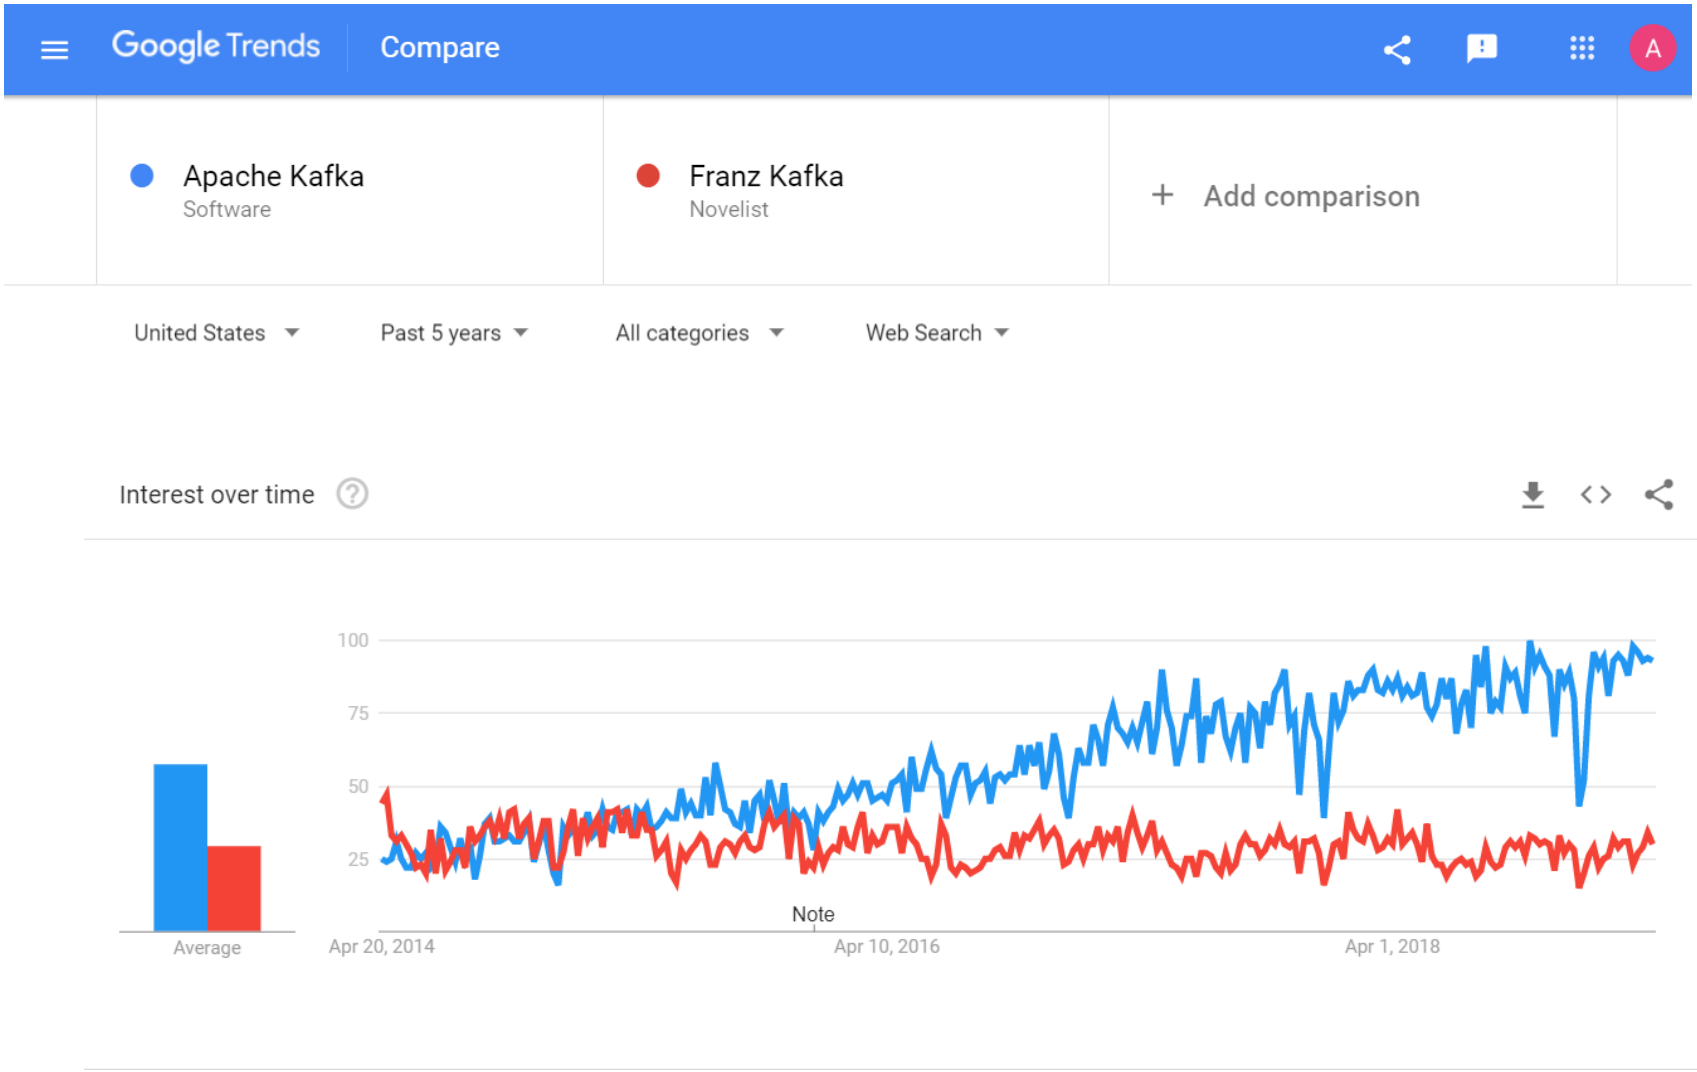

Apache Kafka

- Open Source
- Scala / Java
- Originated in LinkedIn

Kafka Overview

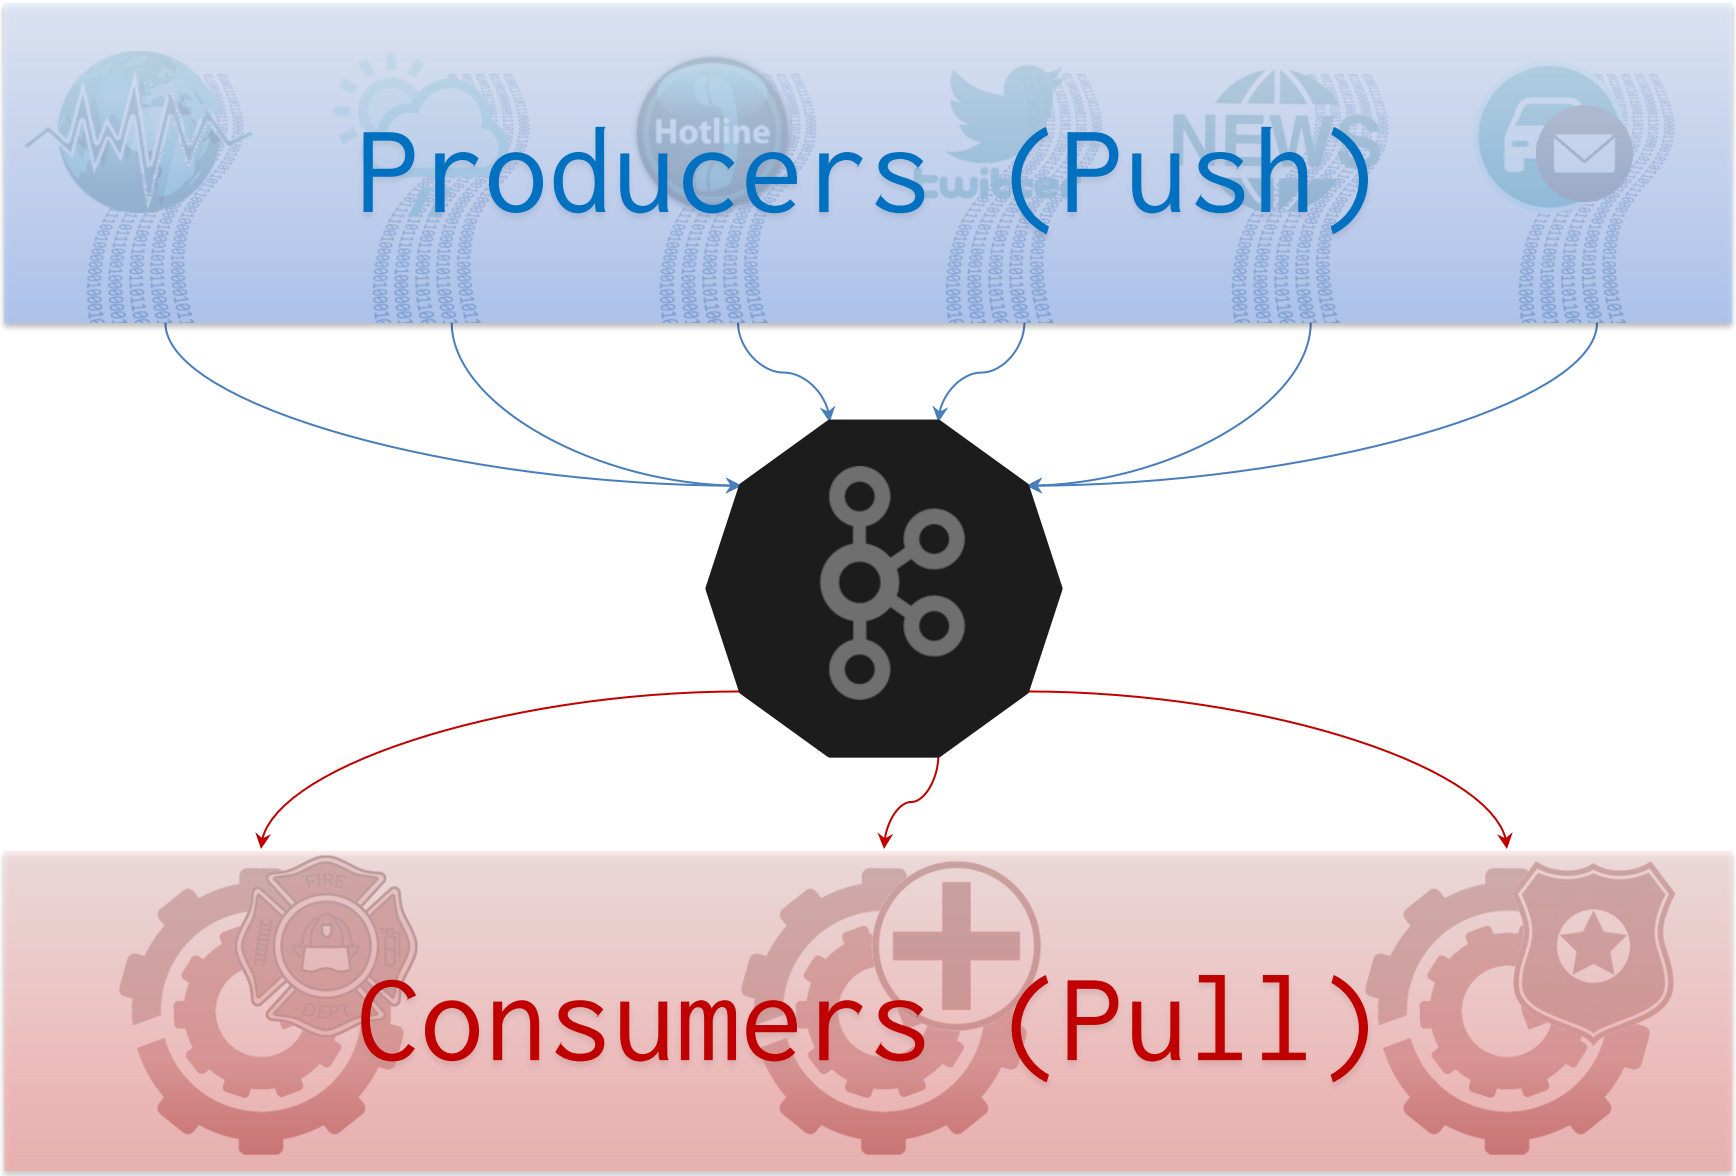

Kafka: Data Model

Kafka Record

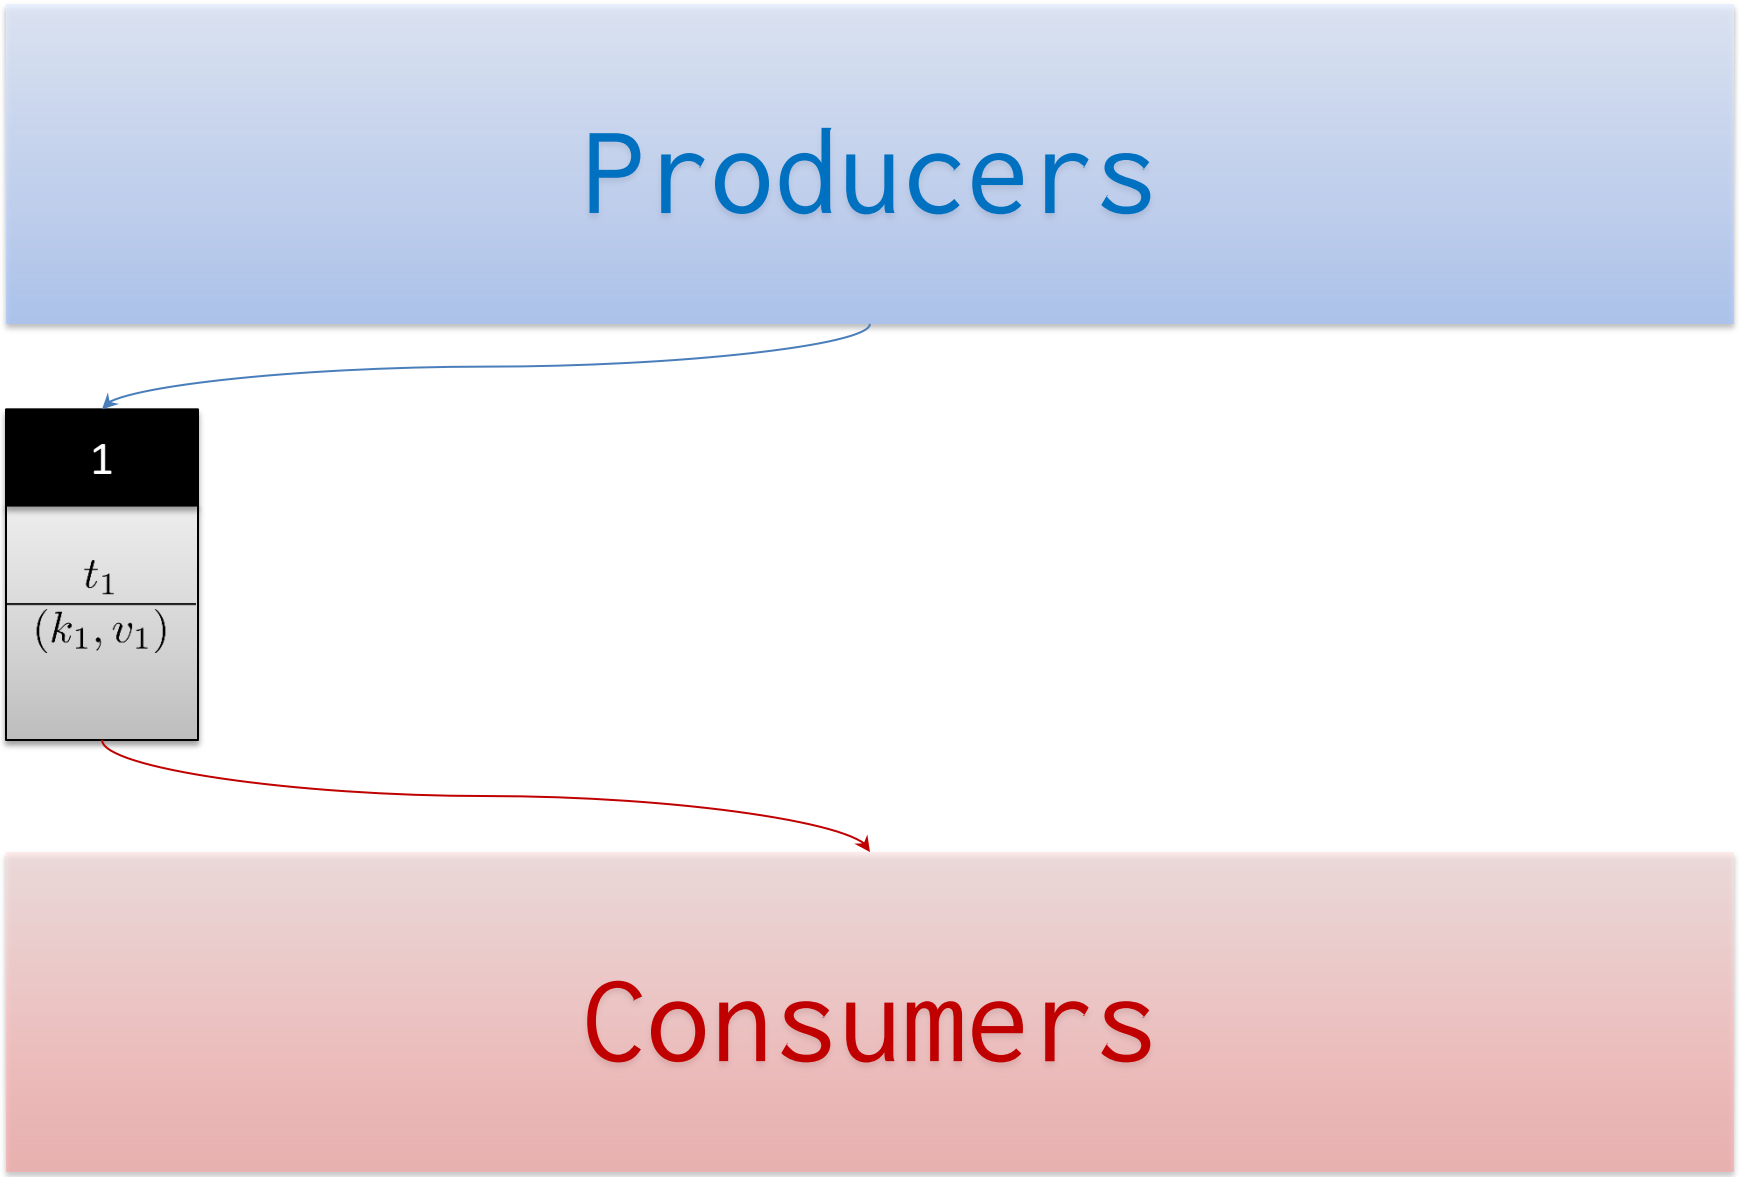

Kafka Record

- Records represent "events"
- Records are immutable
- Contain id (offset), timestamp, key and value
  - Timestamp assigned by application or Kafka

