

# Containerizing Distributed Pipes

Hagen Tönnies

[www.linkedin.com/in/hagen-toennies](http://www.linkedin.com/in/hagen-toennies)

Special Thanks to @CQnib

# Agenda

- **Background**
- **Distributed Tools**
- **Containerizing**
- **Recap**

# Background

# BUZZWORDS

🌀 Kafka, **Samza** and the

Unix philosophy

of distributed data

Martin Kleppmann @martinkl

# PIPES

# Unix Pipe Recap

- [...]In Unix-like computer operating systems, **a pipeline** is a **sequence** of **processes chained together** by their standard streams, so that the **output** of each **process** (stdout) feeds directly as **input** (stdin) **to the next one**.

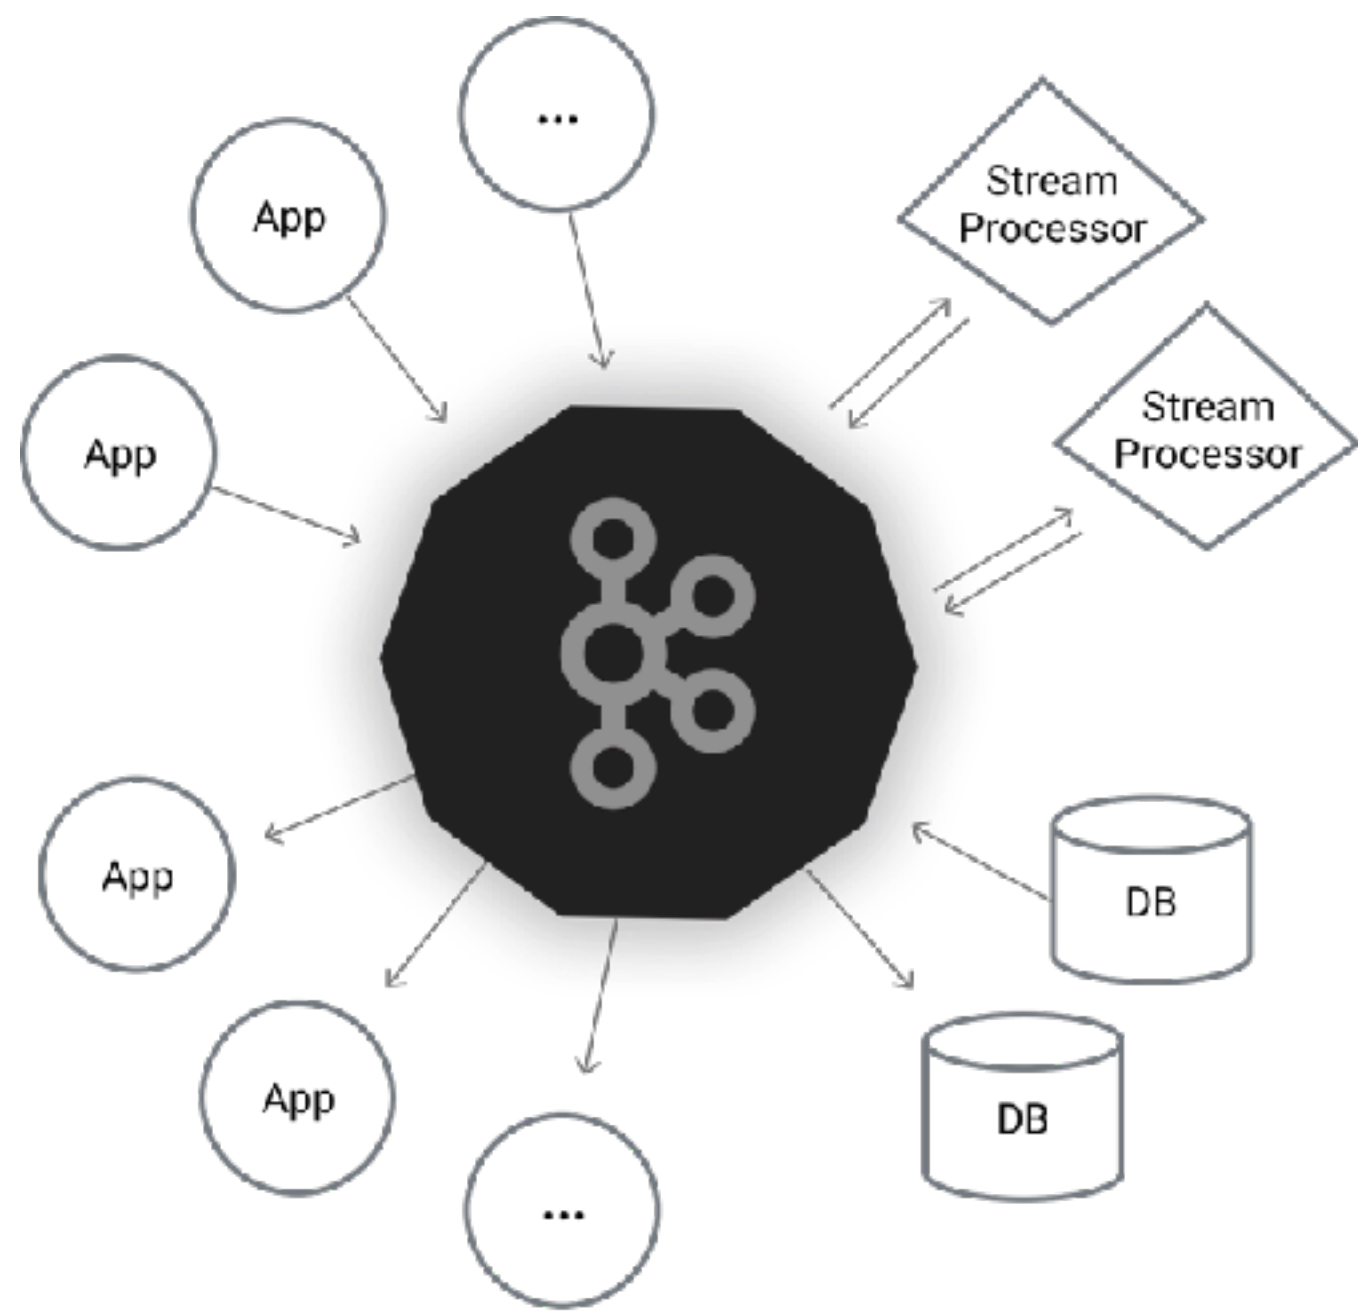
# Unix Philosophy

- Write **programs** that **do one thing** and **do it well**.
- Write **programs** to **work together**.
- Write