Kafka Ledger

Producers

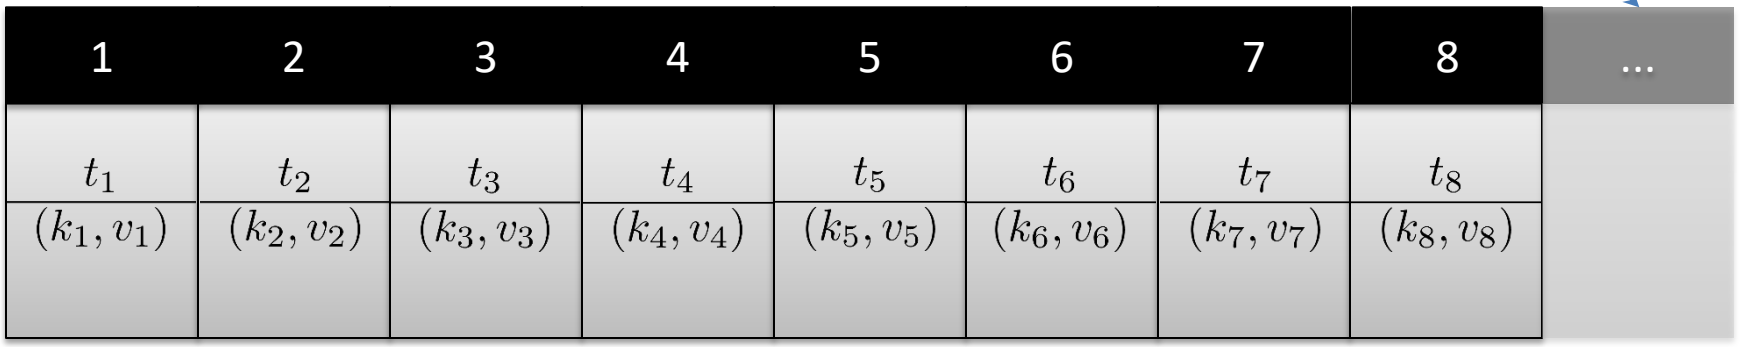

Consumers

Kafka Ledger

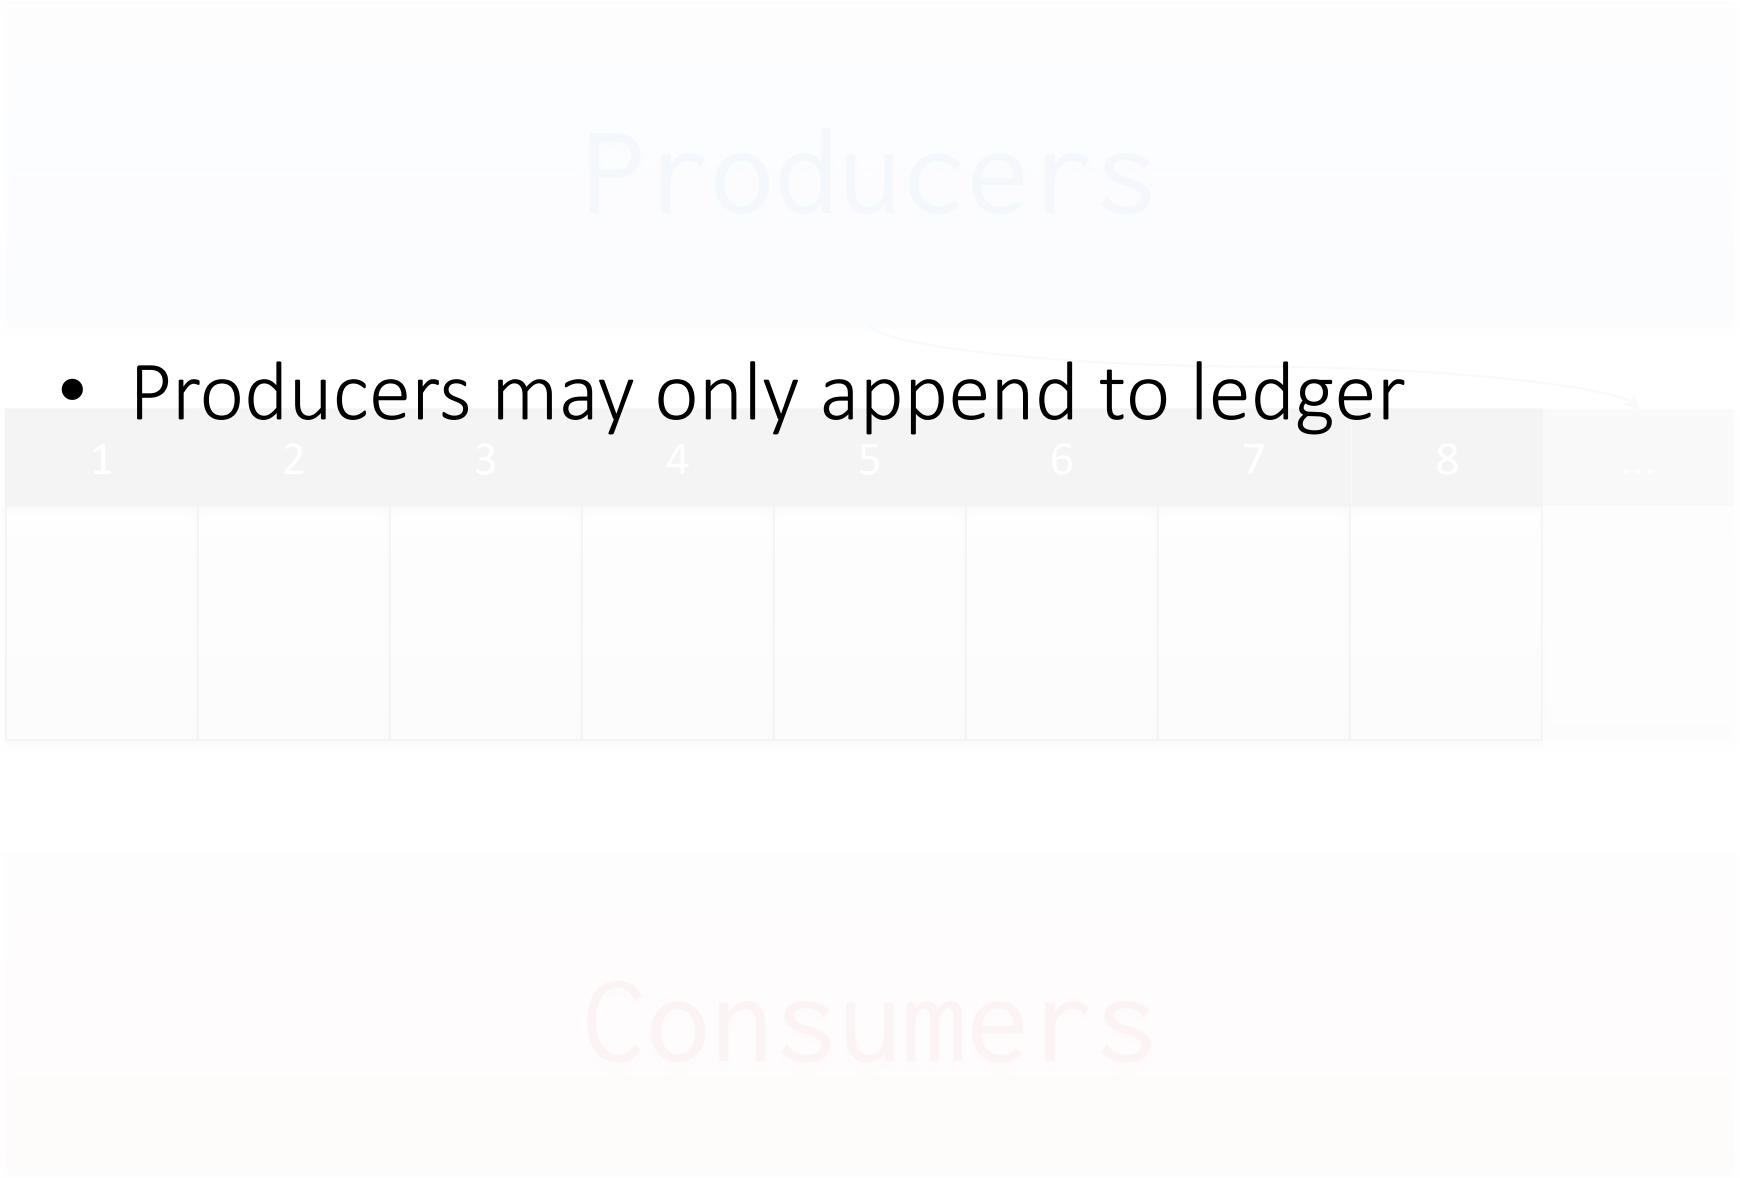

- Producers may only append to ledger

Kafka Log

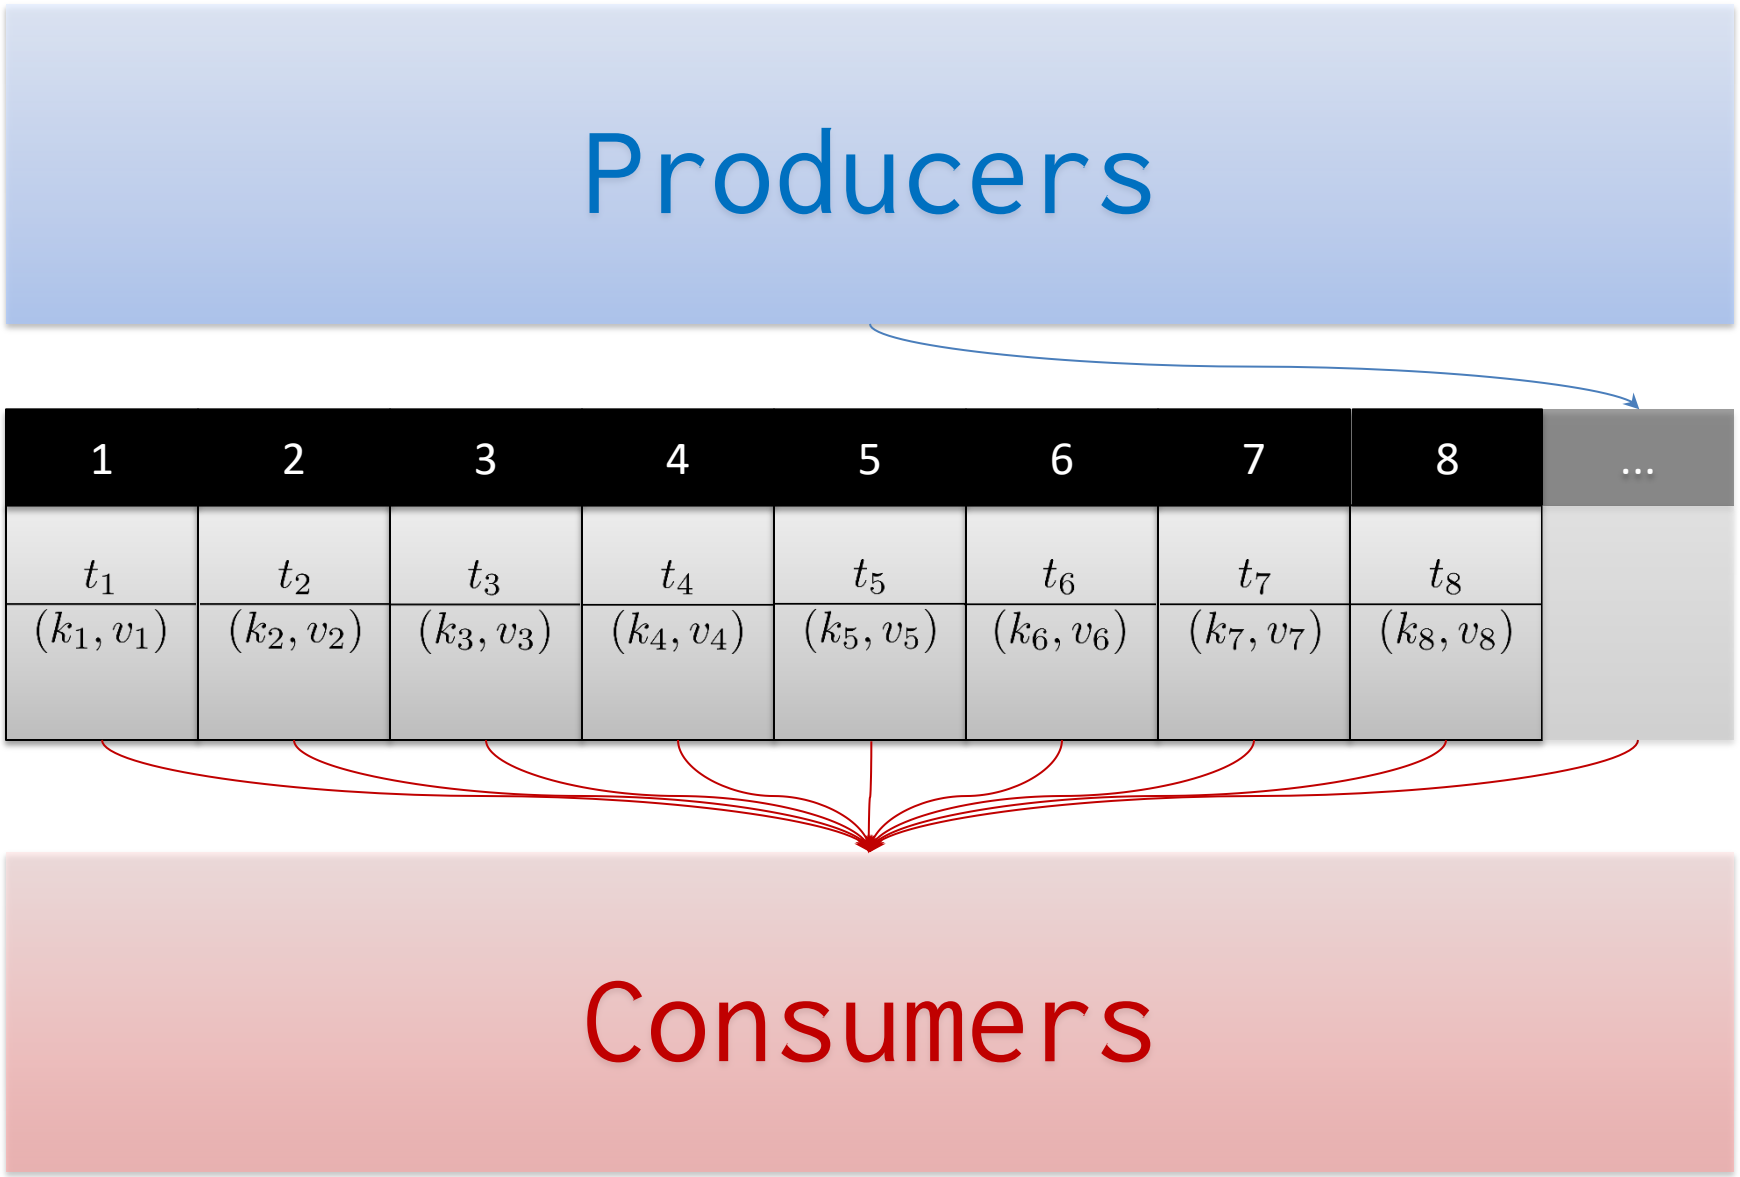
Producers
1
2
3
4
5
6
7
8
...
t_1 (k_1, v_1)
t_2 (k_2, v_2)
t_3 (k_3, v_3)
t_4 (k_4, v_4)
t_5 (k_5, v_5)
t_6 (k_6, v_6)
t_7 (k_7, v_7)
t_8 (k_8, v_8)
Consumers

Kafka Log

- Producers may only append to log
- Consumers can read from anywhere*

  * kind of

Kafka: Topics

Kafka Topics

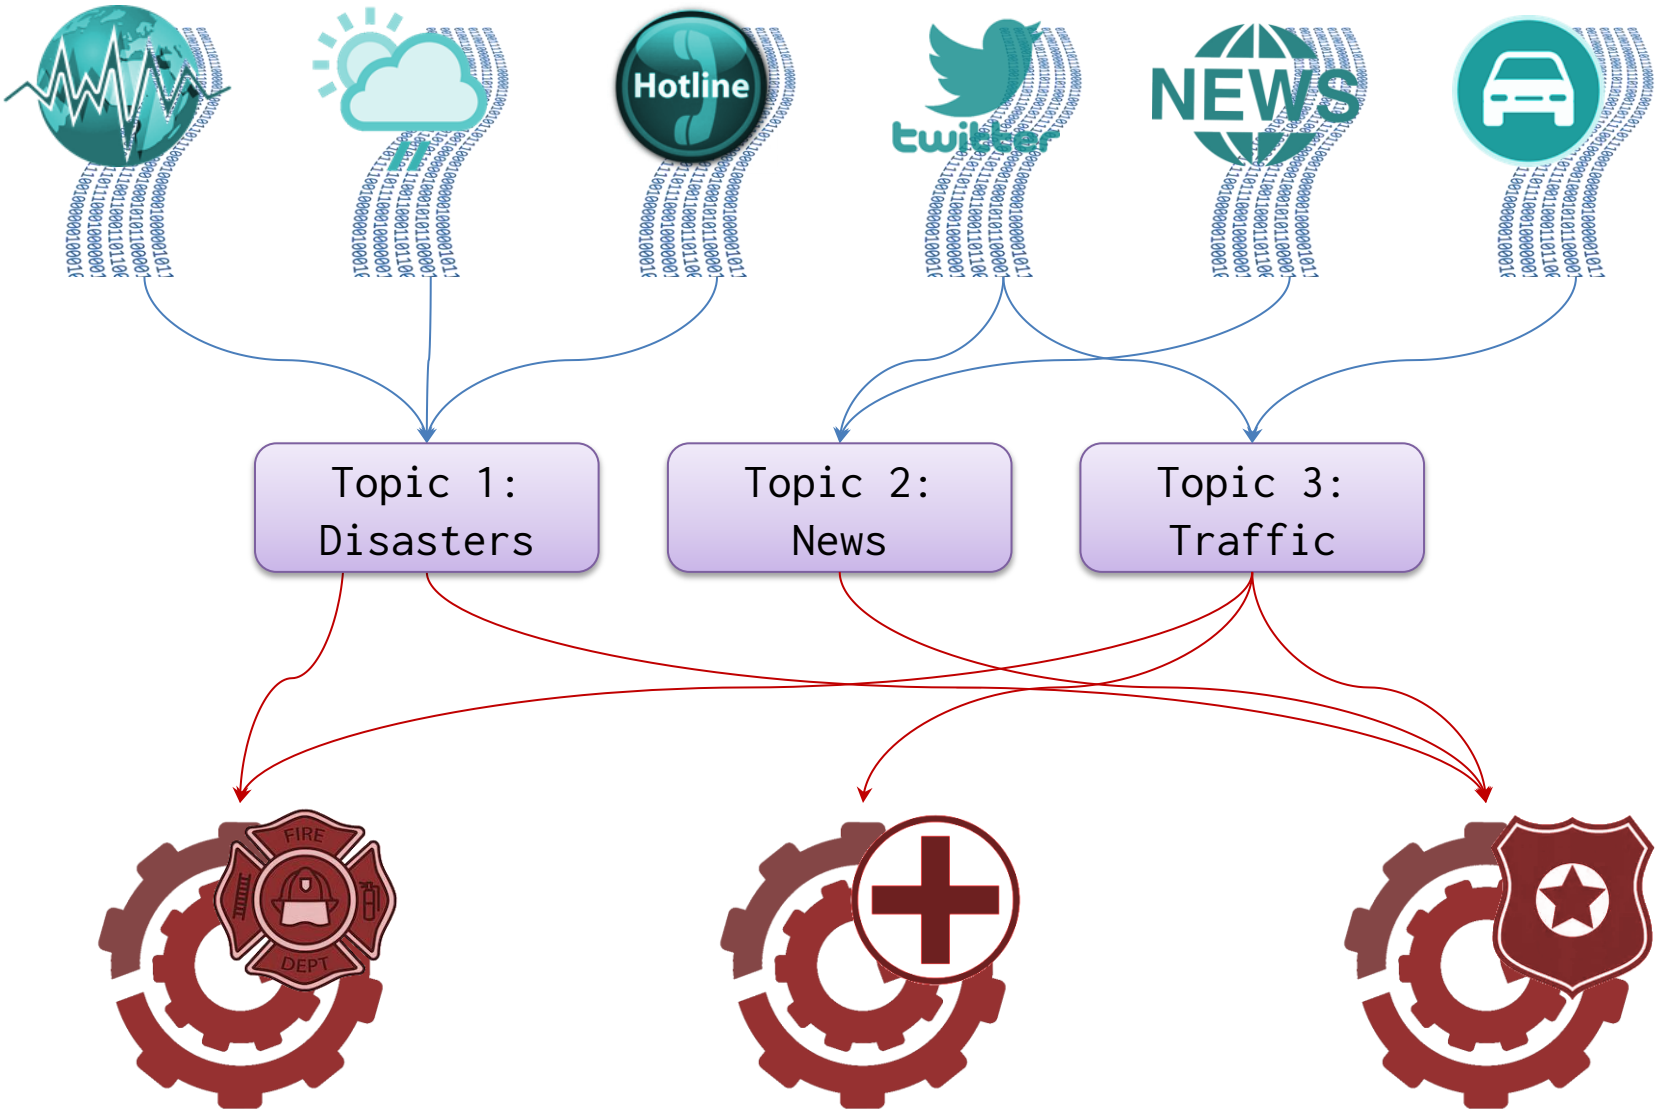

Kafka Topics

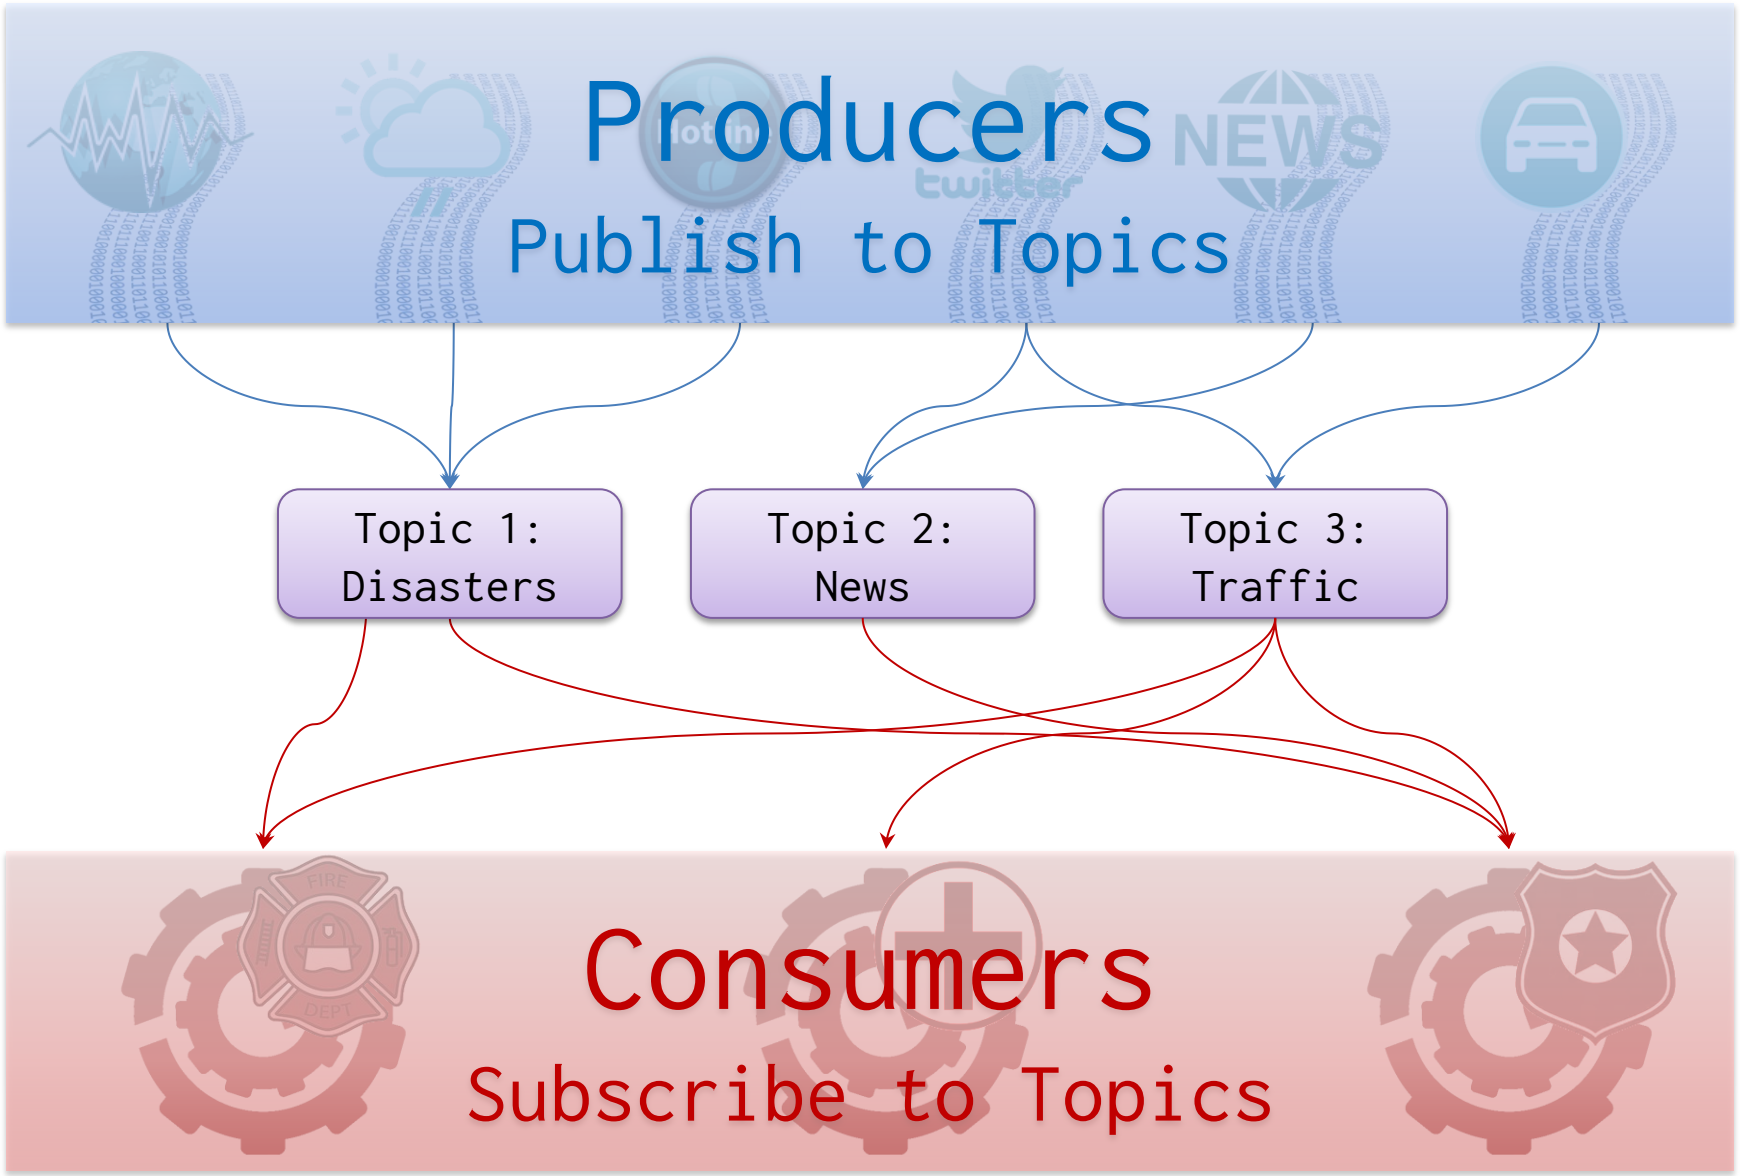

Topic

- Topics are persistent (on disk)
  - Configurable retention policy
    - Keep everything
    - Delete once consumed
    - Keep for a period of time
    - Use fixed amount of space

Topic: Default Partitioning by Key

Partition 1

1	2	3	4	5	6	...
t_1^1	t_2^1	t_3^1	t_4^1	t_5^1	t_6^1	
(k_1^1, v_1^1)	(k_2^1, v_2^1)	(k_3^1, v_3^1)	(k_4^1, v_4^1)	(k_5^1, v_5^1)	(k_6^1, v_6^1)	

Partition 2

1	2	3	4	5	6	7	8	...
t_1^2	t_2^2	t_3^2	t_4^2	t_5^2	t_6^2	t_7^2	t_8^2	
(k_1^2, v_1^2)	(k_2^2, v_2^2)	(k_3^2, v_3^2)	(k_4^2, v_4^2)	(k_5^2, v_5^2)	(k_6^2, v_6^2)	(k_7^2, v_7^2)	(k_8^2, v_8^2)	

Partition 3

1	2	3	4	...
t_1^3	t_2^3	t_3^3	t_4^3	
(k_1^3, v_1^3)	(k_2^3, v_2^3)	(k_3^3, v_3^3)	(k_4^3, v_4^3)	

Topic: Default Partitioning by Key

- Ordering (offset) guaranteed per partition
  - Not across partitions!
  - For ordering across partitions, use timestamp

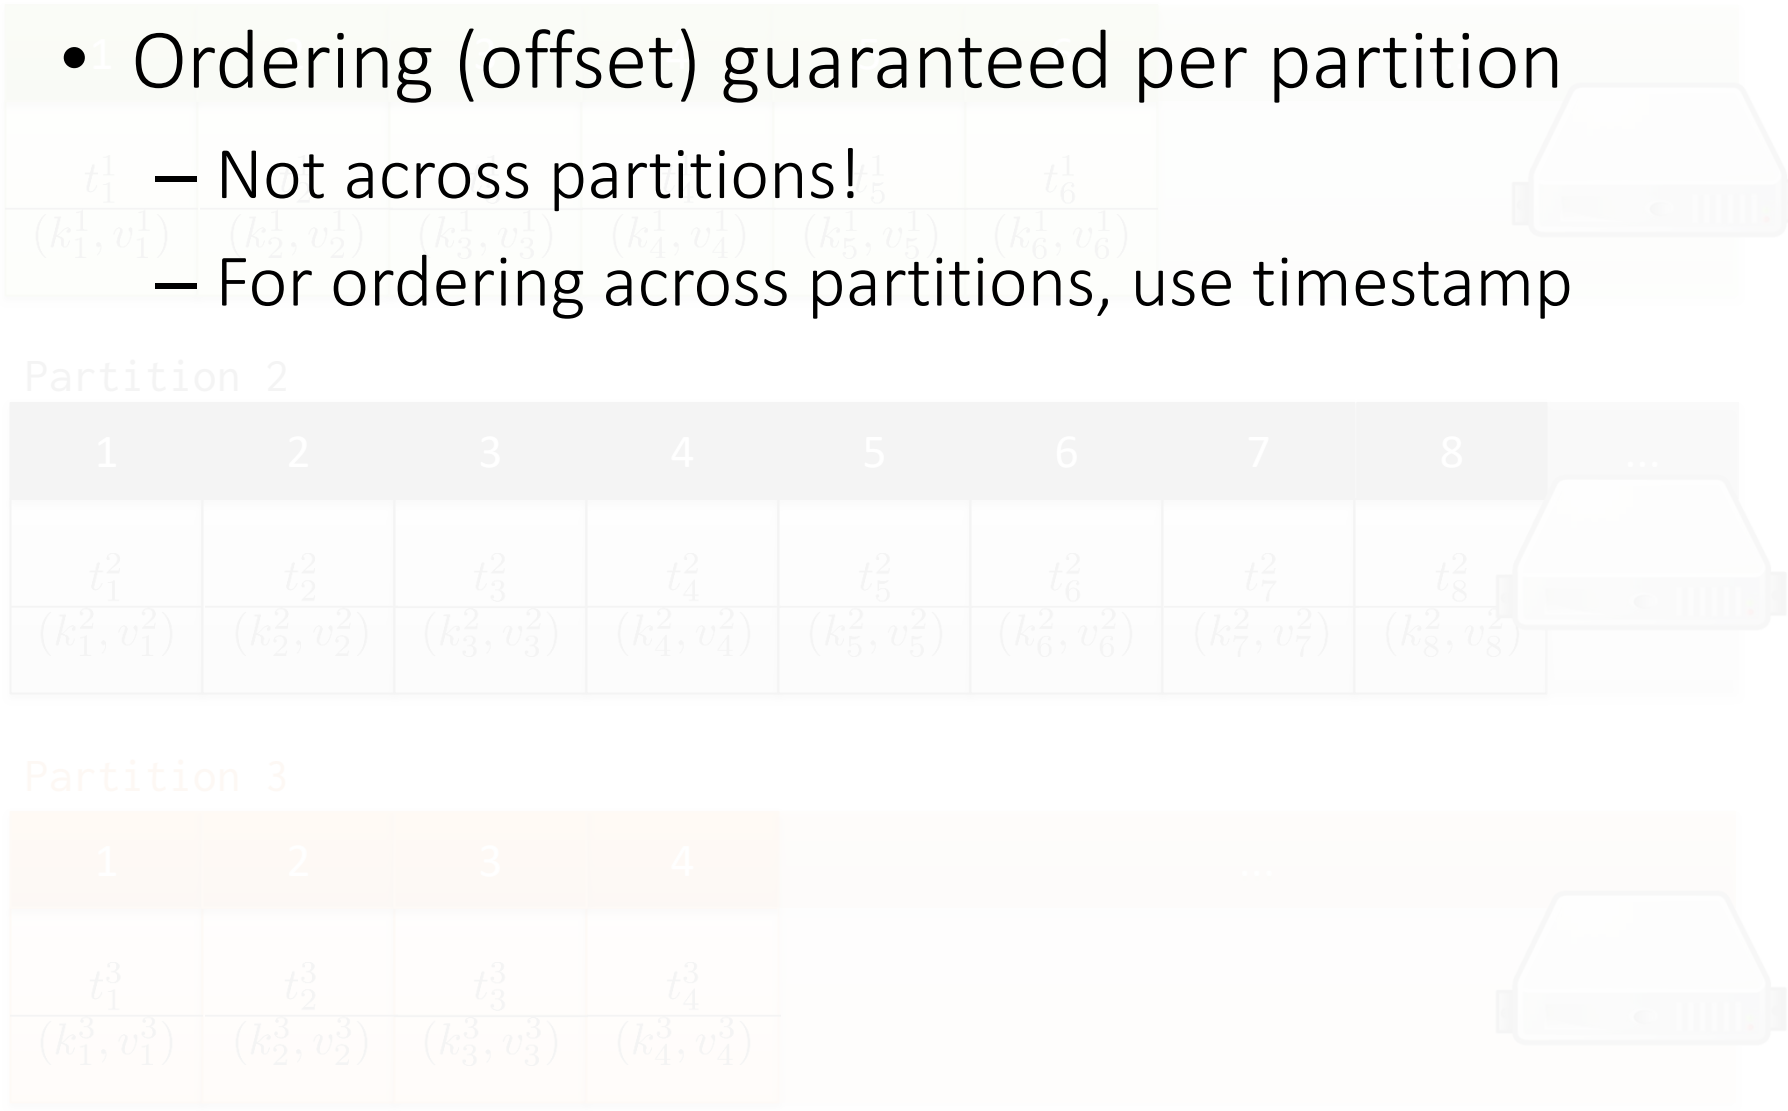

Replication

- Topics can be replicated
  - Choose factor per topic
  - Automatic load balancing

Problem?

Order?

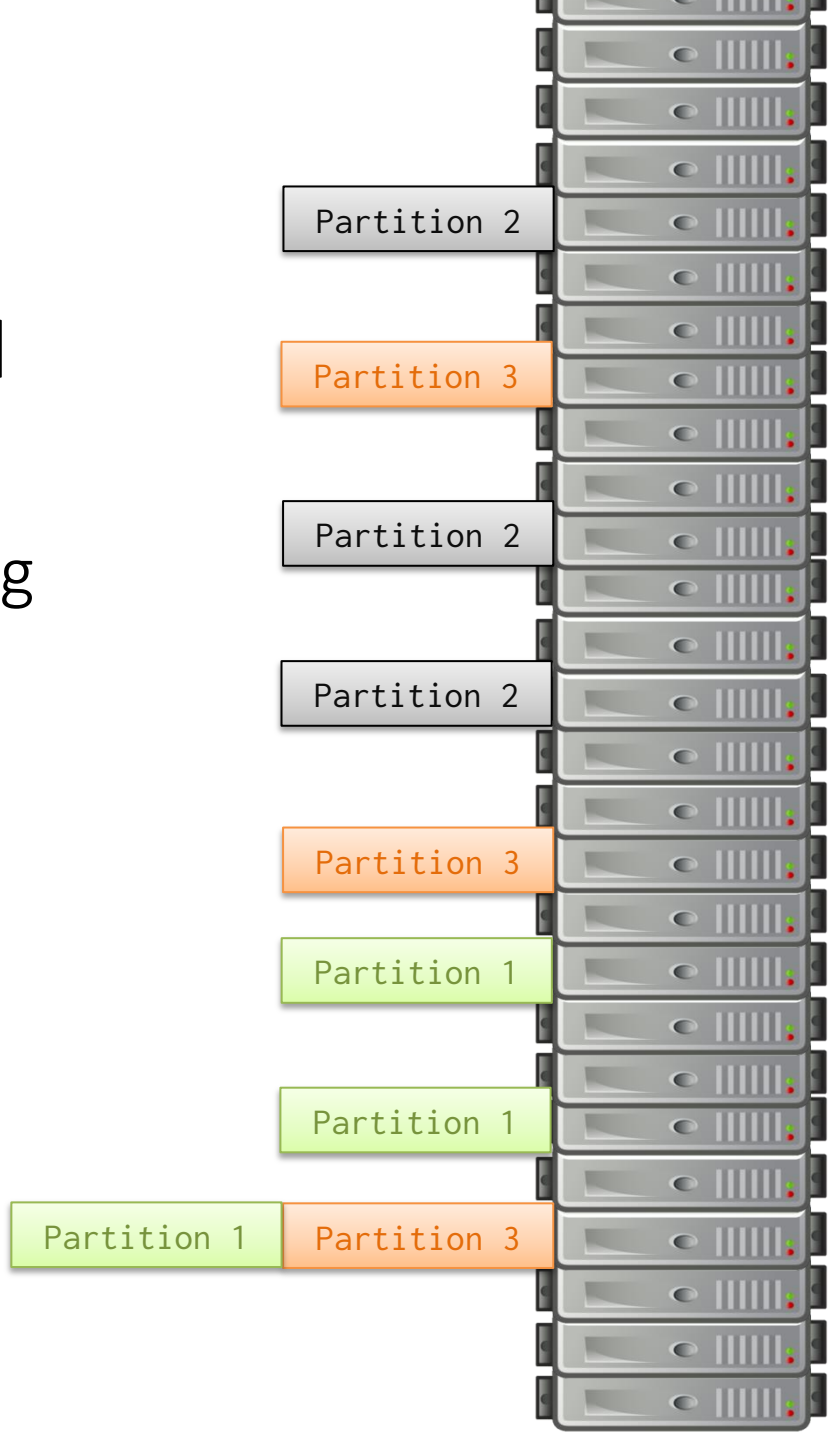

Leader

- Topics can be replicated
  - Choose factor per topic
  - Automatic load balancing
- One machine is the **leader**
  - The others are followers
  - Leader automatically elected
  - Ensures order per partition
  - Reads/writes to leader

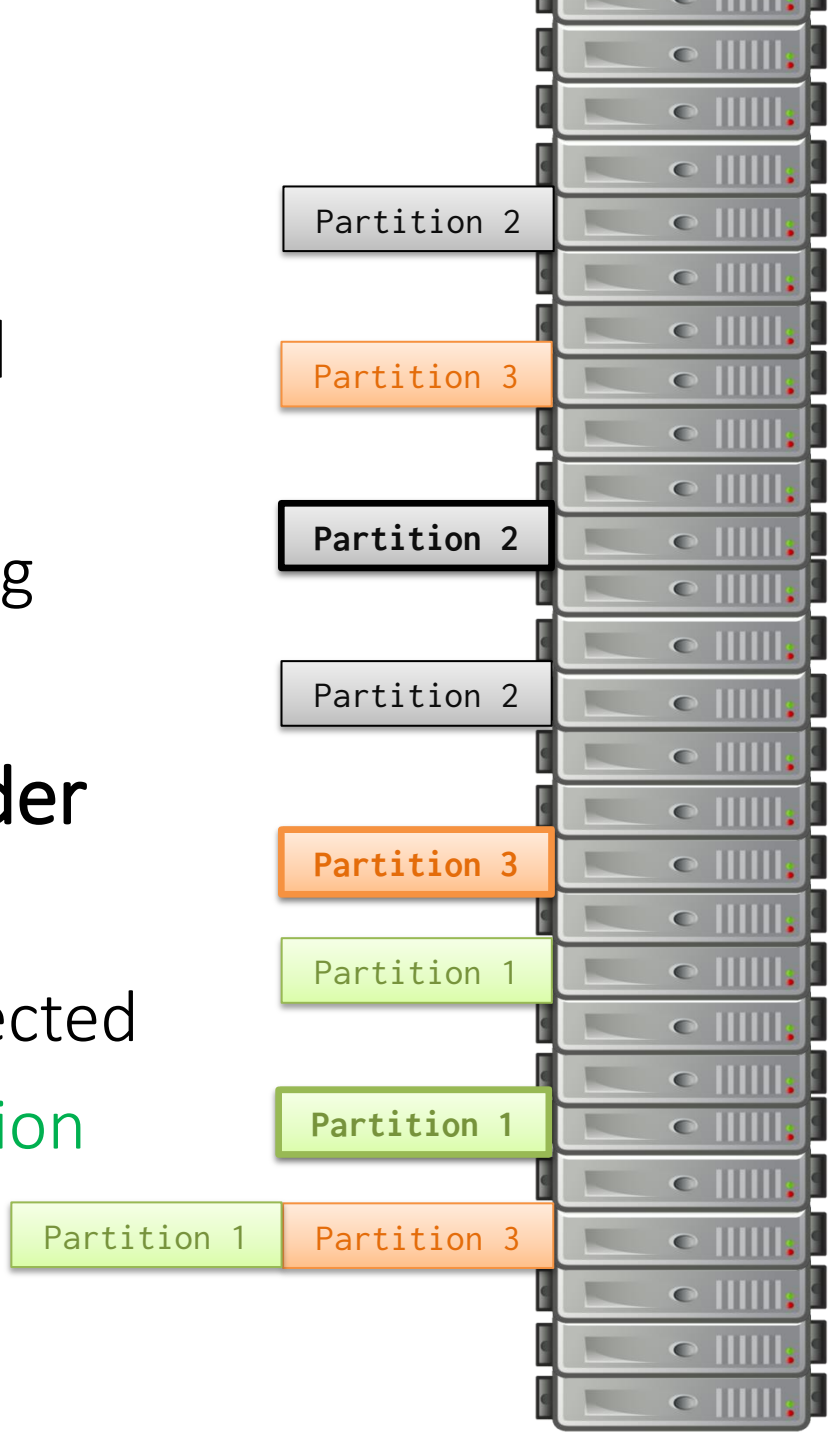

Kafka: Write Guarantees

Writes: Asynchronous (No Guarantee)

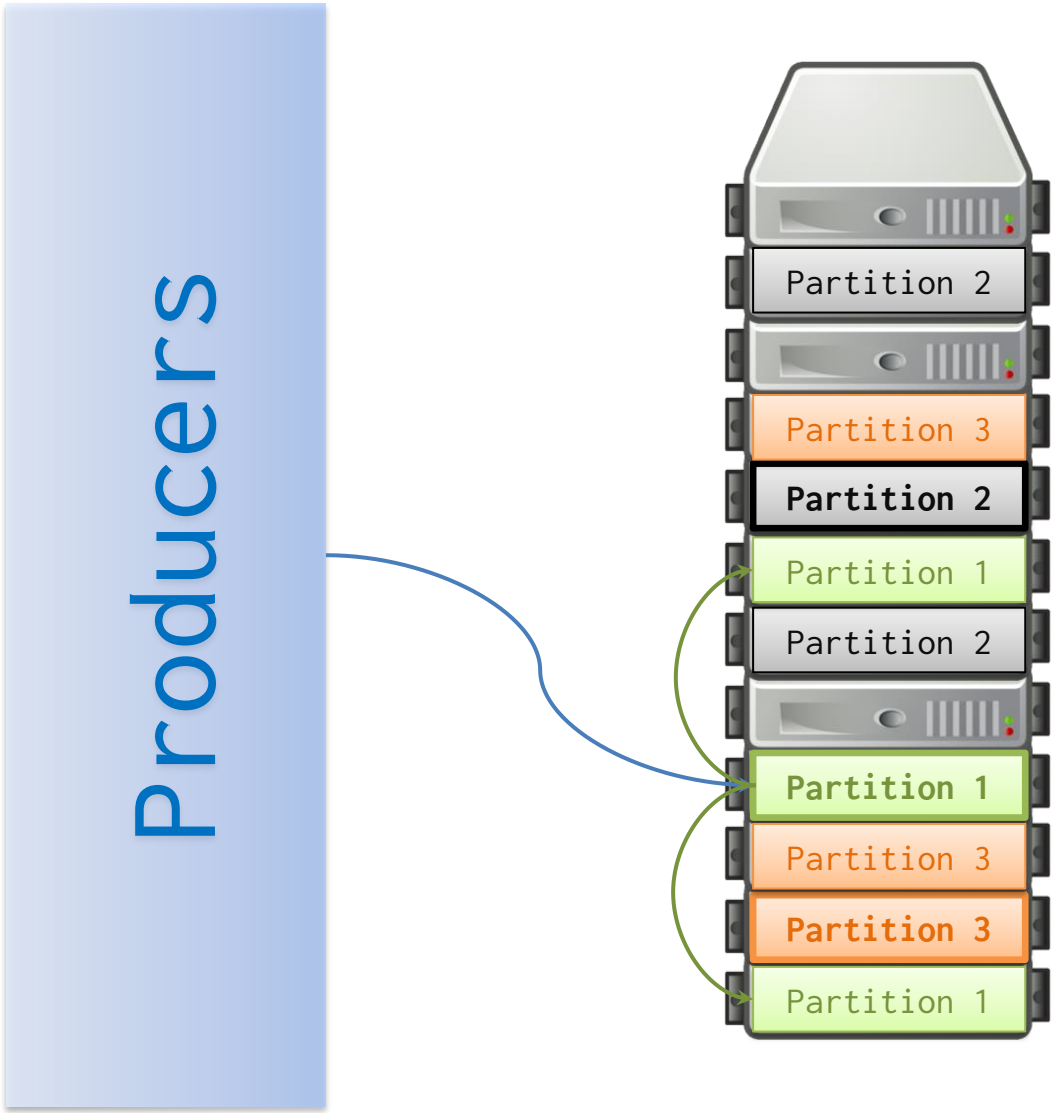

Writes: Leader Commit

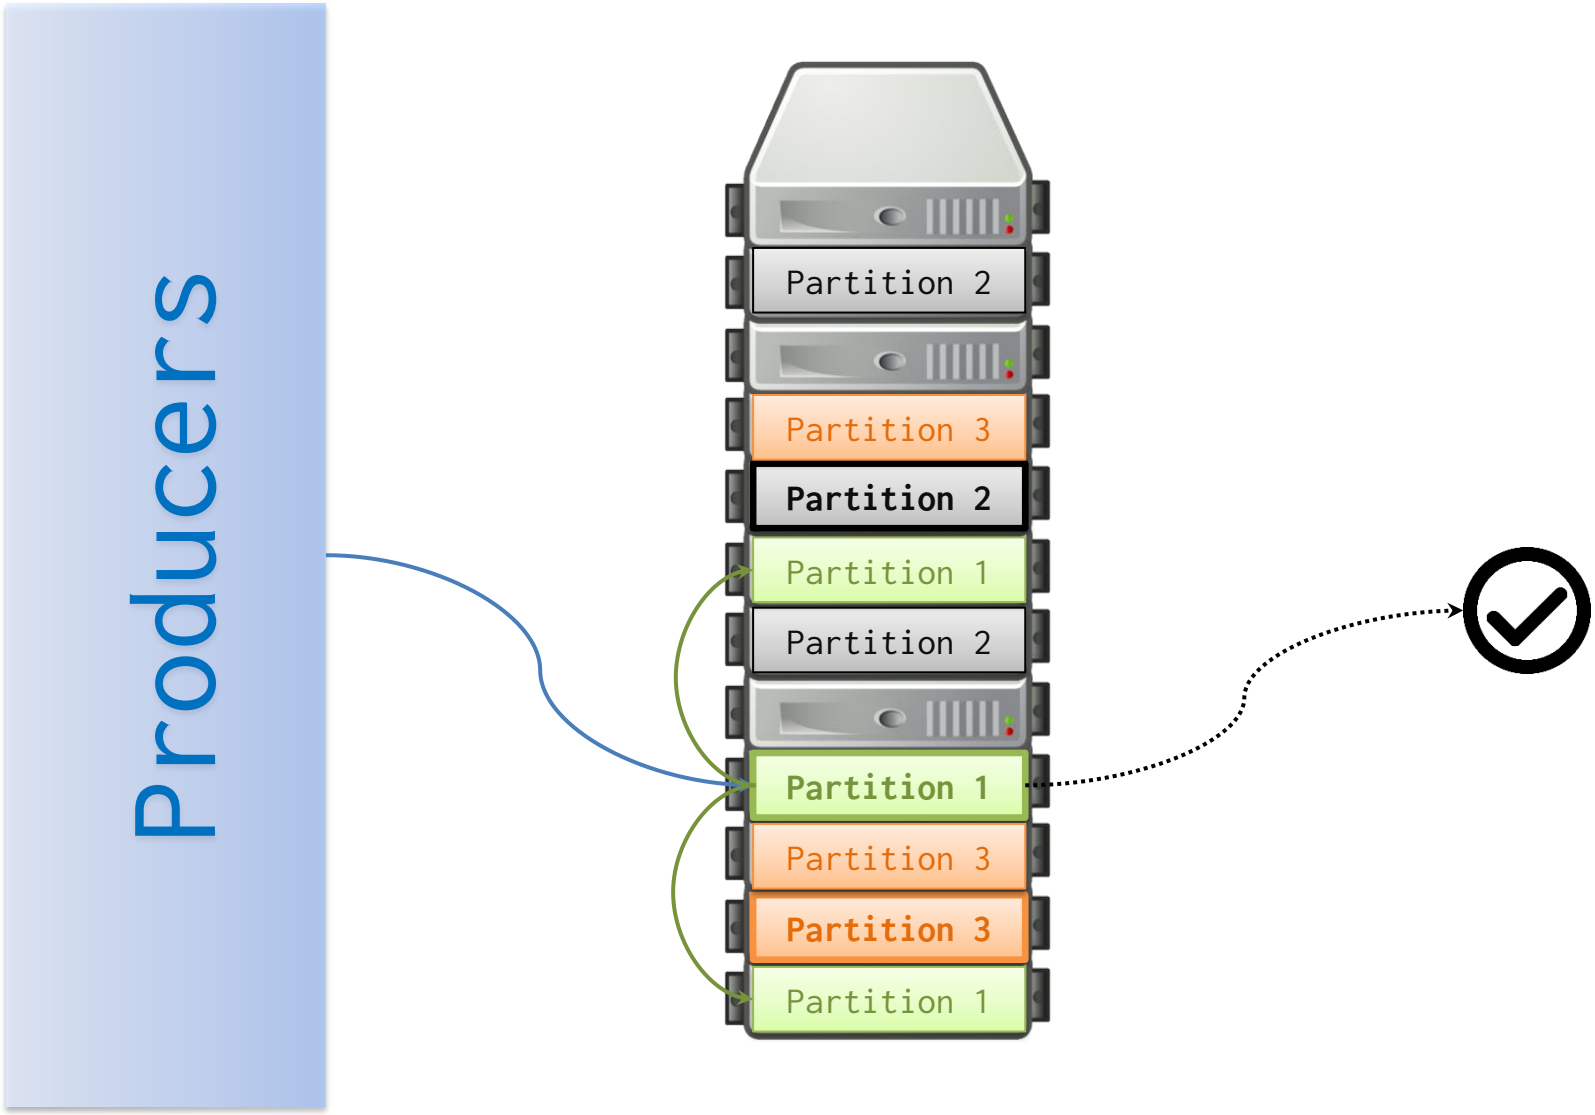

Writes: Leader Commit + Quorum (2)

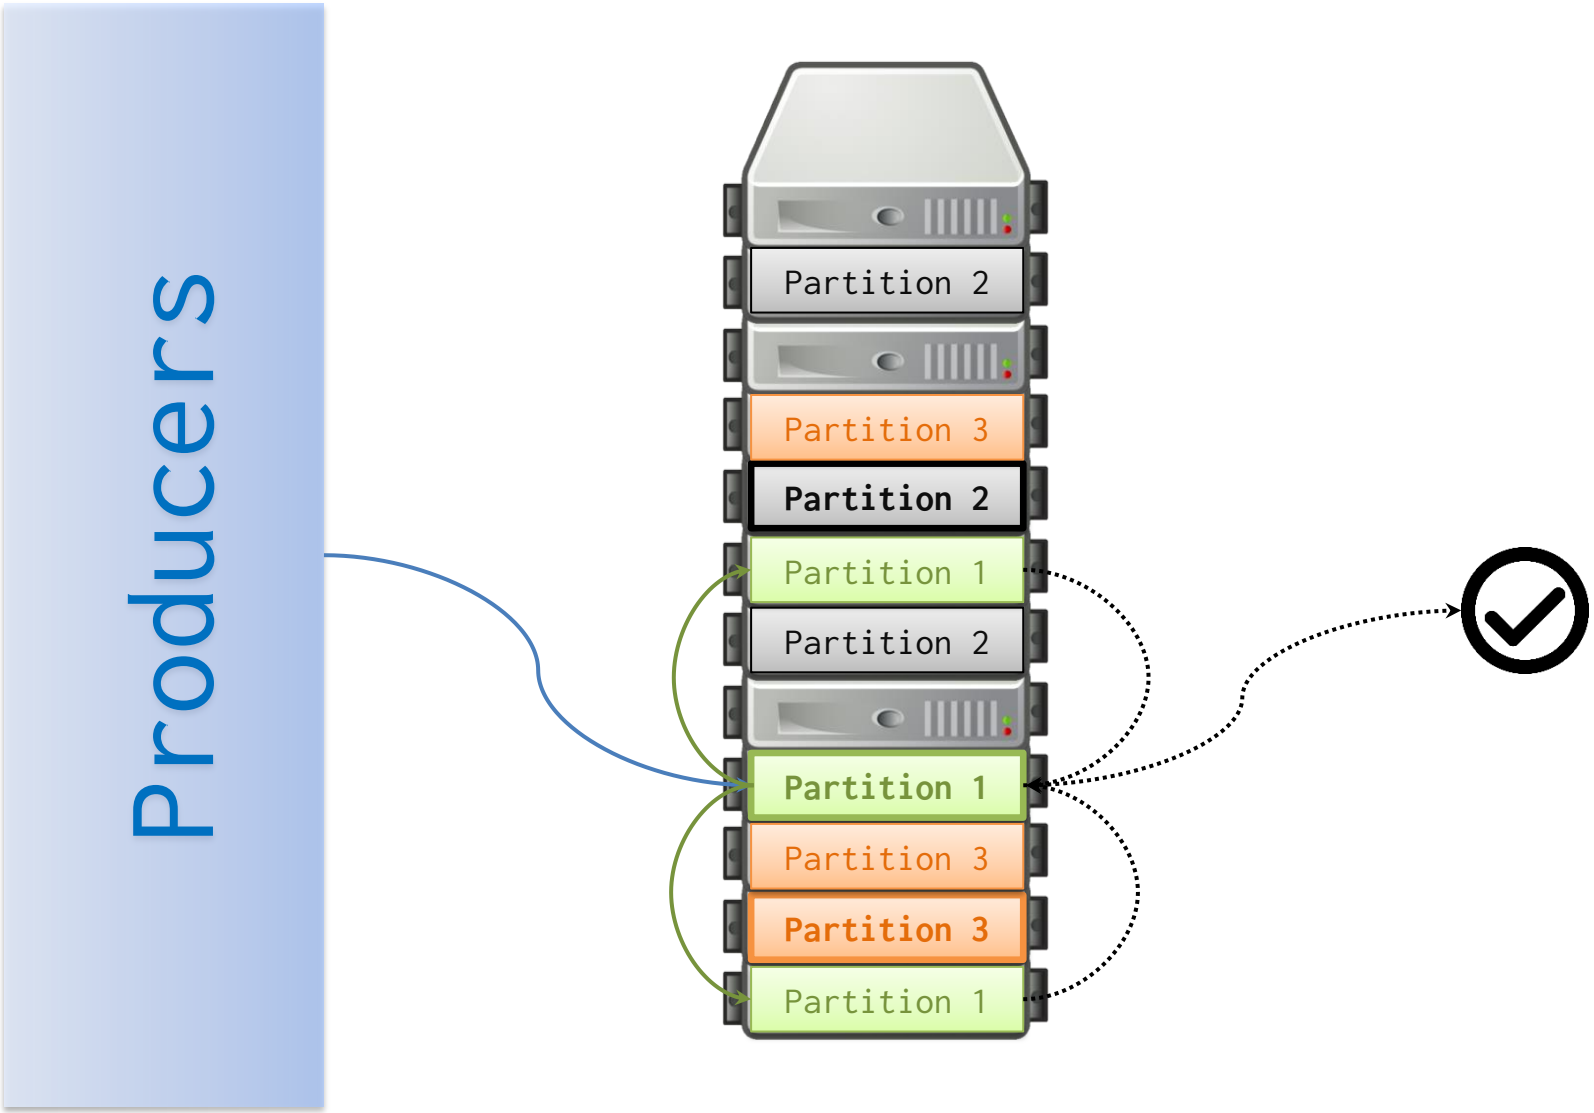

Write Guarantees

- Asynchronous
  - No guarantee
  - Very low latency
- Leader Commit
  - Persistent on leader
  - Medium latency (disk write + network ack)
- Leader Commit + Quorum *n*
  - Persistent on leader + *n* machines
  - High latency (disk writes + network acks)

Kafka: Reads

Kafka tracks consumer offset

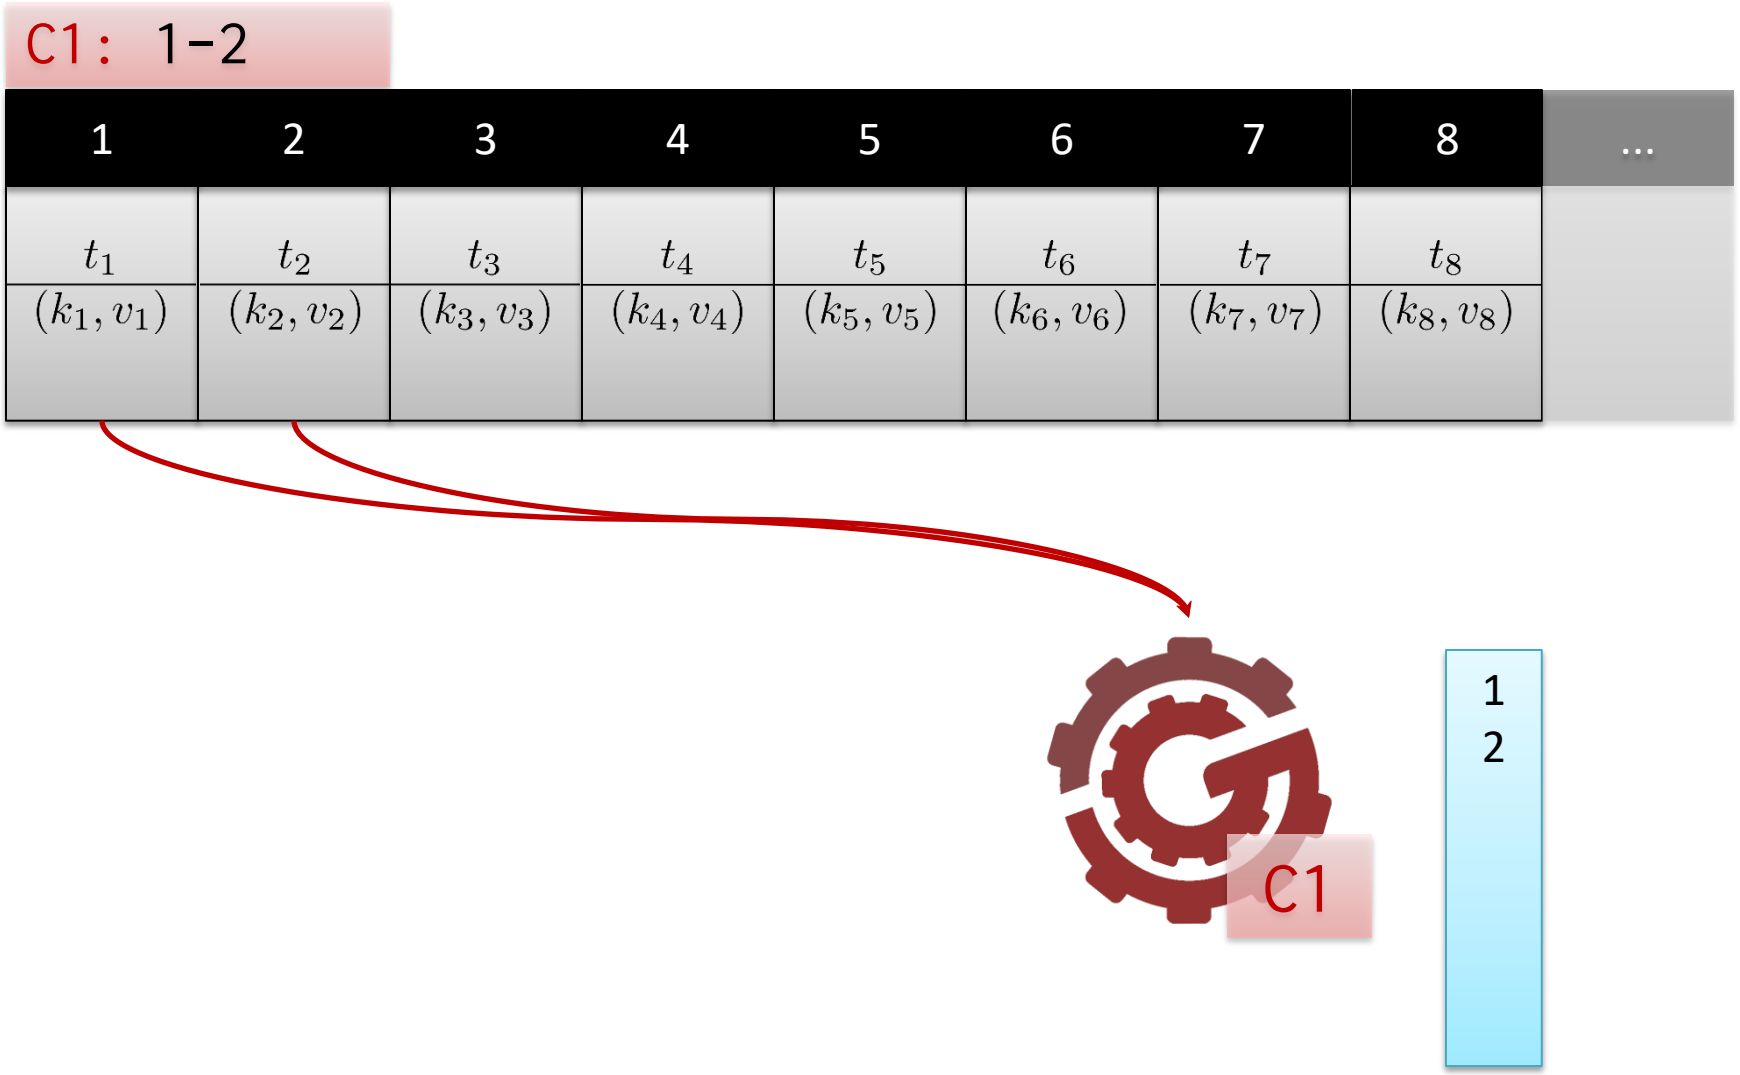

Kafka tracks consumer offset

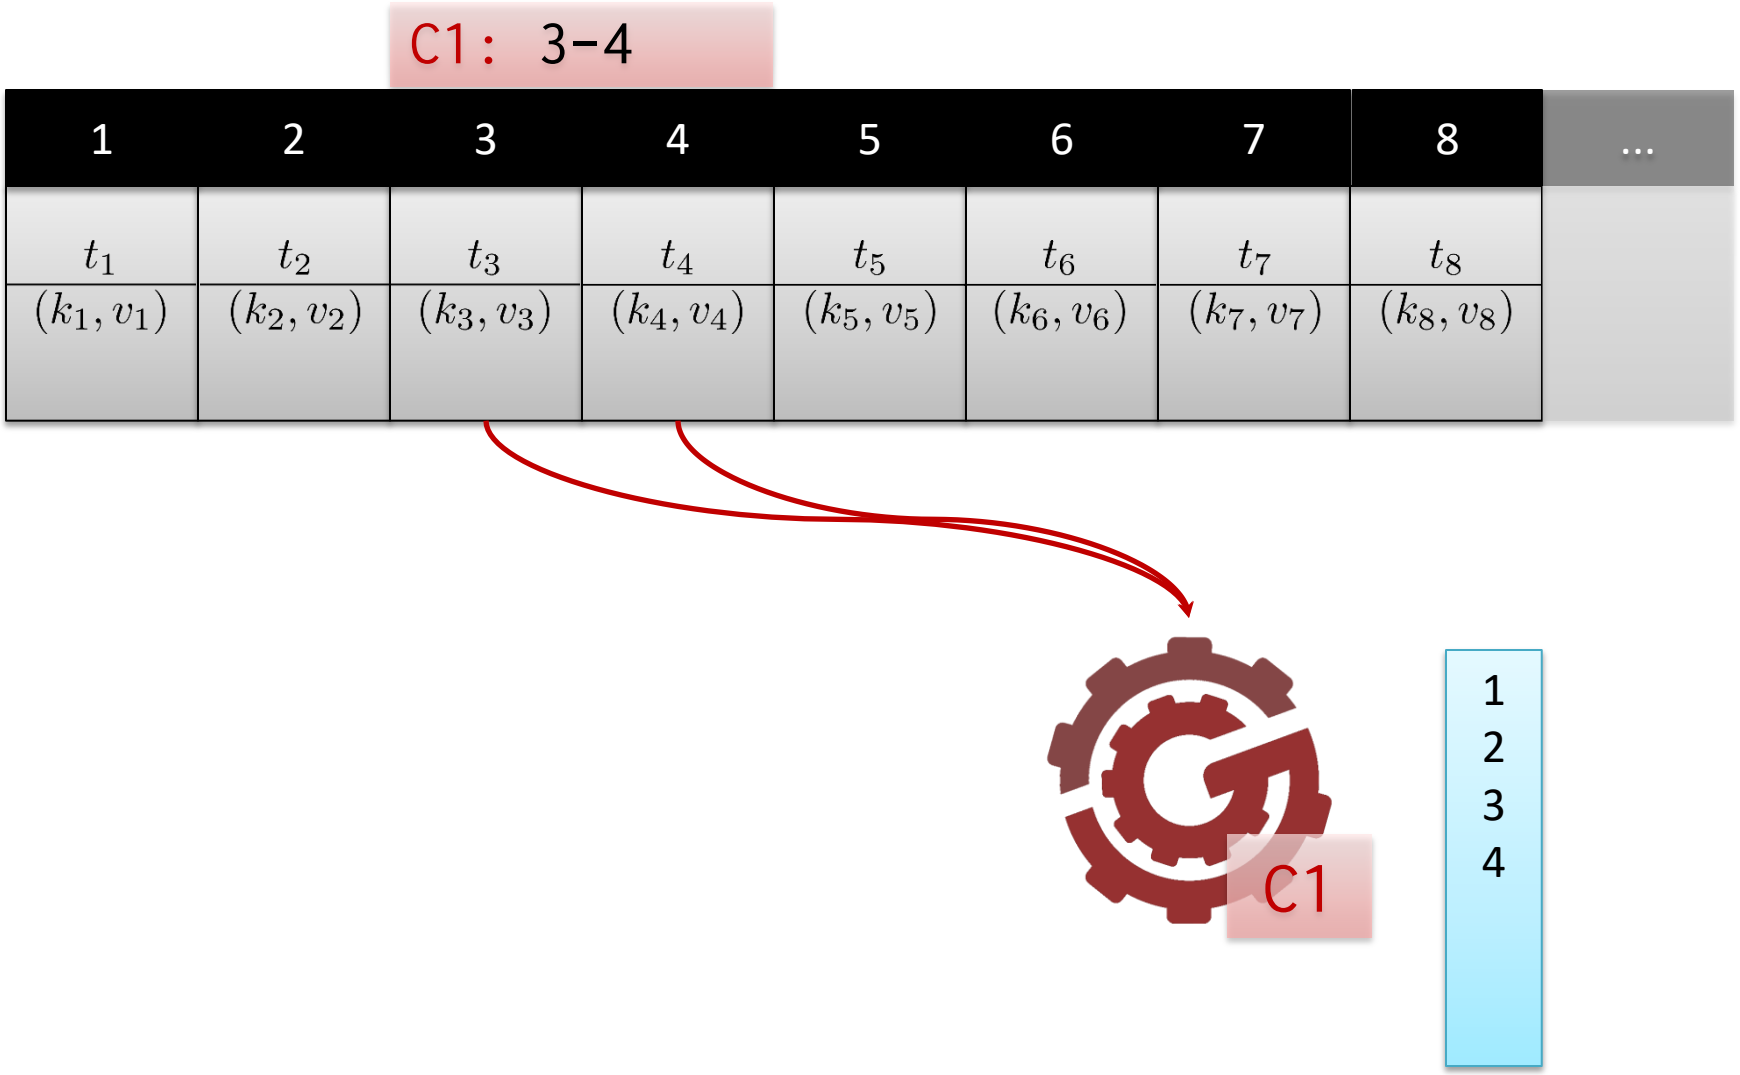

Kafka tracks consumer offset

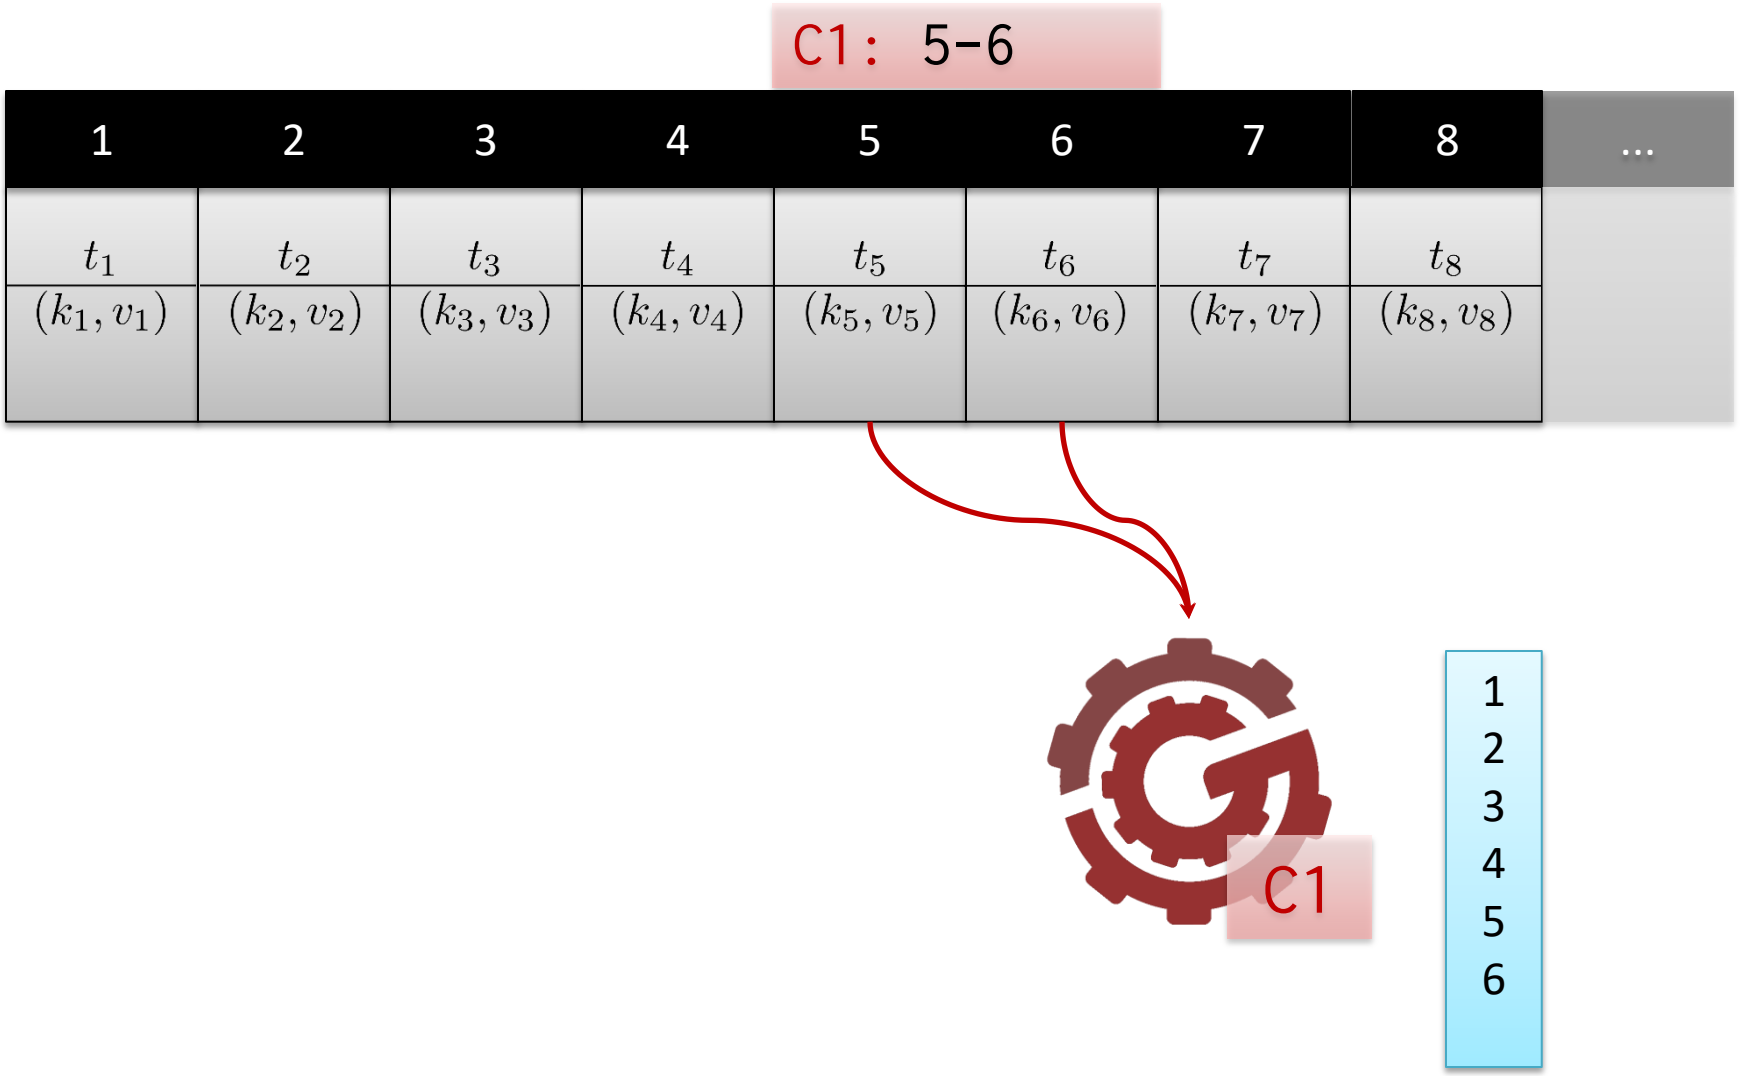

Failures?

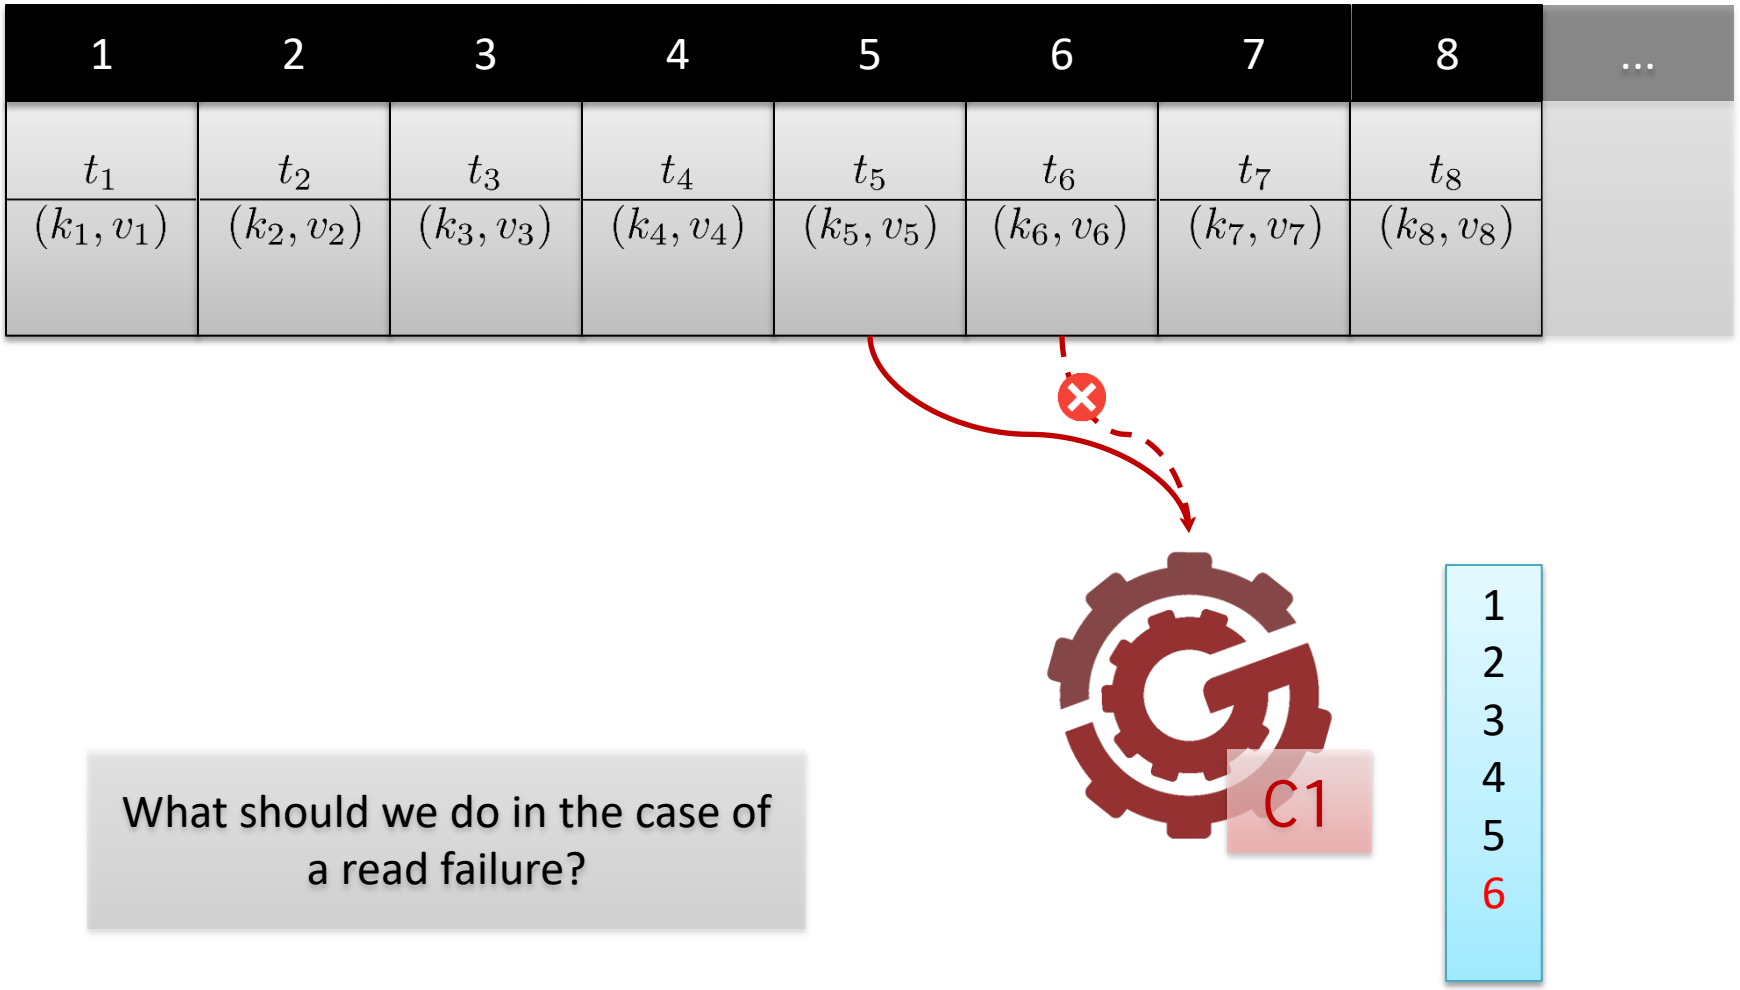
1
2
3
4
5
6
7
8
...
t_1 (k_1, v_1)
t_2 (k_2, v_2)
t_3 (k_3, v_3)
t_4 (k_4, v_4)
t_5 (k_5, v_5)
t_6 (k_6, v_6)
t_7 (k_7, v_7)
t_8 (k_8, v_8)
C1
1
2
3
4
5
6
What should we do in the case of a read failure?

Kafka: Read Guarantees

Read Guarantees

- At least once
  - Each value processed at least once
  - Consumer offset updated on consumer ACK
- At most once
- Effectively once
- Exactly once

Read: At Least Once (Default)

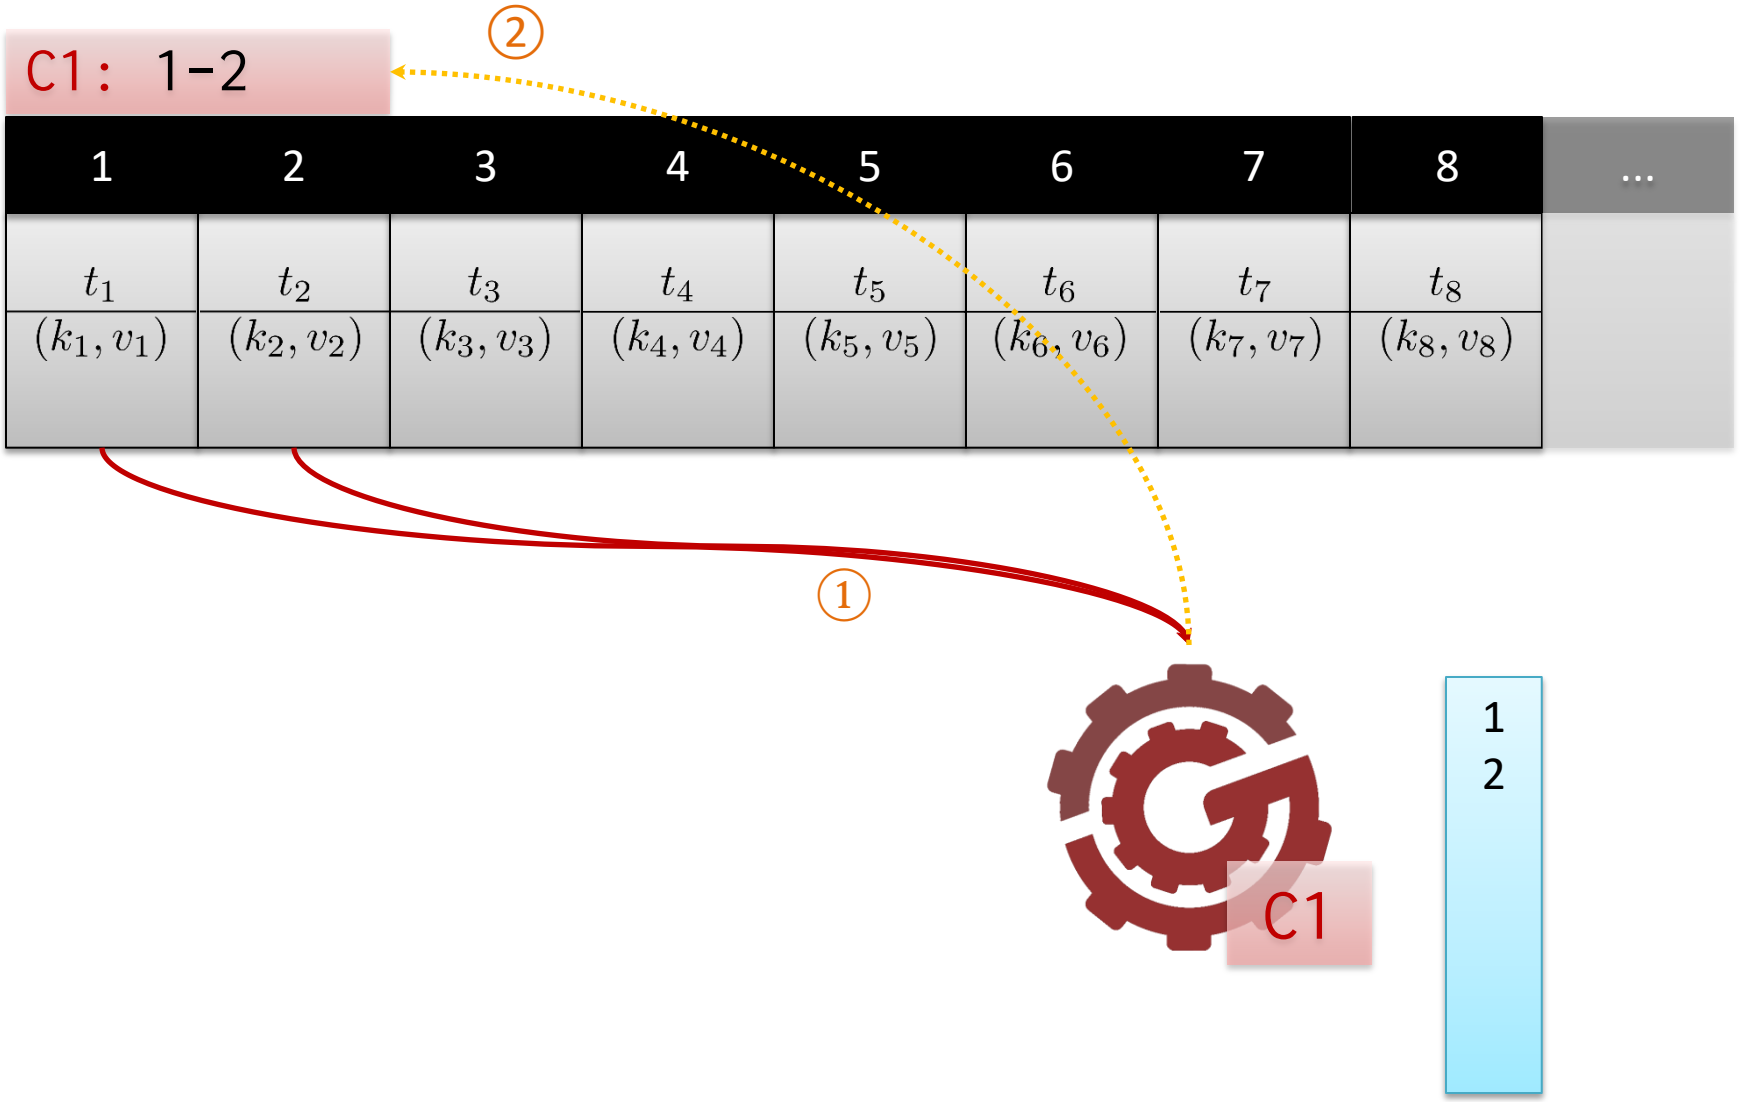

Read: At Least Once (Default)

C1: 1-2

1	2	3	4	5	6	7	8	...
t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	
(k_1, v_1)	(k_2, v_2)	(k_3, v_3)	(k_4, v_4)	(k_5, v_5)	(k_6, v_6)	(k_7, v_7)	(k_8, v_8)	

1
2

Read: At Least Once (Default)

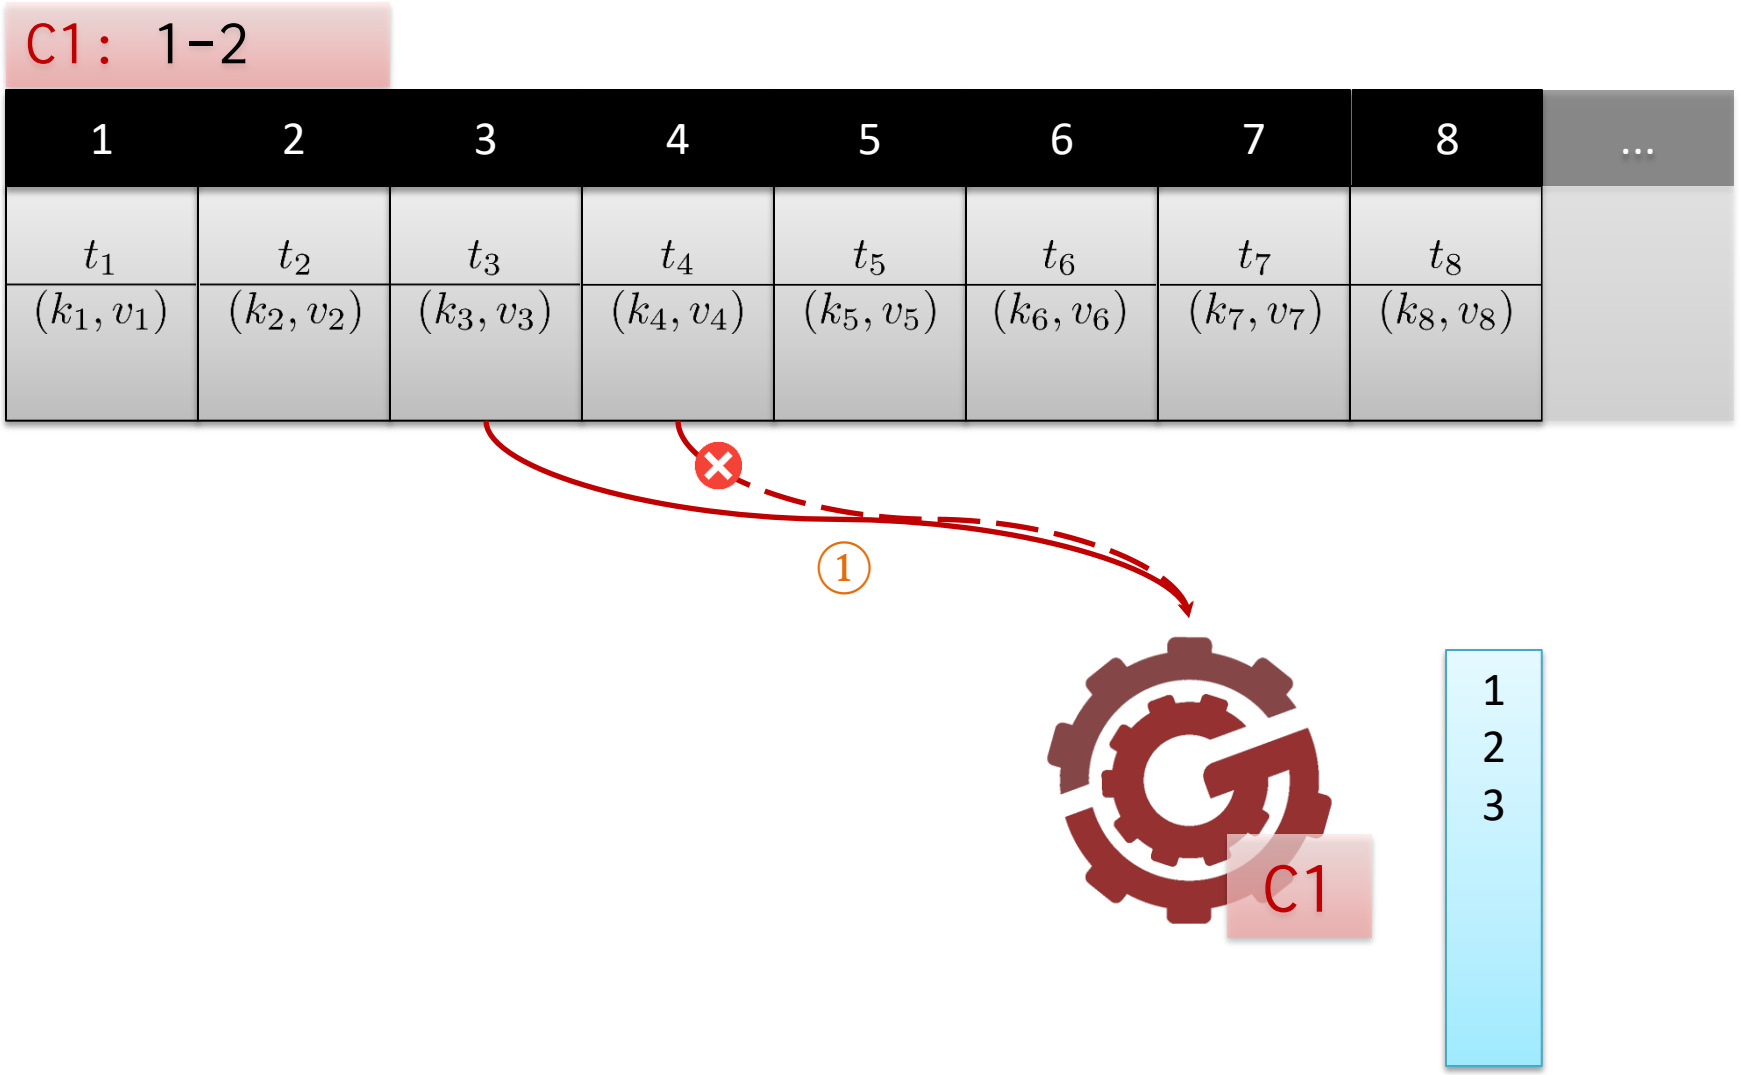

Read: At Least Once (Default)

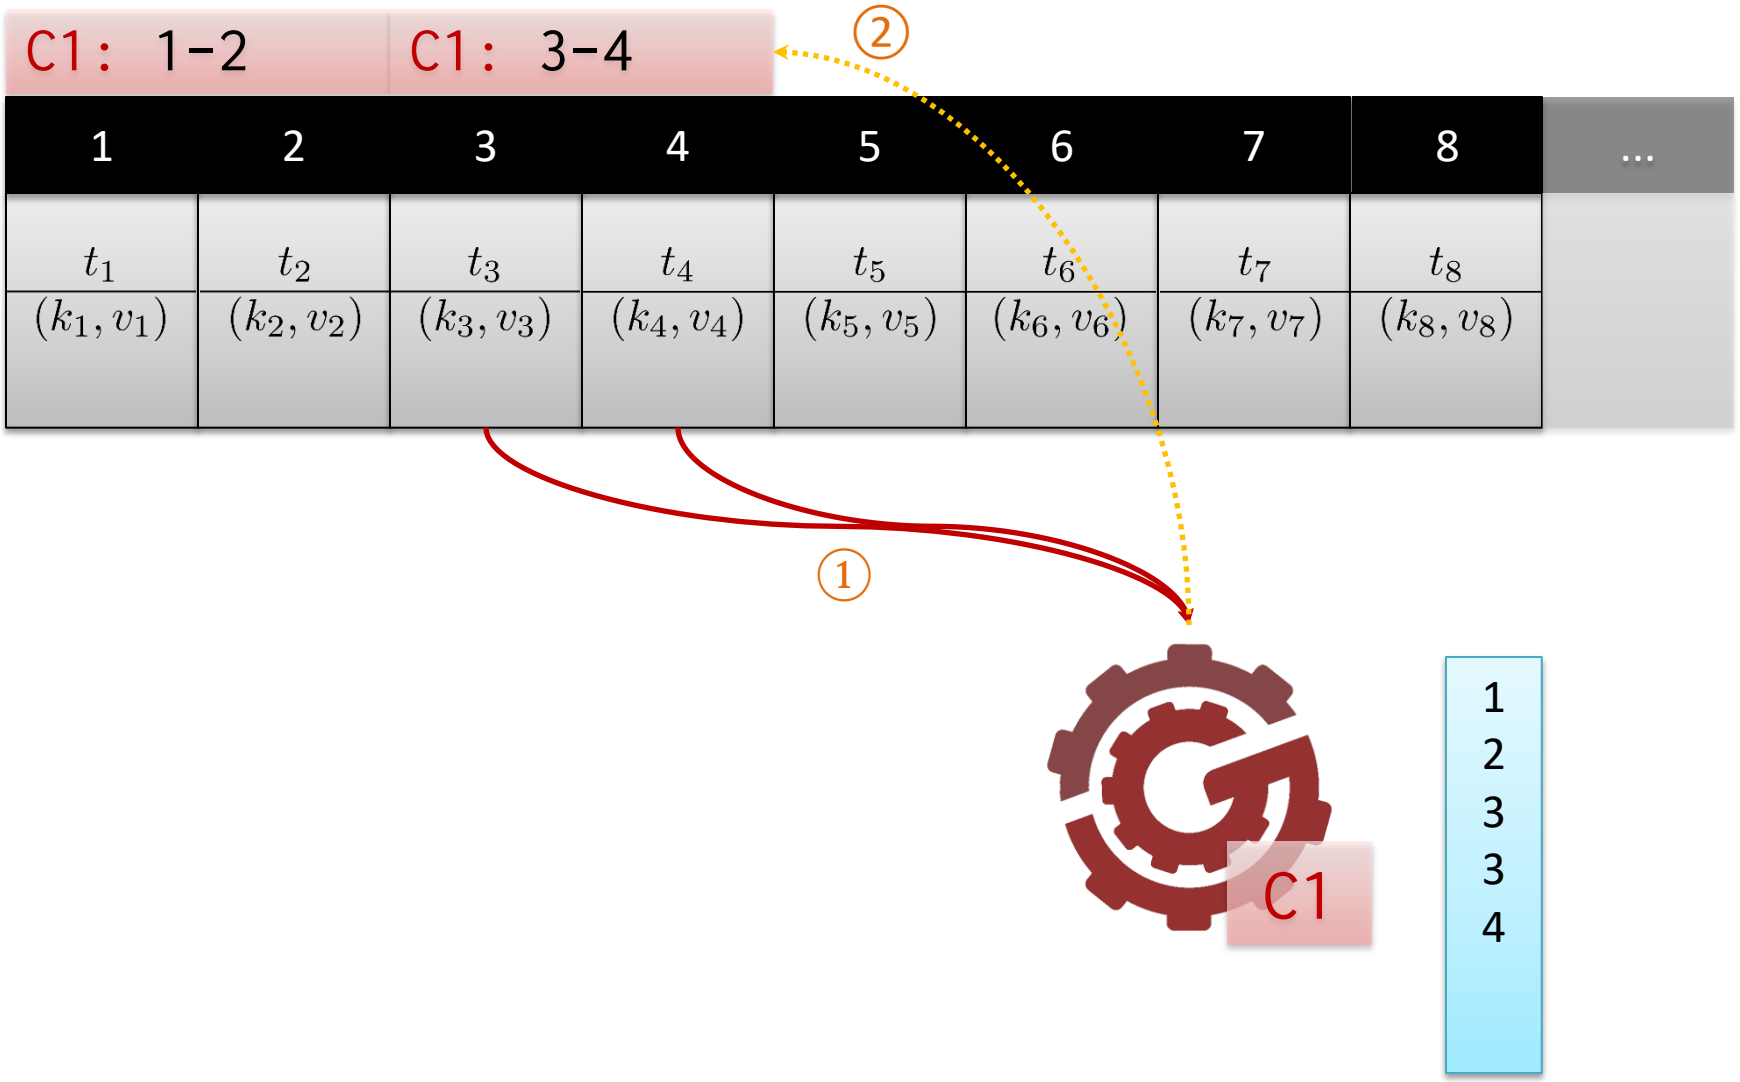

Read: At Least Once (Default)

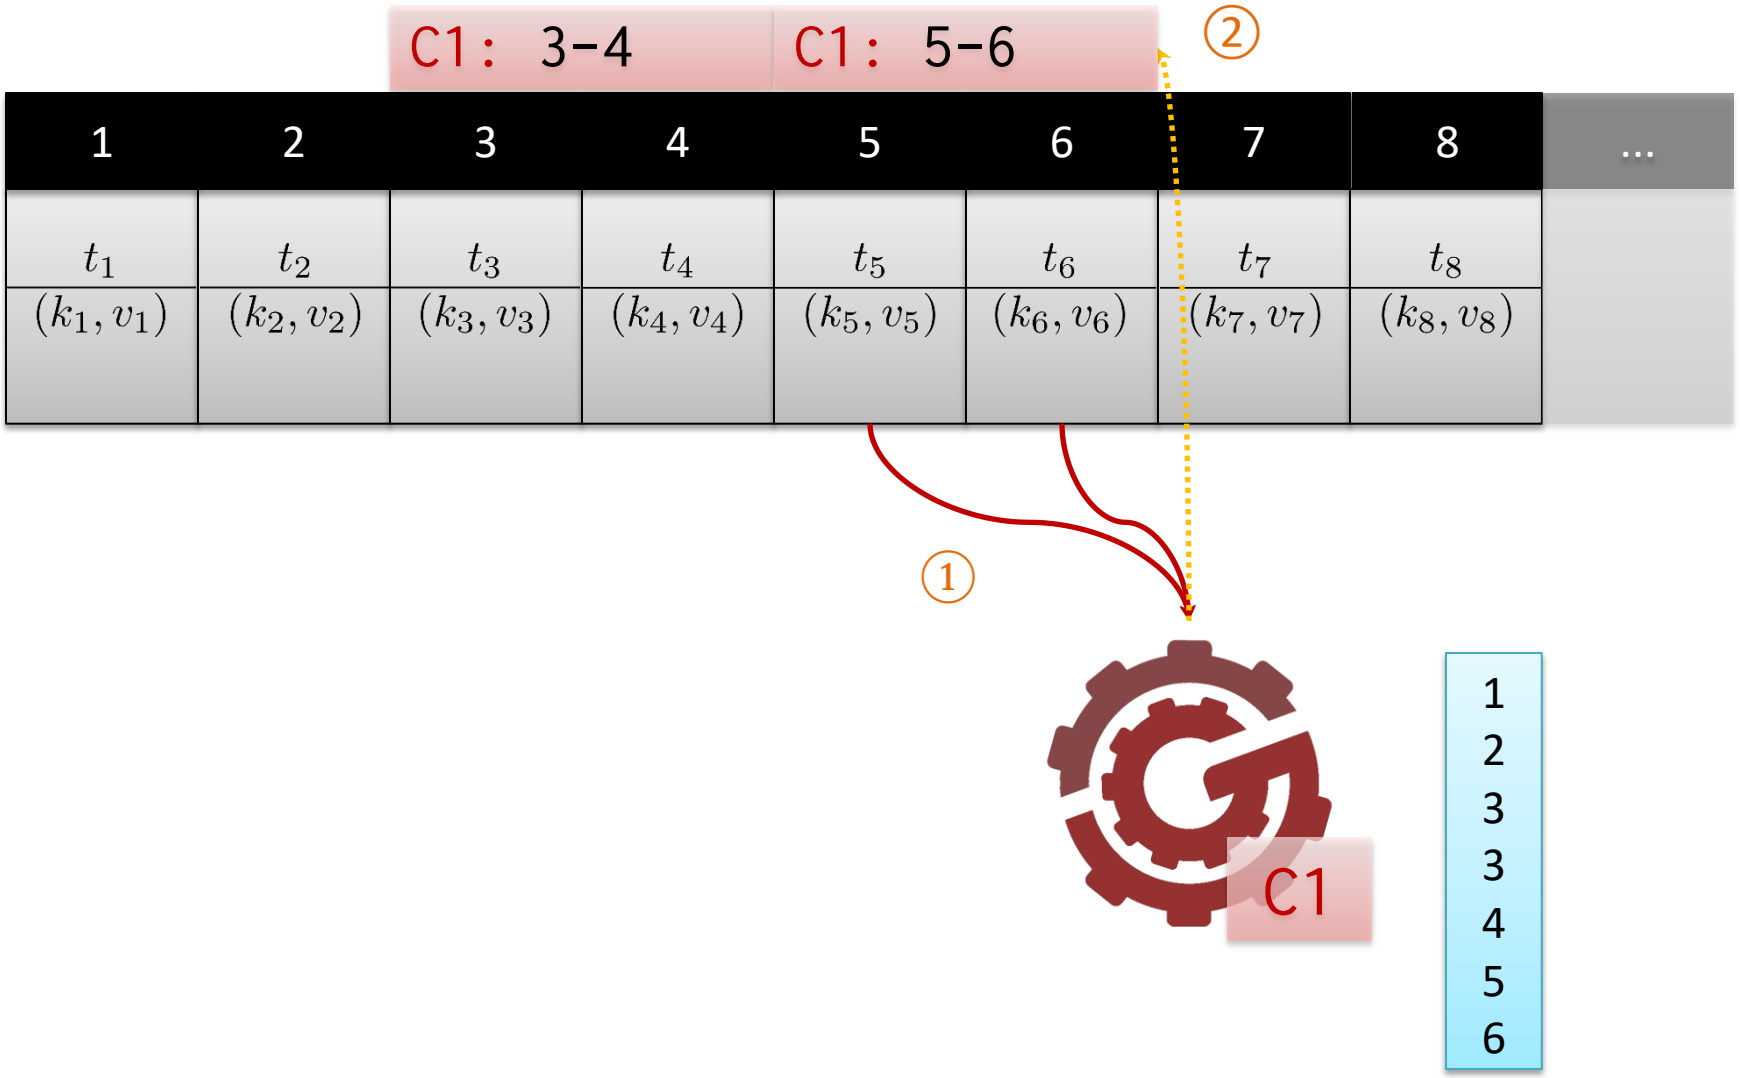

Read Guarantees

- At least once
- At most once
  - Each value processed at most once
  - Consumer offset updated immediately
- Effectively once
- Exactly once

Read: At Most Once

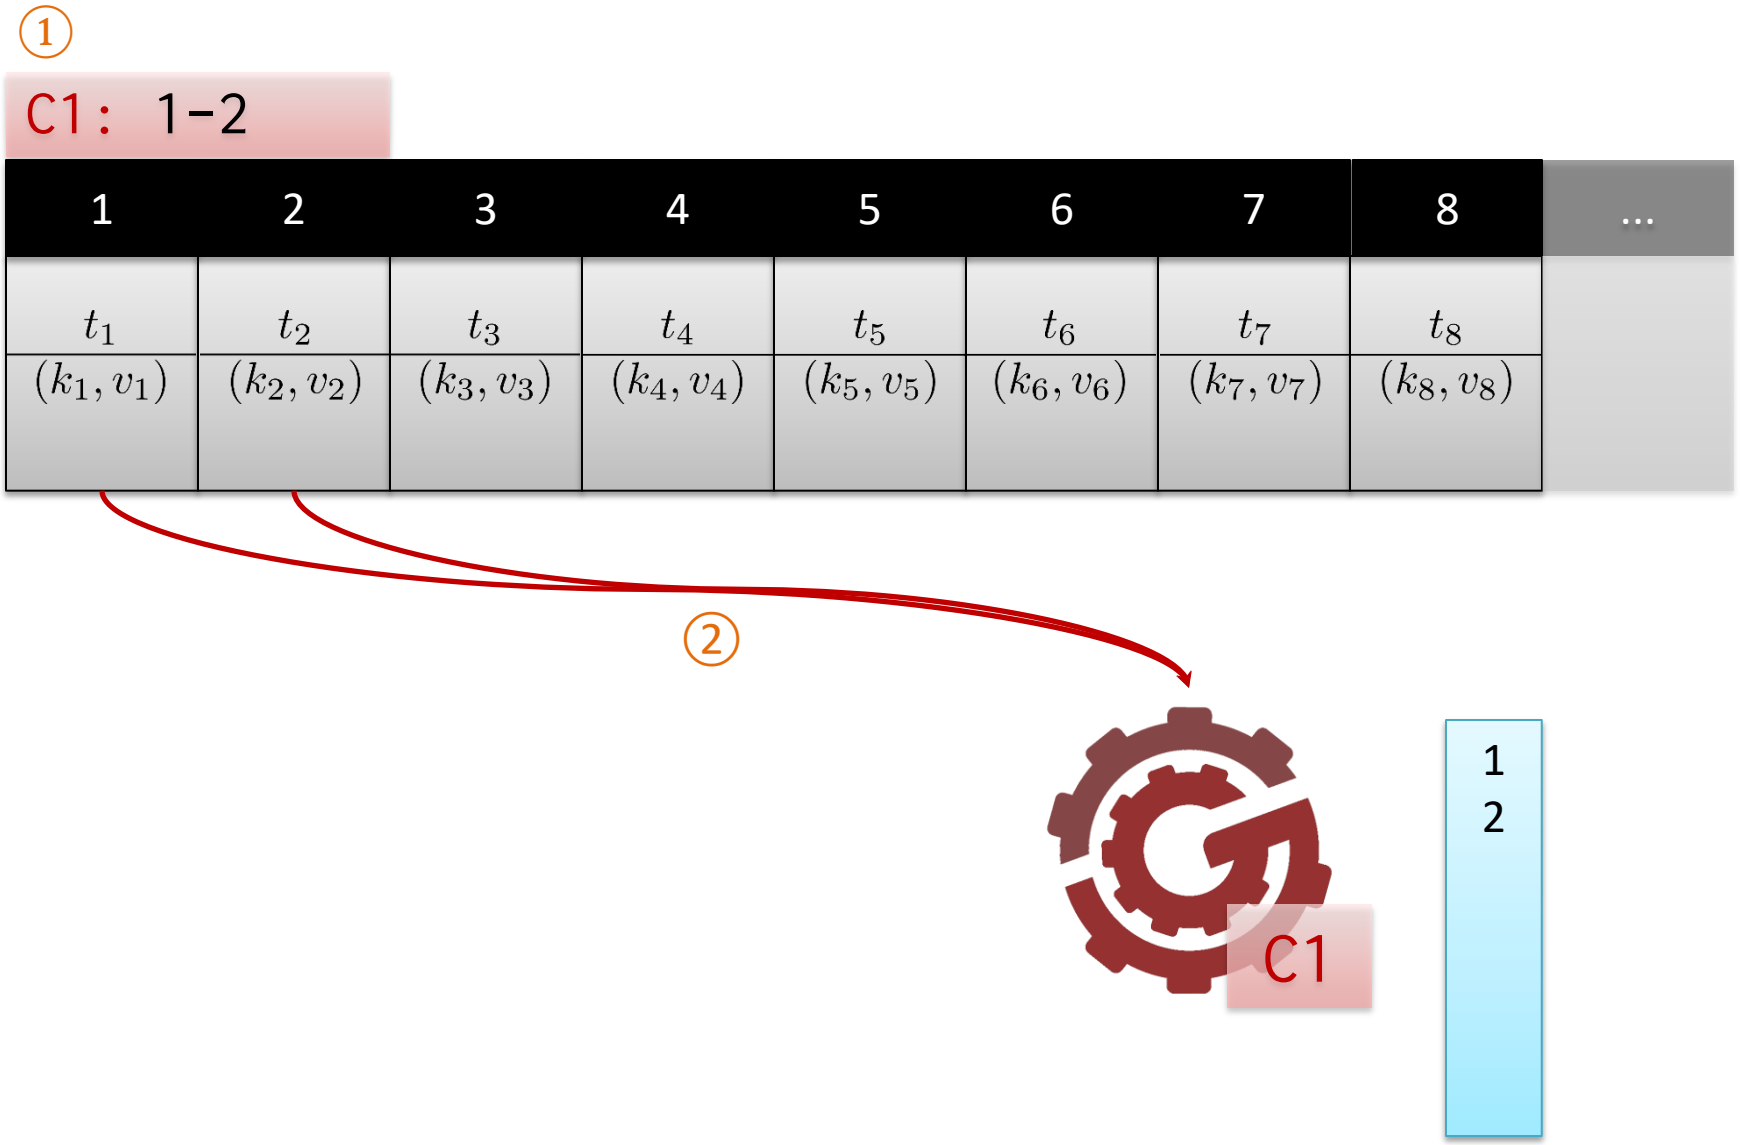

Read: At Most Once

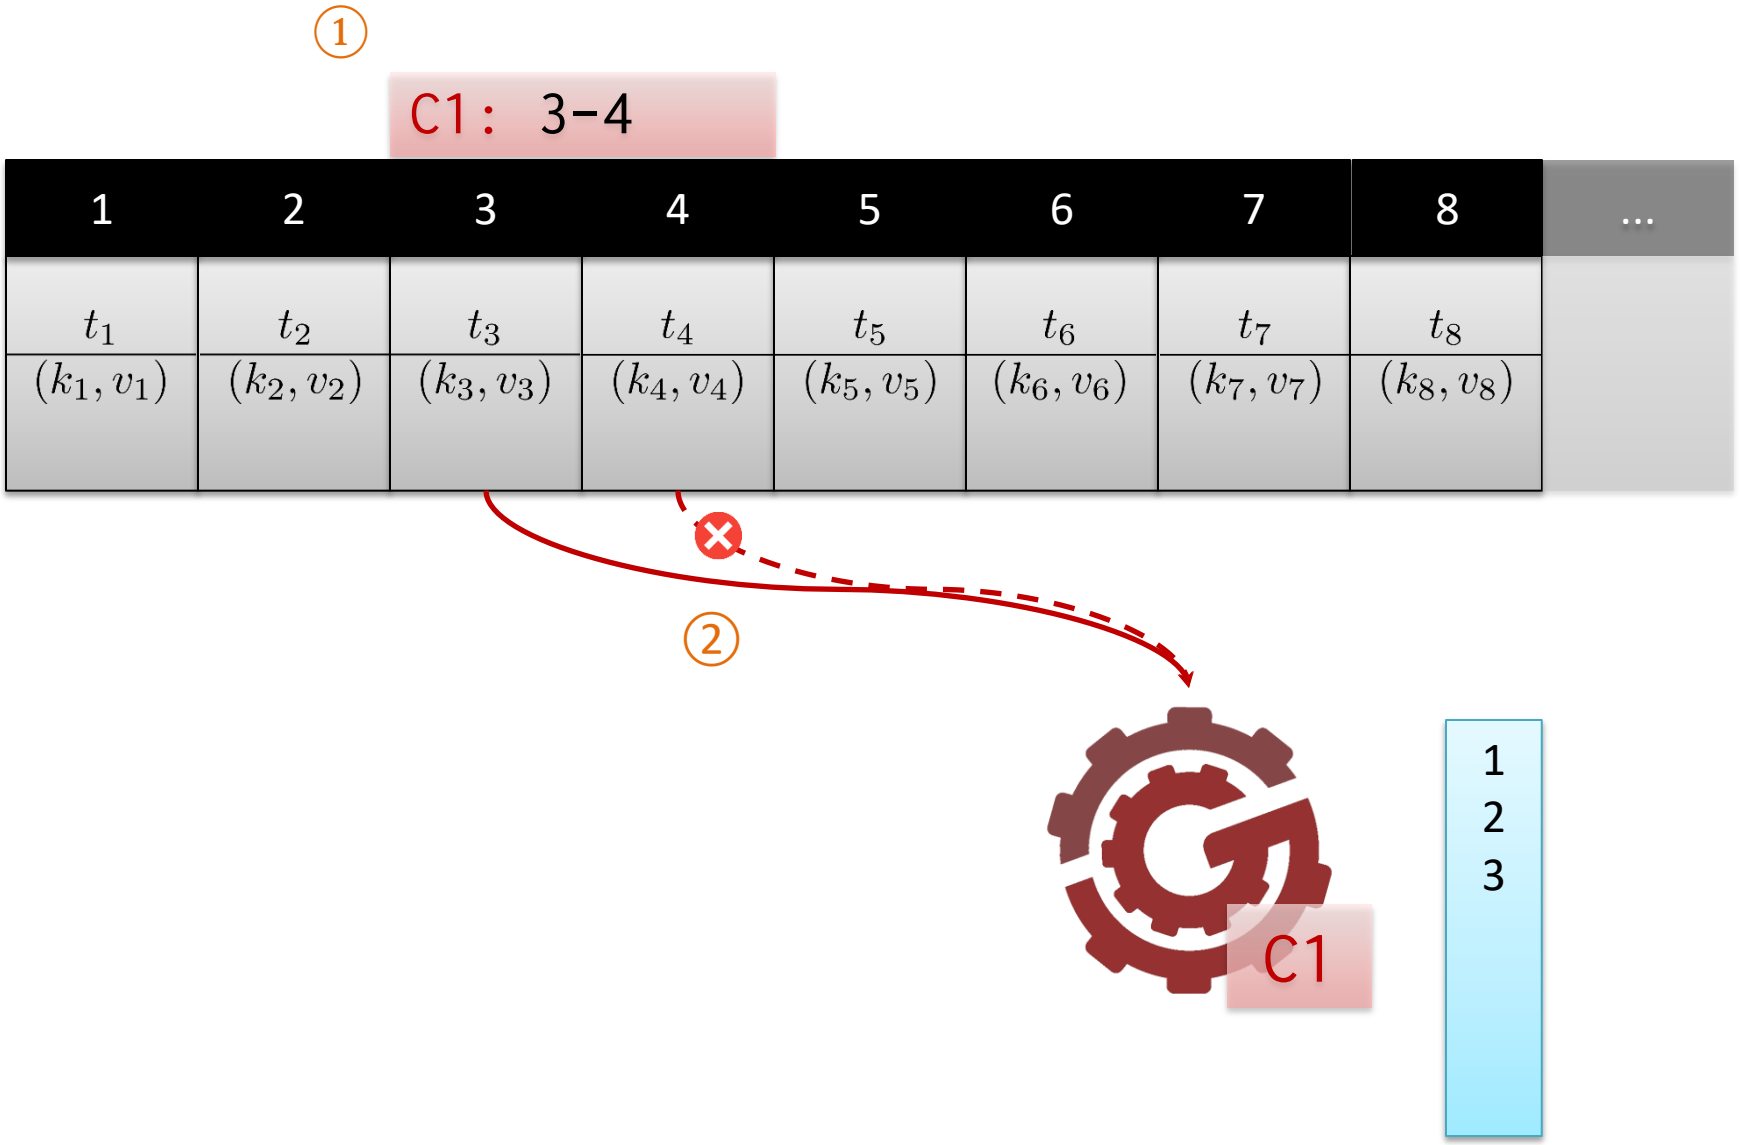

Read: At Most Once

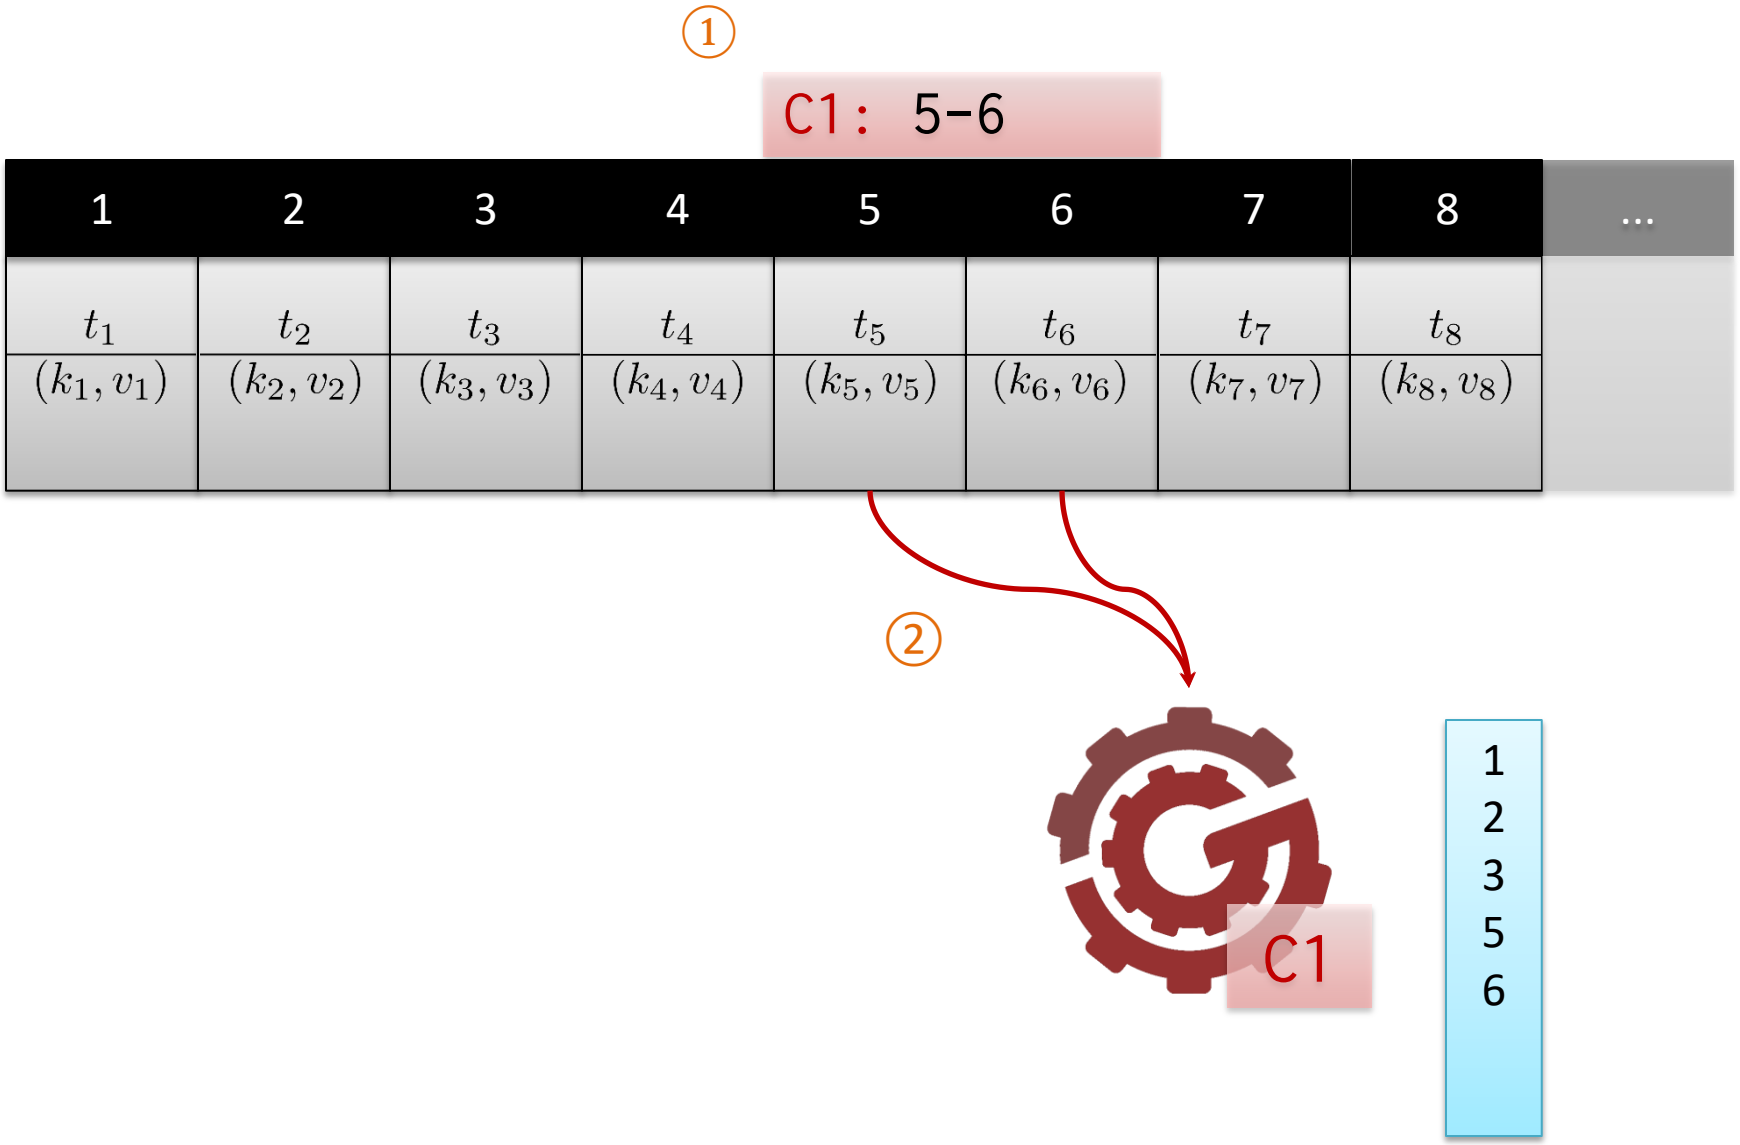

Read Guarantees

- At least once
- At most once
- Effectively once
  - At least once but …
  - Consumer takes care of duplicates
- Exactly once

Read: Effectively Once

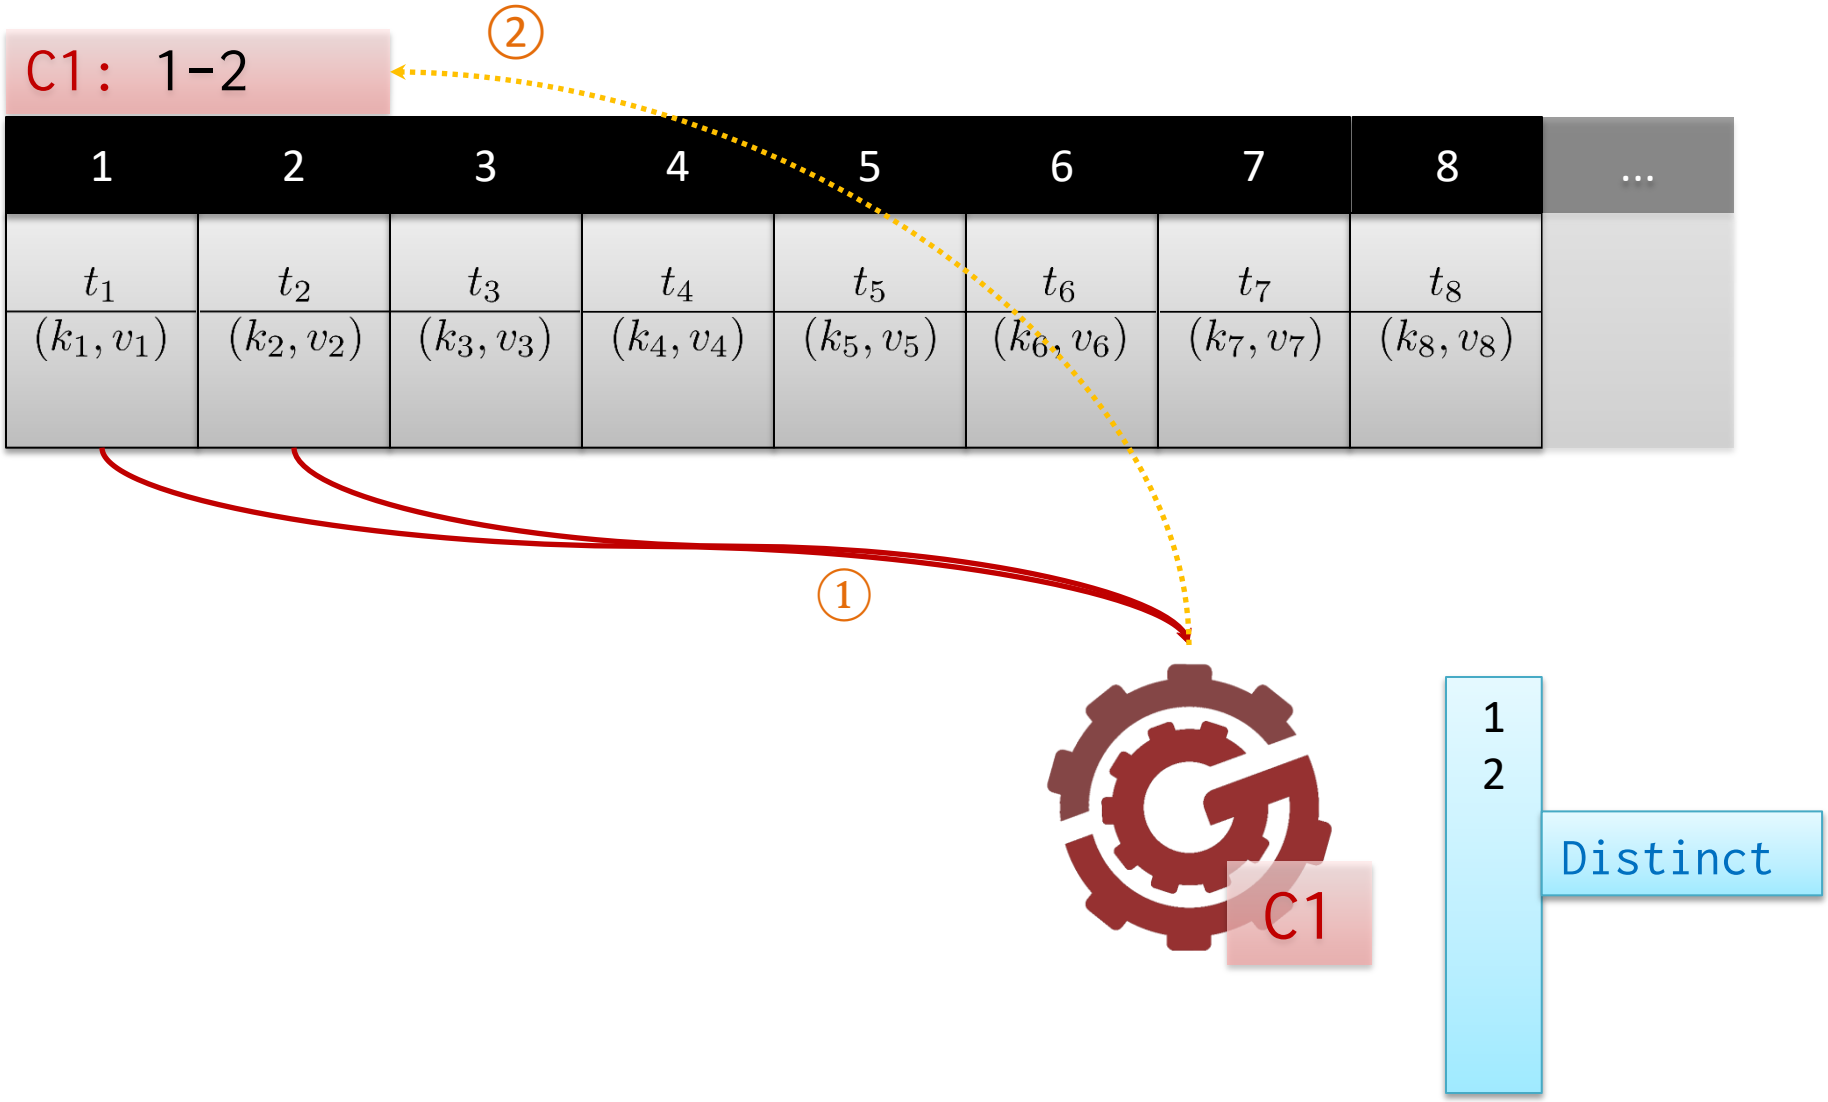

Read: Effectively Once

C1: 1-2

1	2	3	4	5	6	7	8	...
t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	
(k_1, v_1)	(k_2, v_2)	(k_3, v_3)	(k_4, v_4)	(k_5, v_5)	(k_6, v_6)	(k_7, v_7)	(k_8, v_8)	

1
2

Read: Effectively Once

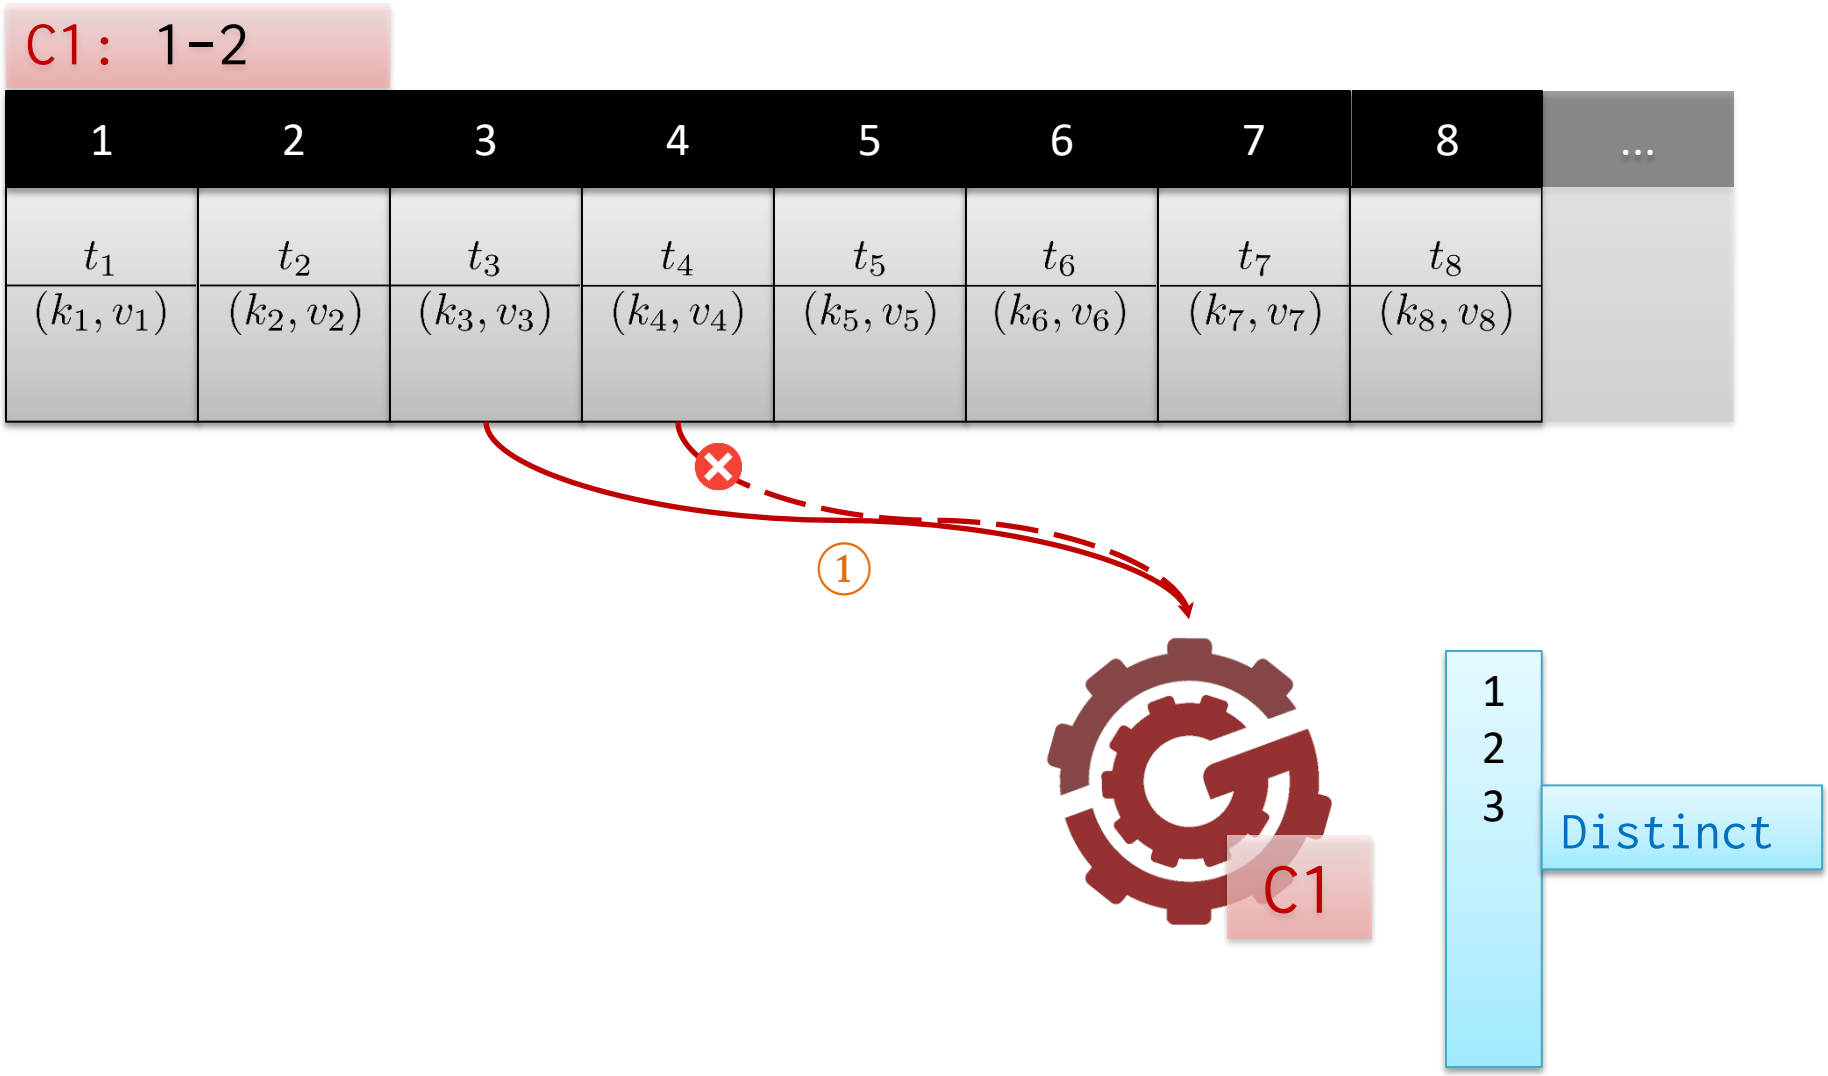

Read: Effectively Once

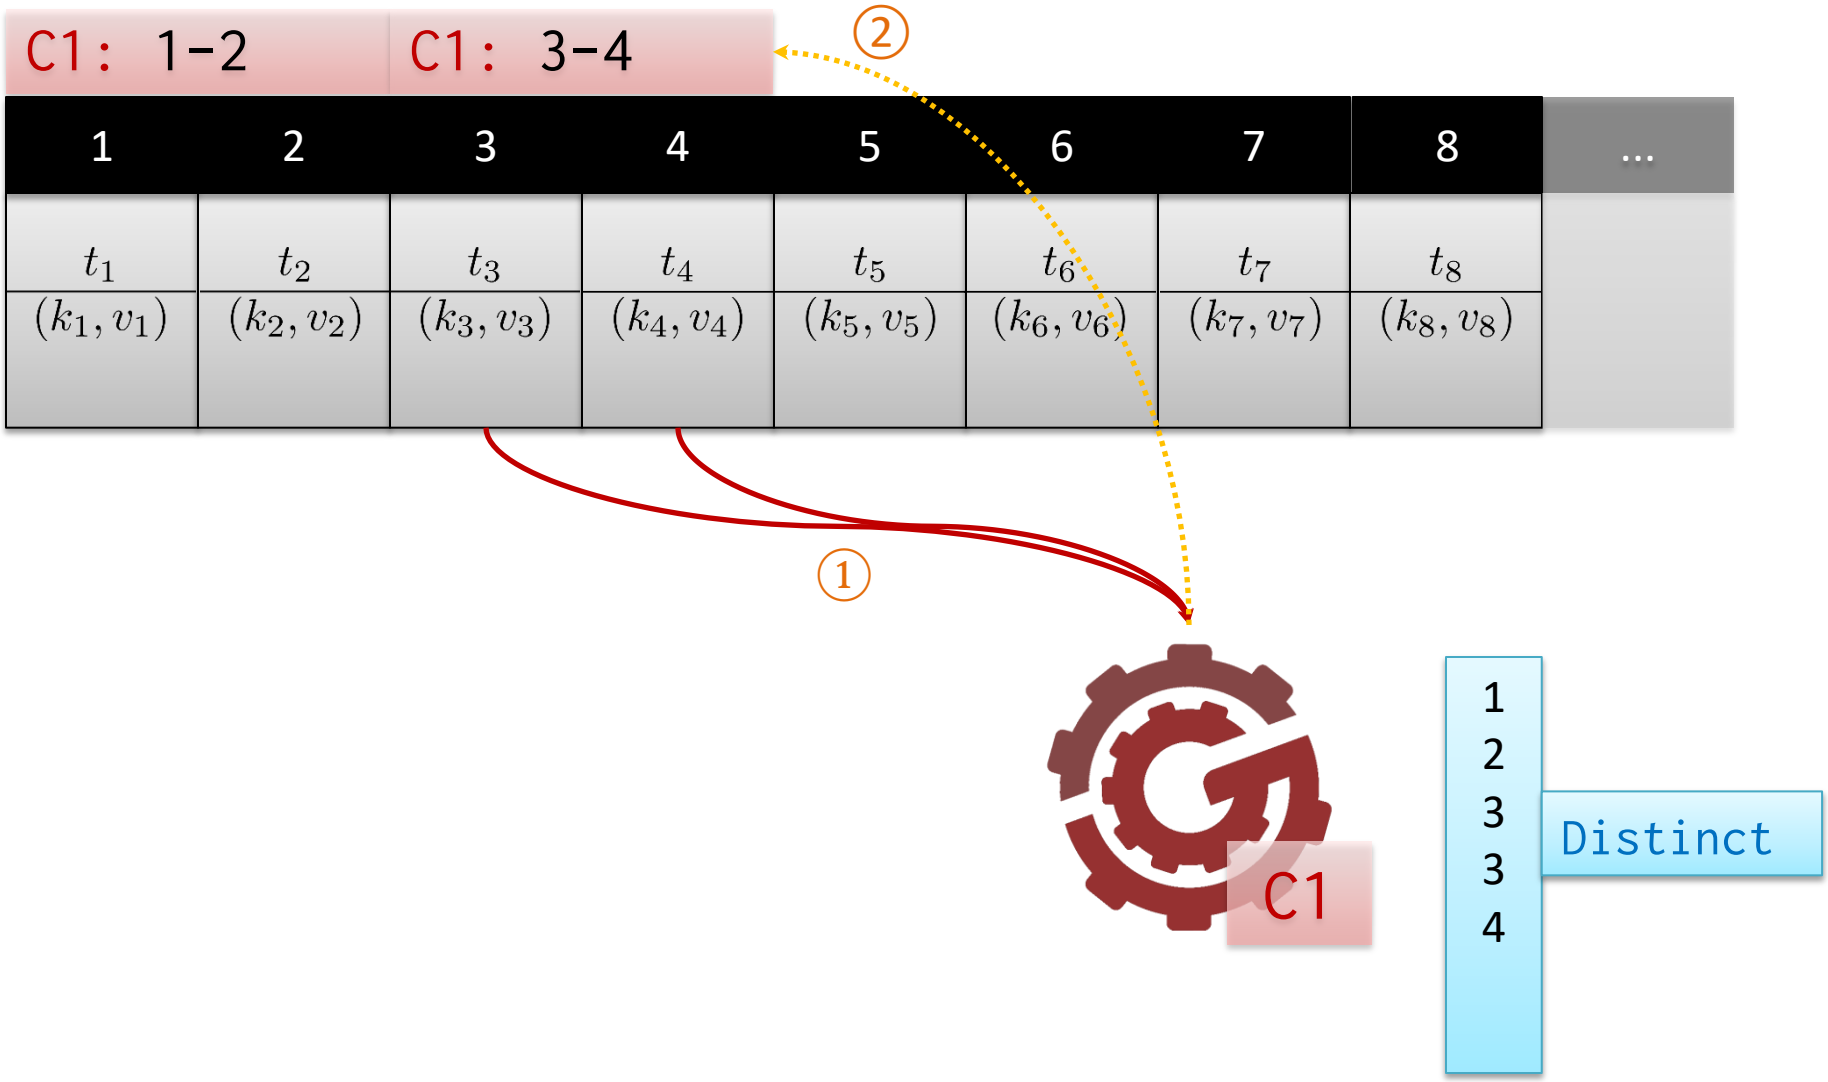

Read: Effectively Once

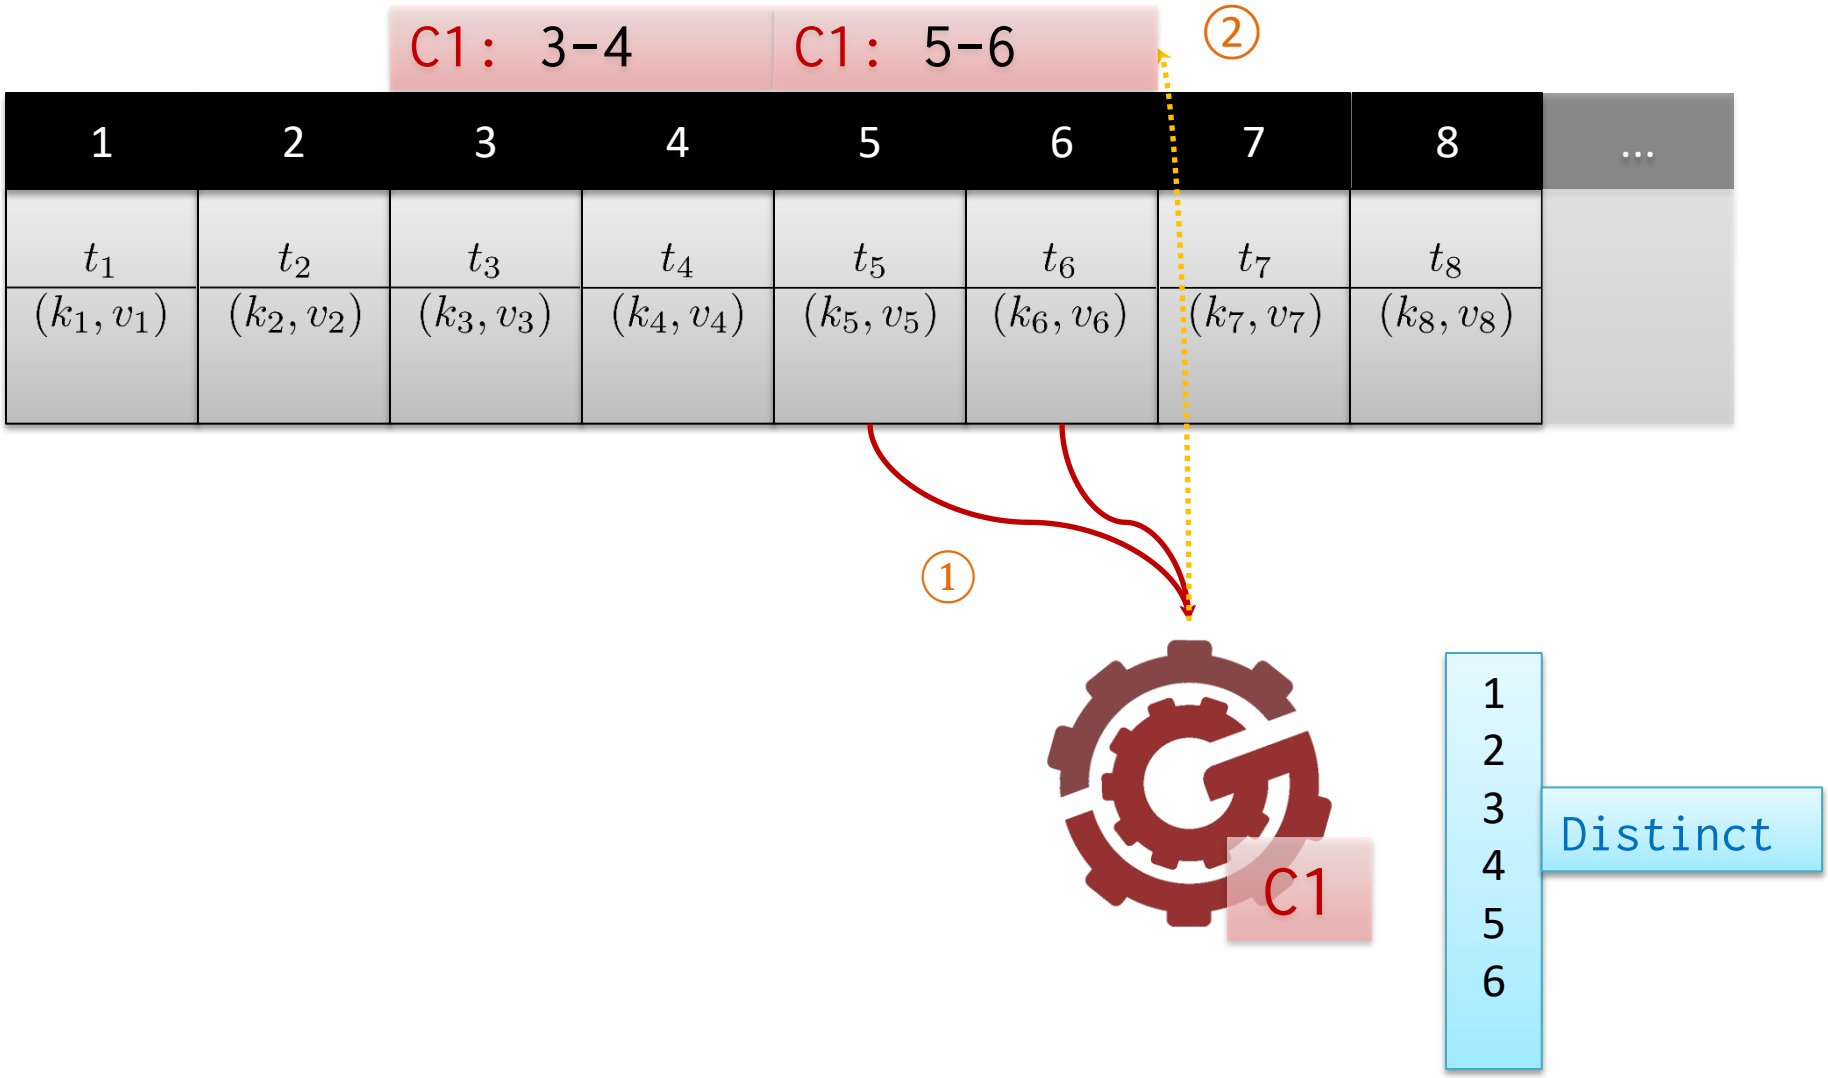

Read Guarantees

- At least once
- At most once
- Effectively once
- Exactly once
  - Data and offset updated as a single transaction

Read: Exactly Once

Transaction 1: ①, ②

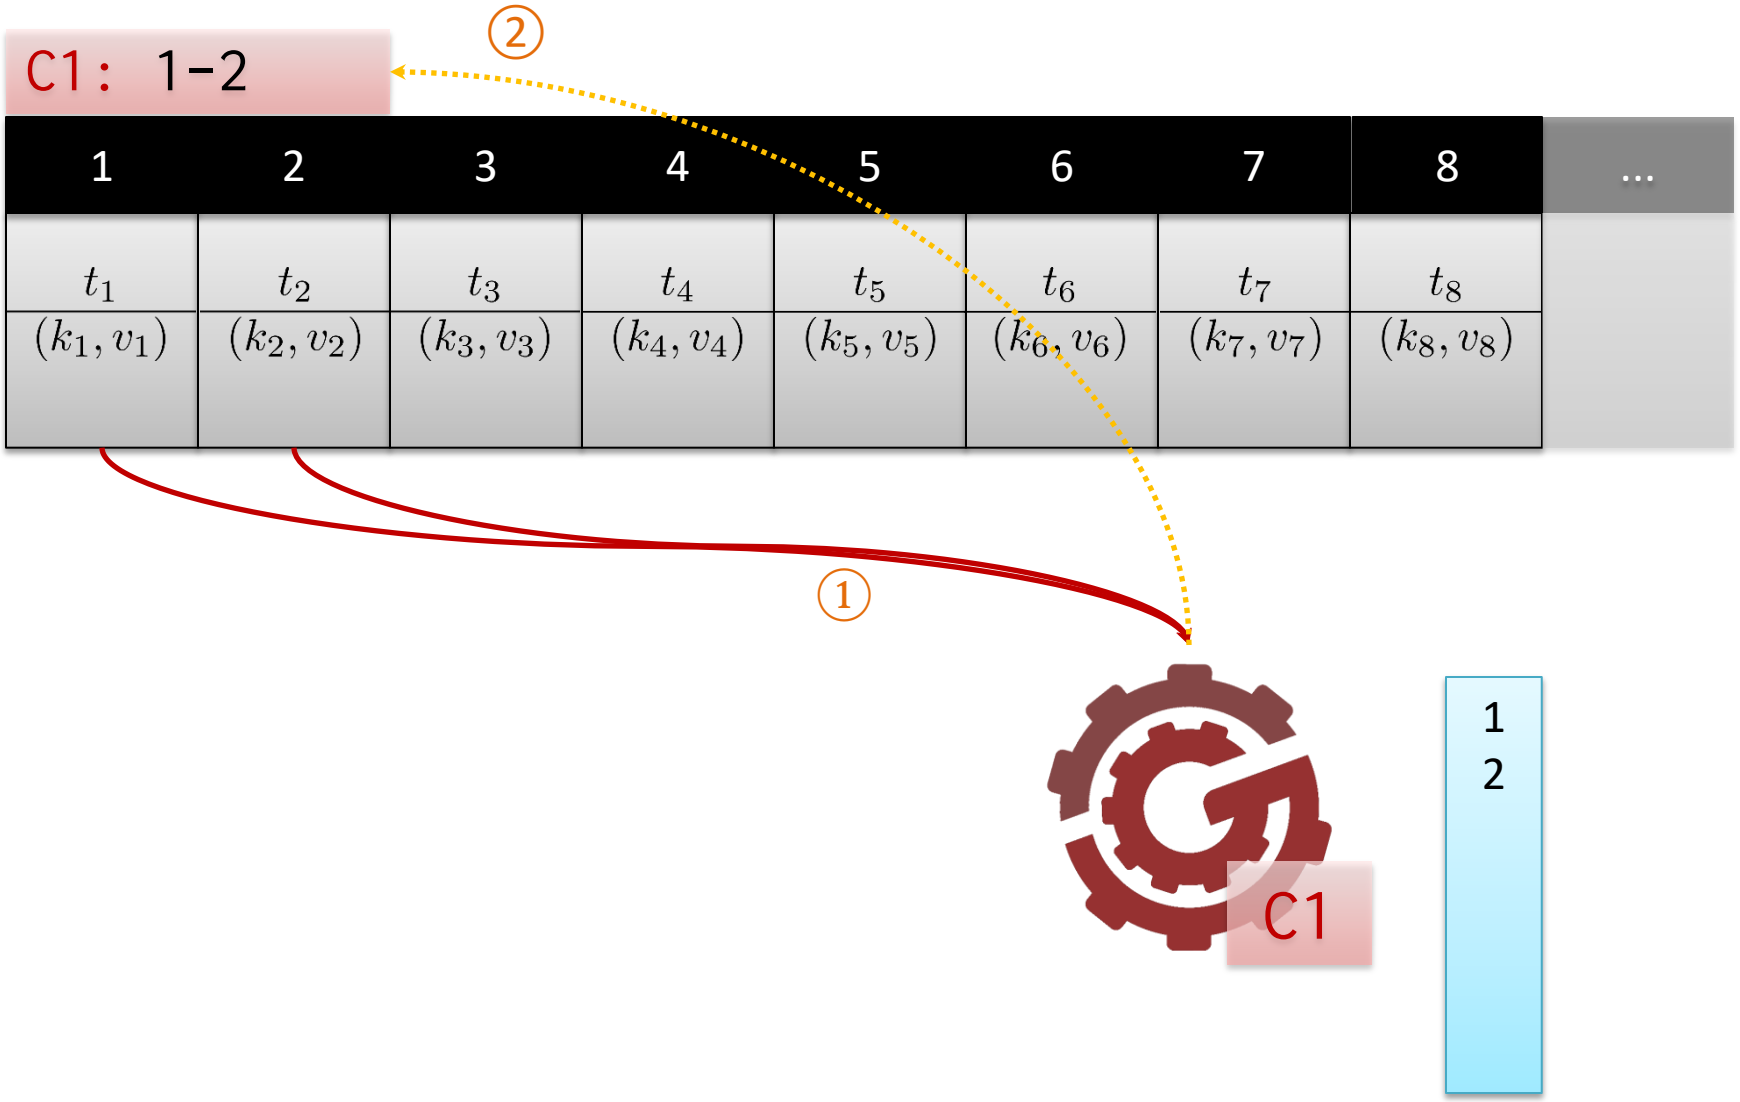

Read: Exactly Once

C1: 1-2

1	2	3	4	5	6	7	8	...
t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	
(k_1, v_1)	(k_2, v_2)	(k_3, v_3)	(k_4, v_4)	(k_5, v_5)	(k_6, v_6)	(k_7, v_7)	(k_8, v_8)	

1
2

Read: Exactly Once

Transaction 2: ①, ②

C1: 1-2

1	2	3	4	5	6	7	8	...
t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	
(k_1, v_1)	(k_2, v_2)	(k_3, v_3)	(k_4, v_4)	(k_5, v_5)	(k_6, v_6)	(k_7, v_7)	(k_8, v_8)	

C1

1
2

Read: Exactly Once

Transaction 3: ①, ②

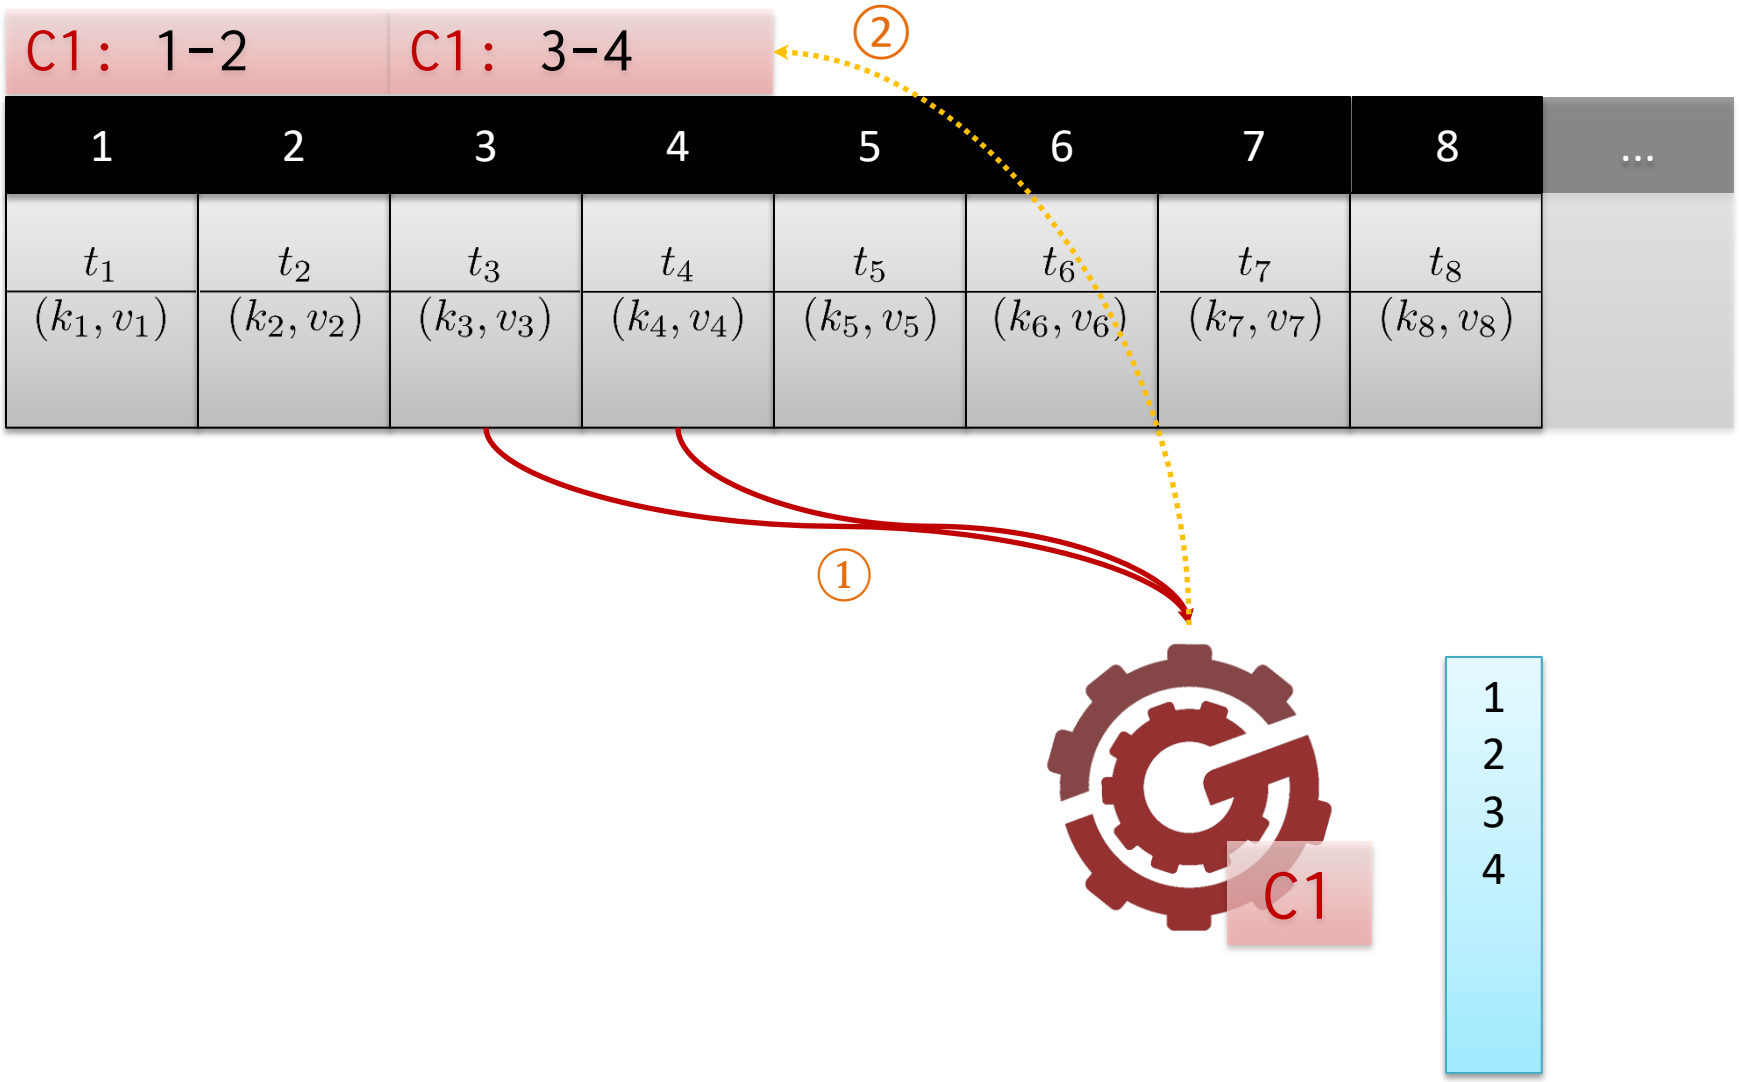

Read: Exactly Once

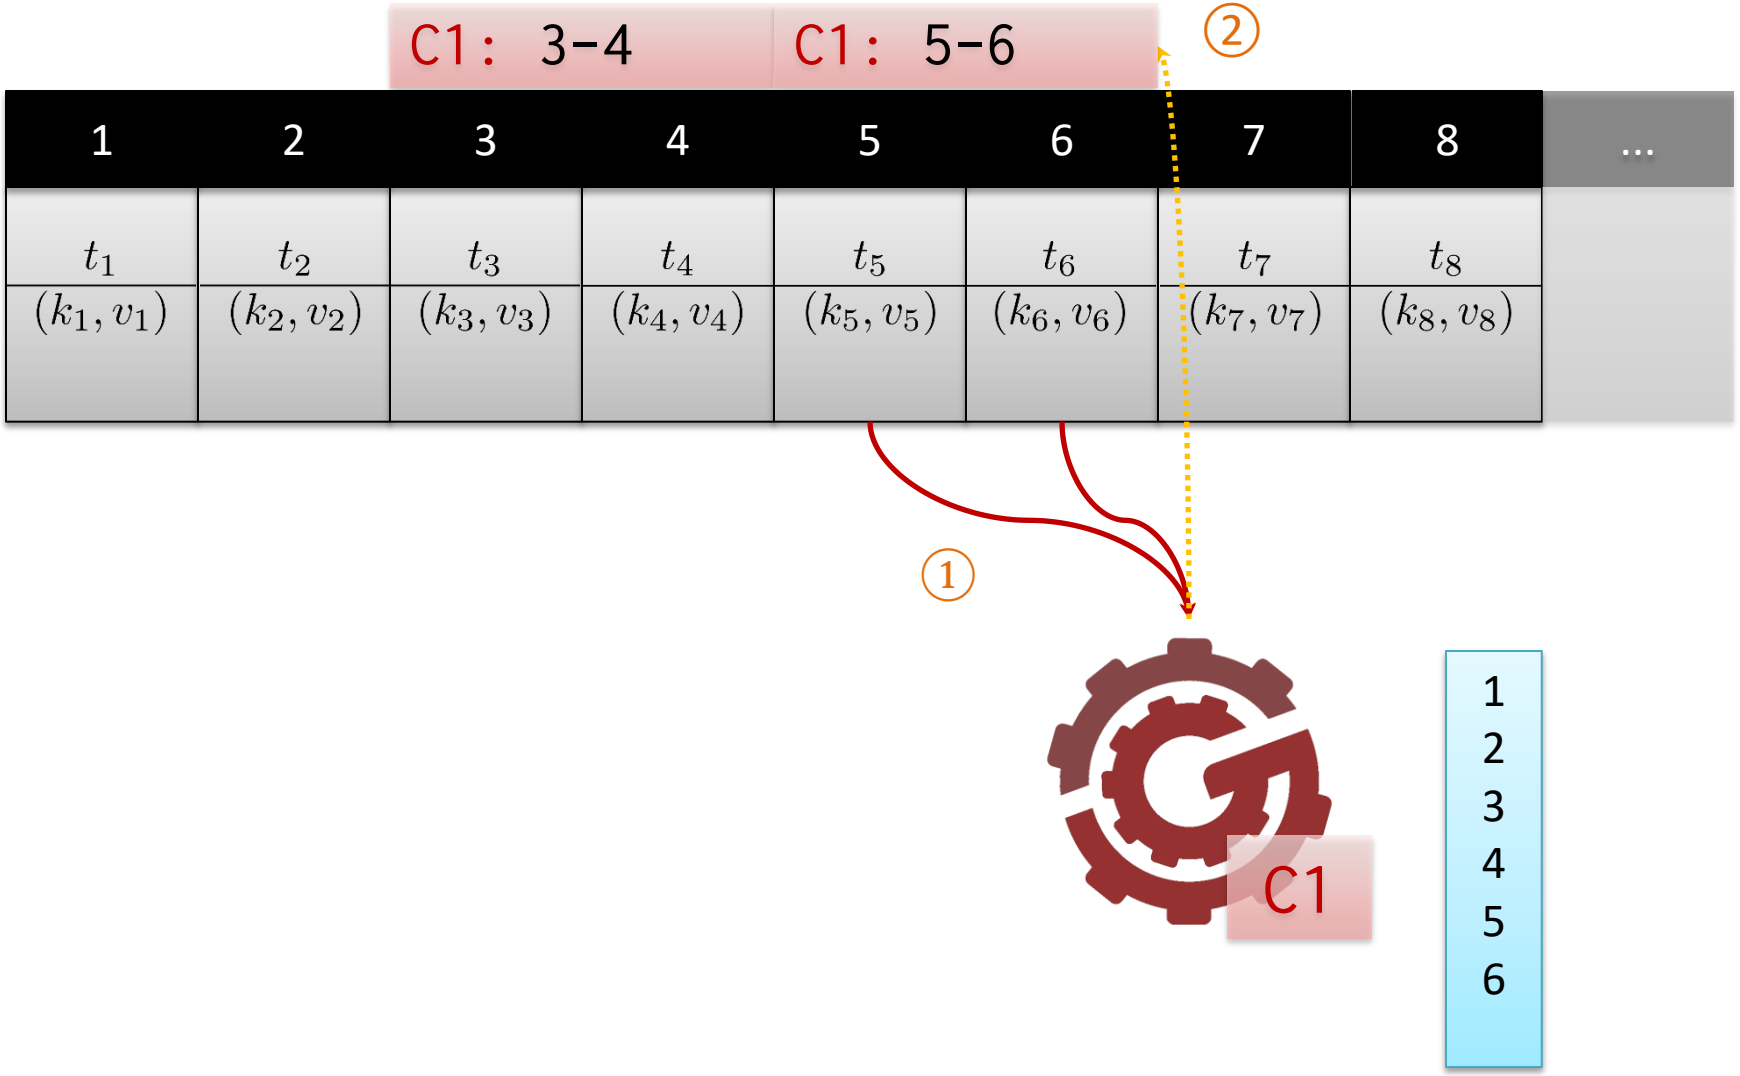

Leader Replication and Reads

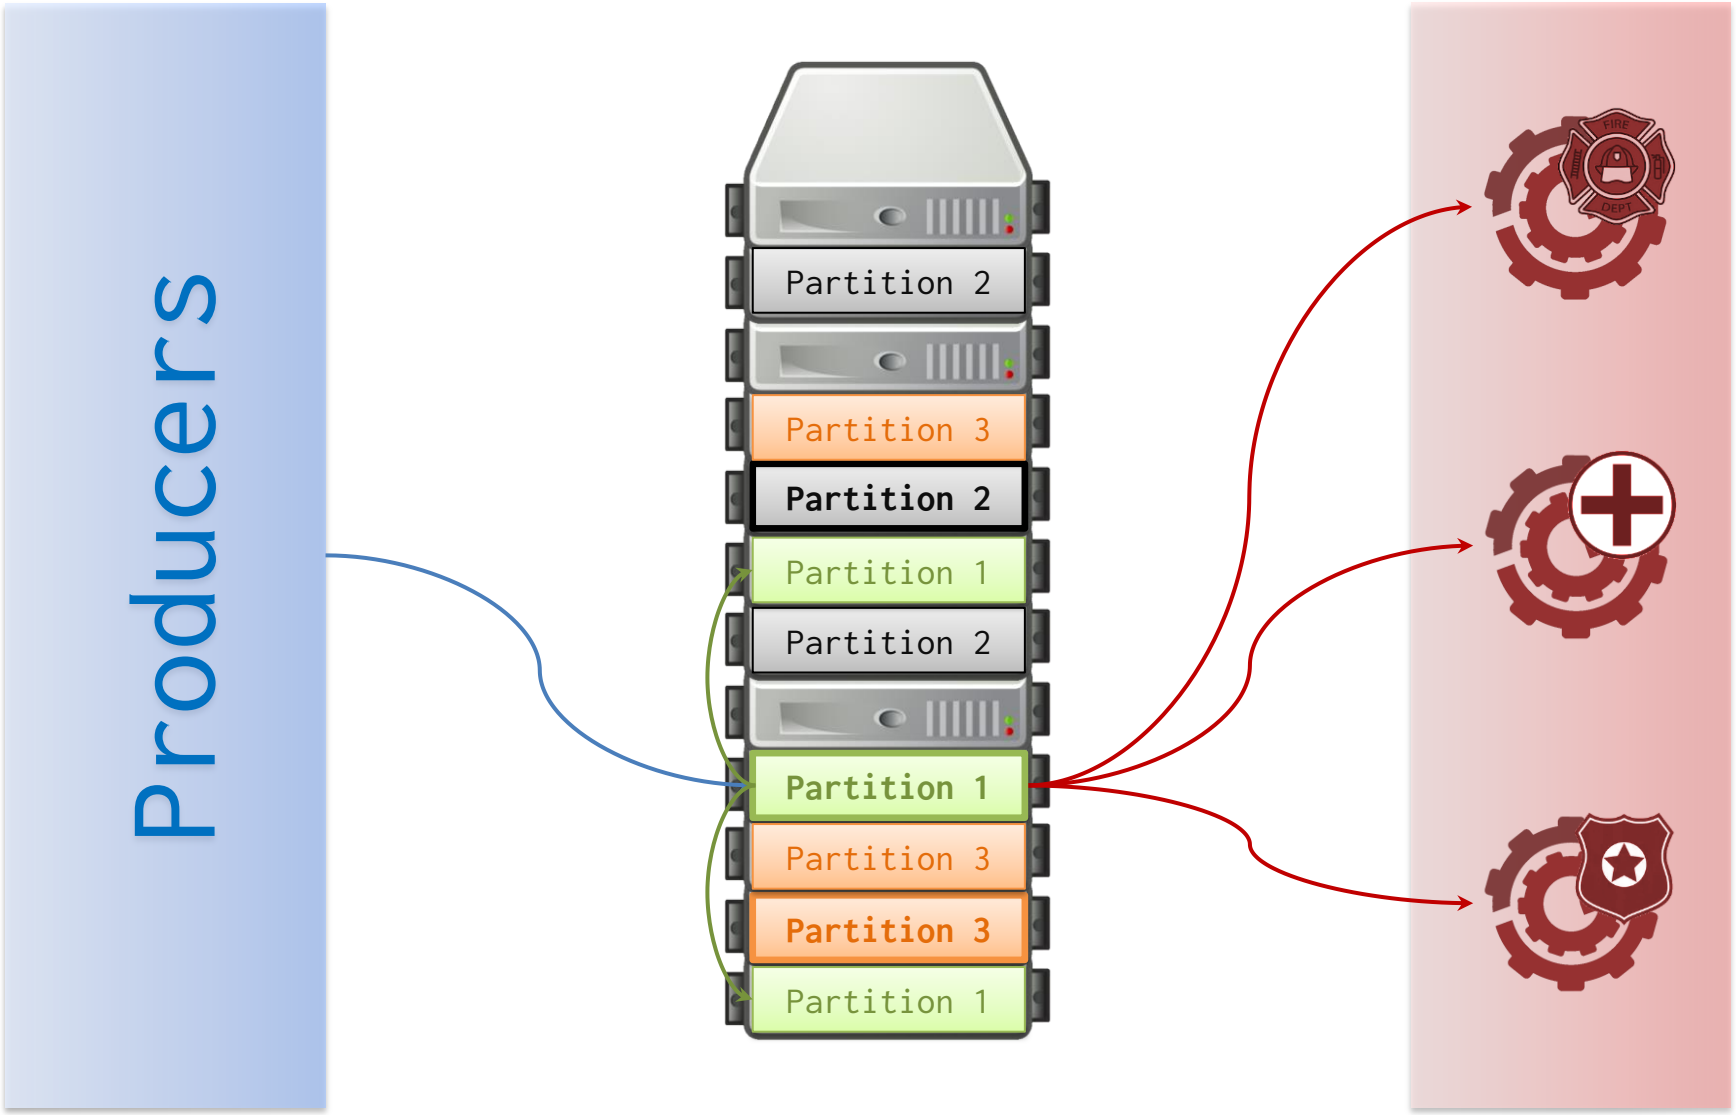

KAFKA: CONSUMER GROUPS

Consumer Groups

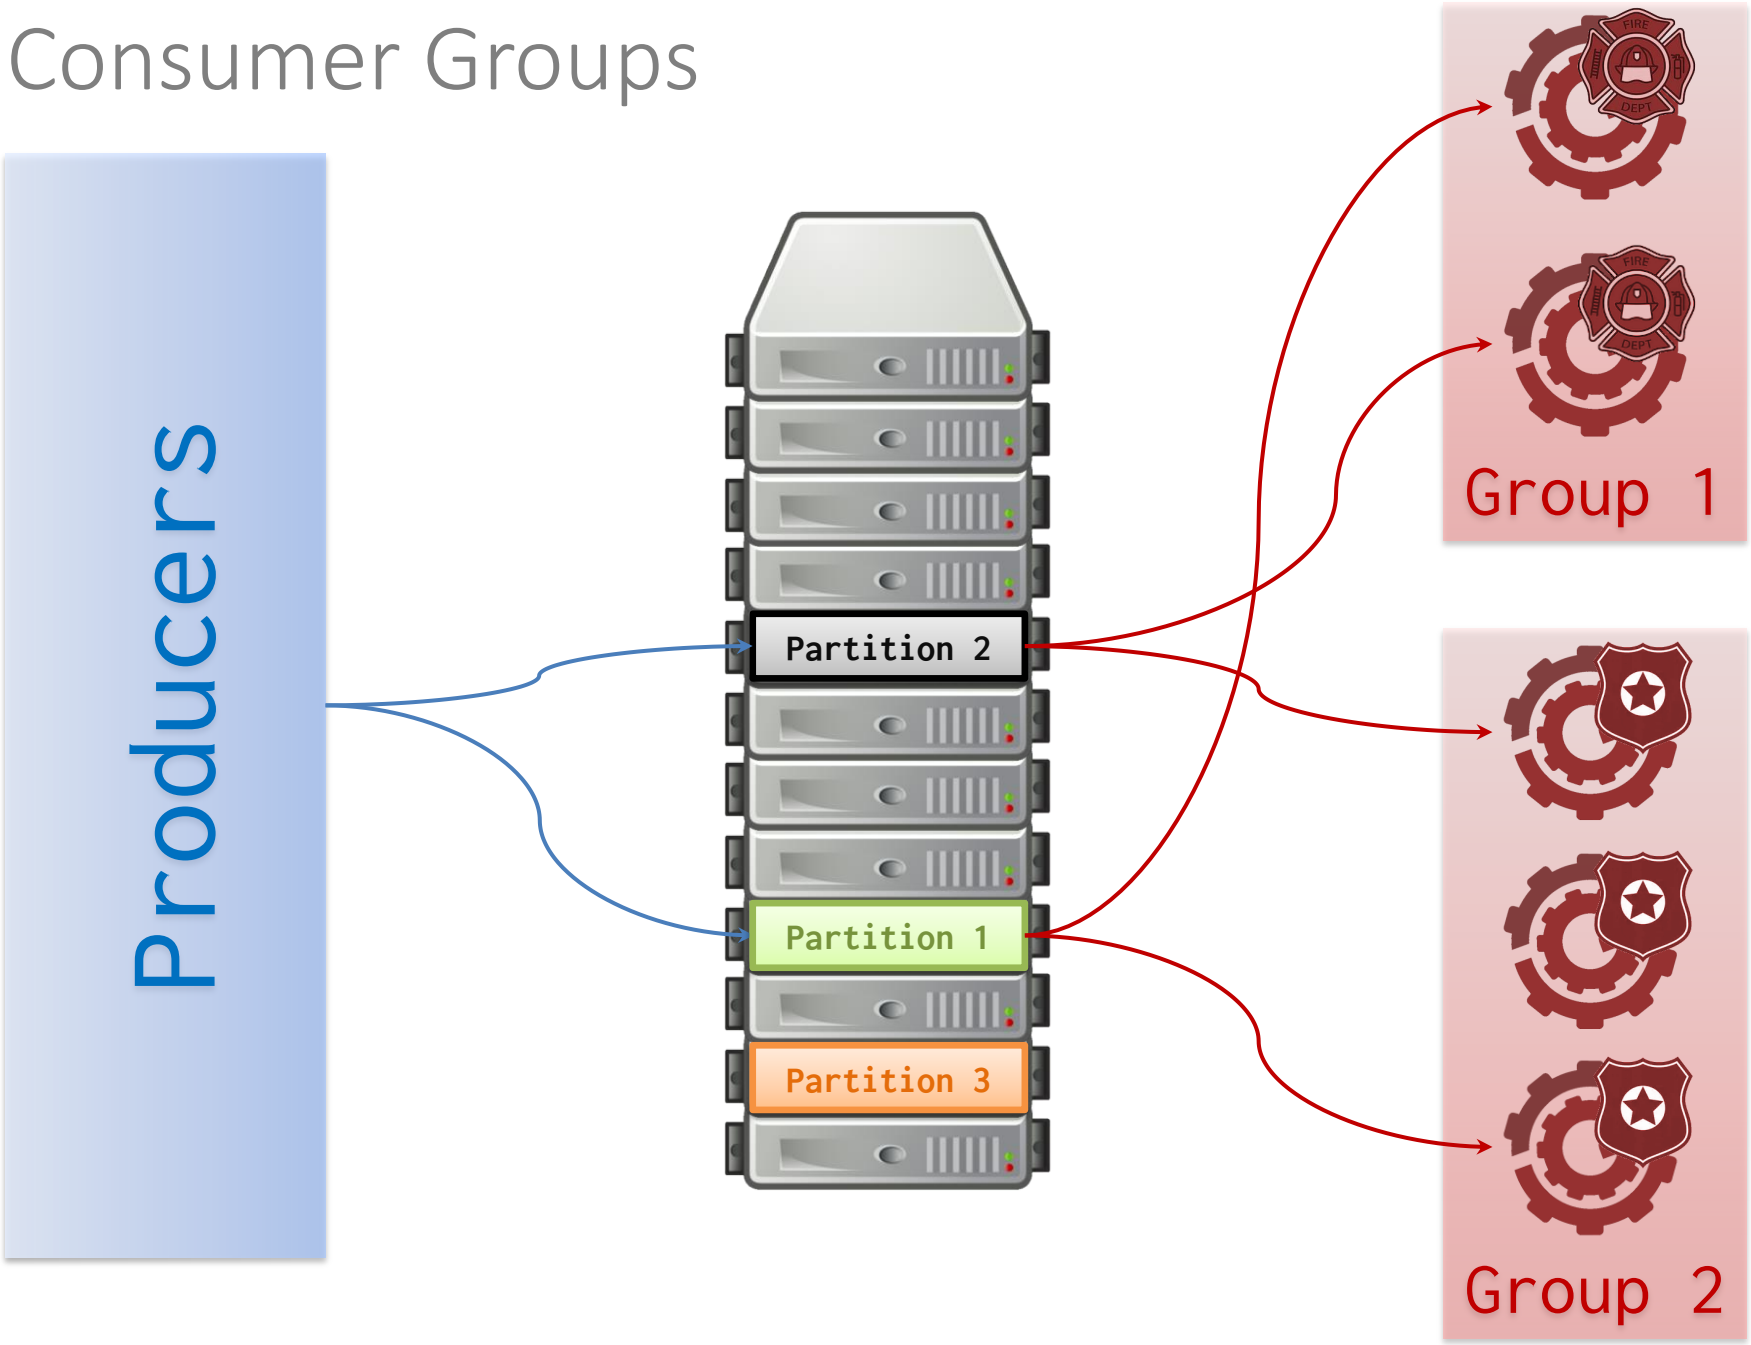

Consumer Groups

- Write to one consumer in each group
  - Allows to partition consumers
  - Load balancing within each group

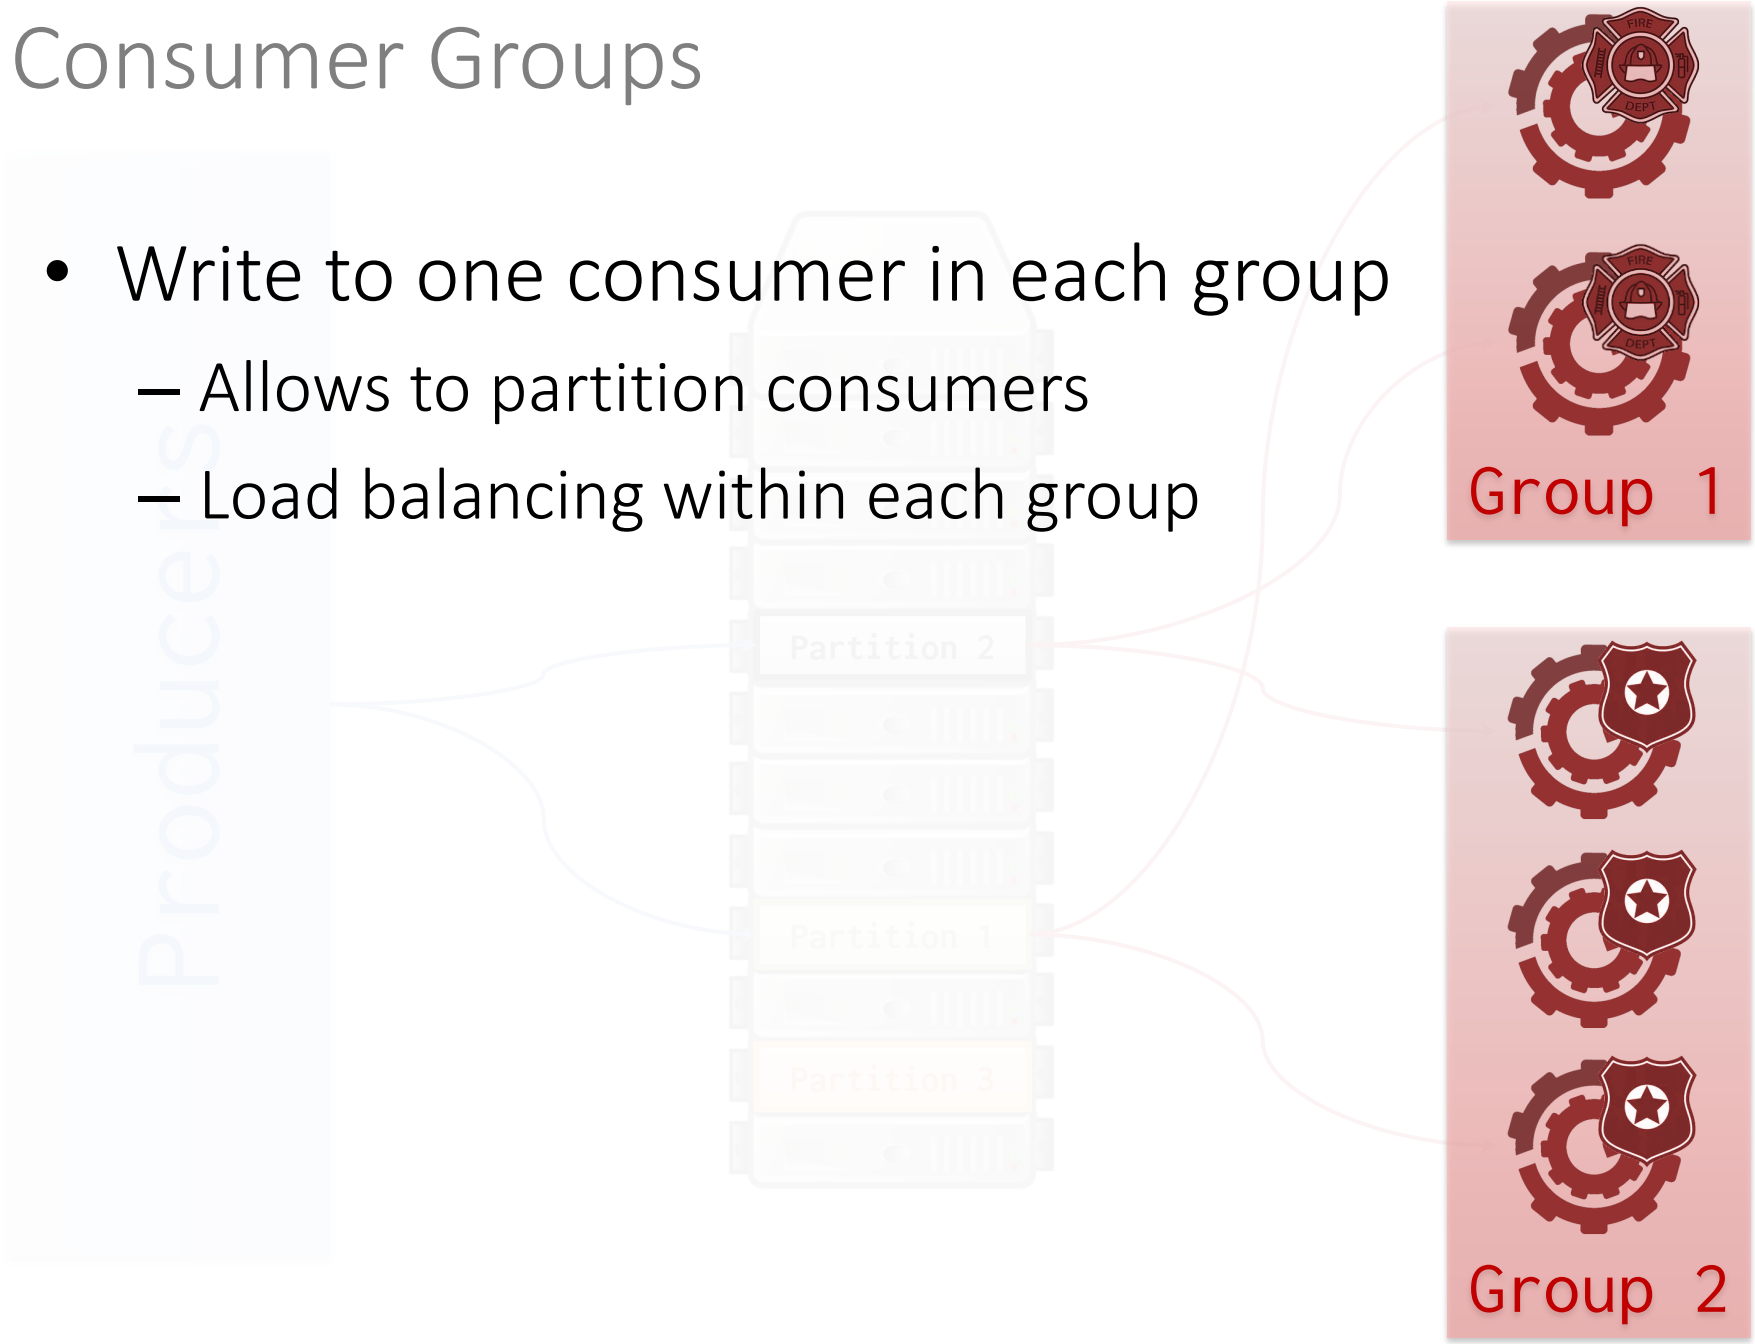

Kafka: Streams and Connectors

Kafka Overview

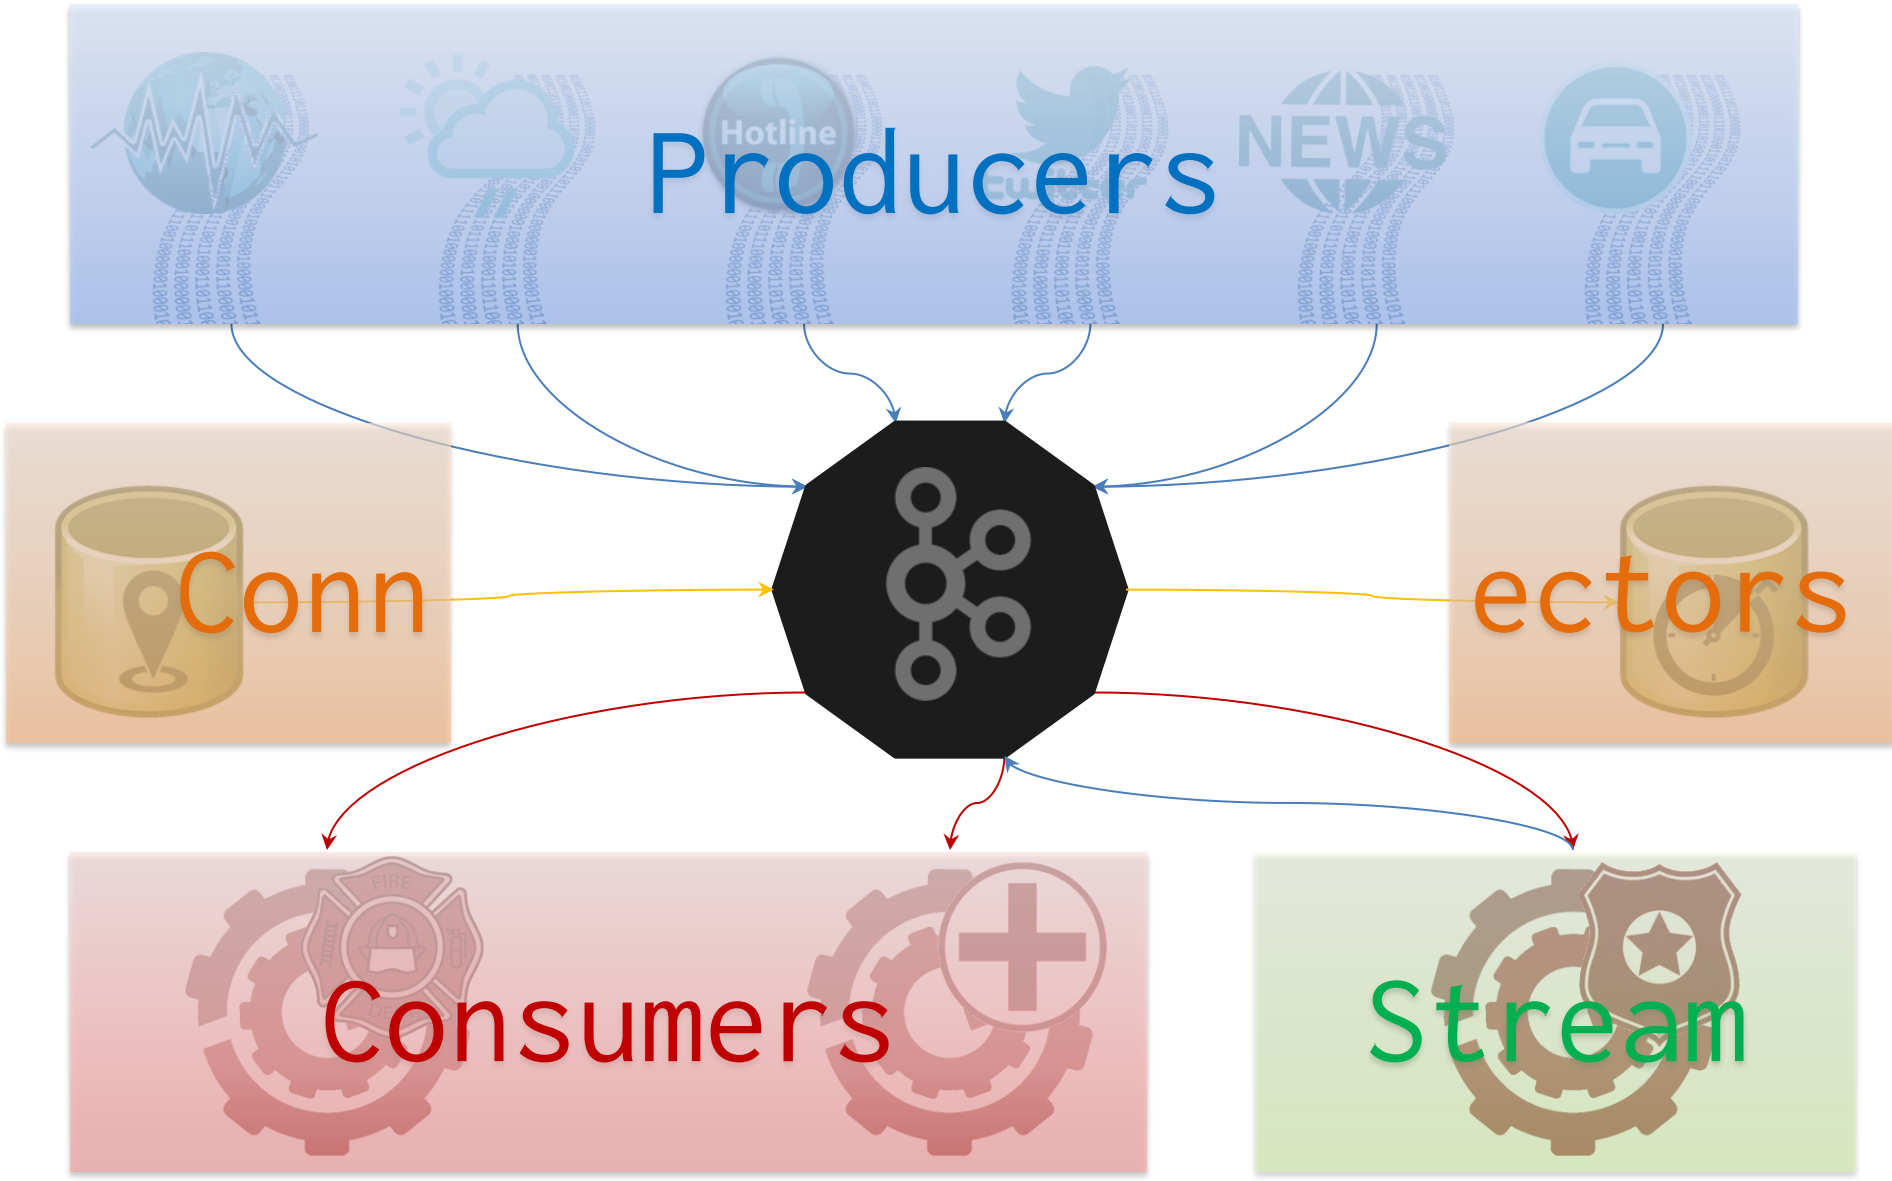

Kafka Overview

- Producer API:
  - Append records to topics (push)
- Consumer API:
  - Read records from topics (pull)
- Connector API:
  - Read/write to external components
    - For example, a database or other streaming platforms
- Stream API (Producer + Consumer):
  - Read records from input topics
  - Append records to output topics

Optimisations and Other Features

Kafka Optimisations

- Log Compaction
  - Repeated sequential values are suppressed
- Direct Disk-to-Network
  - When data don't need to be loaded into JVM
- Consumer / Producer Quotas
  - Set limits to avoid saturating the system
- ...

Kafka Streams API

- Aggregation (e.g., count messages)
- Joins (e.g., "unify" two streams)
- Windowing (define retention period)
- Continuous Querying (KSQL)

Available Frameworks

Questions?

Class Projects

Course Marking

- 80% for Weekly Labs
  - 11 labs total
  - 4 labs will be obligatory
  - 4 best out of the remaining 7 labs will count
- 20% for Class Project

Assignments each week

Working in groups

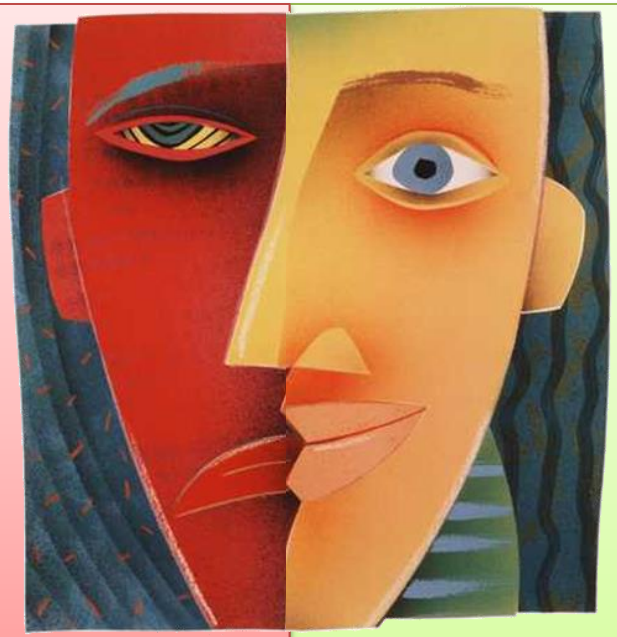

Hands-on each week!

Working in groups!

Class Project

- Done in threes
- Goal: Use what you've learned to do something cool/fun (hopefully)
- Process:
  - Form groups of three or four (in the forum, before April 30th)
  - On April 30th we will assign the rest automatically
  - Start thinking up topics / find interesting datasets!
  - Register topic
  - Work on projects during semester
- Deliverables: 4 minute presentation (video)
- Marked on: Difficulty, appropriateness, scale, good use of techniques, presentation, coolness, creativity, value
  - Ambition is appreciated, even if you don't succeed

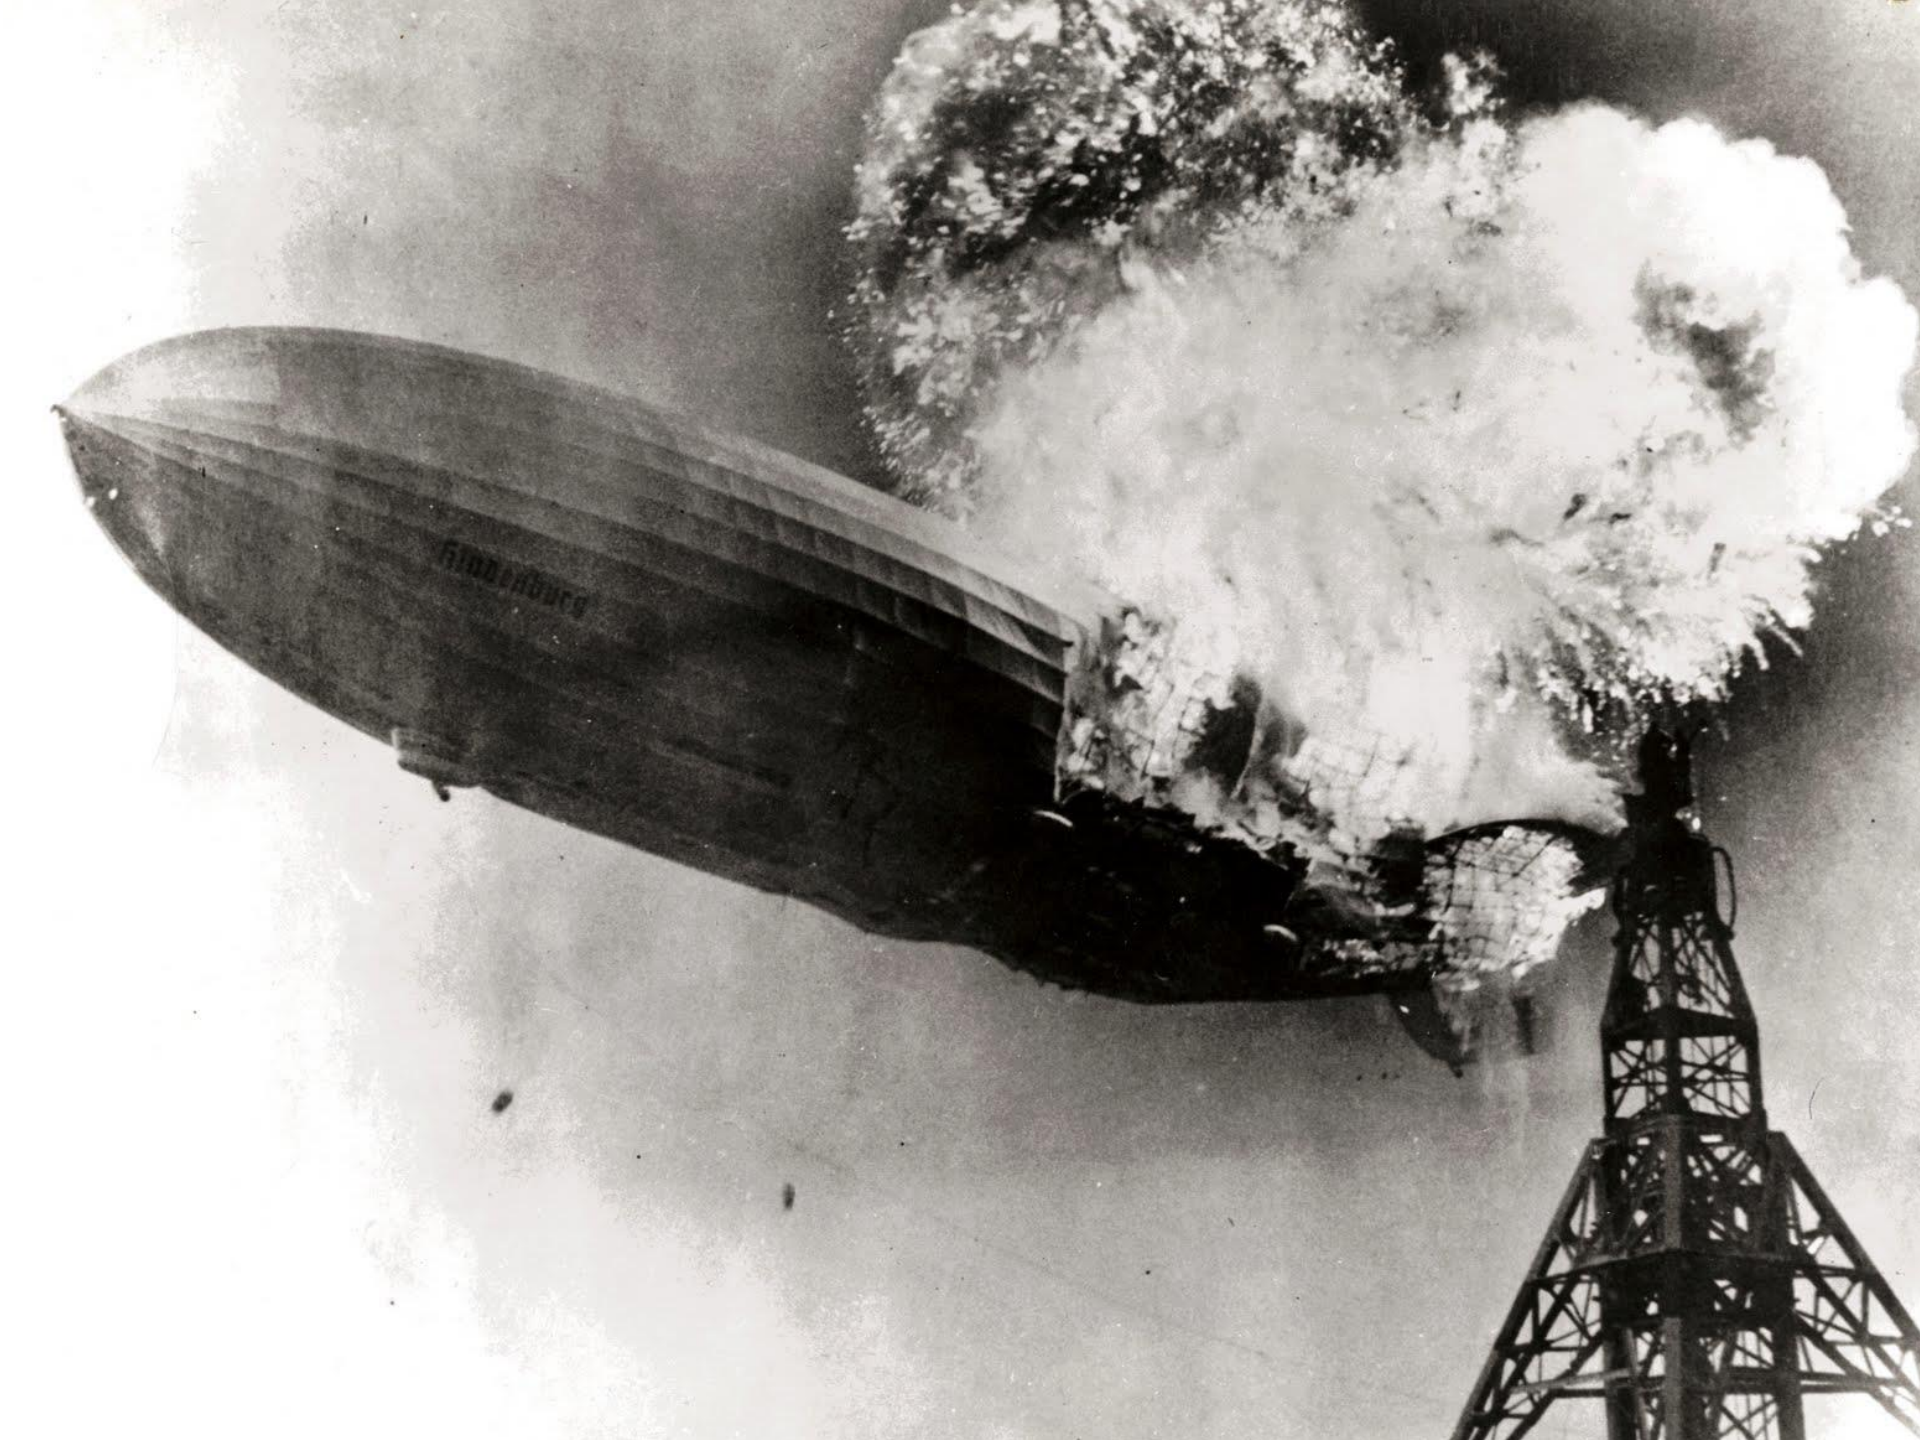
Hindenburg

Desiderata for project

- Must focus around some technique from the course!
- Expected difficulty: similar to a lab, but without any instructions
- Data not too small:
  - Should have >250,000 tuples/entries
- Data not too large:
  - Should have <1,000,000,000 tuples/entries
  - If very large, perhaps take a sample?
- In case of COVID-19 data, we can make exceptions

Where to find/explore data?

- Kaggle:
  - https://www.kaggle.com/
- Google Dataset Search:
  - https://datasetsearch.research.google.com/
- Datos Abiertos de Chile:
  - https://datos.gob.cl/
  - https://es.datachile.io/
- …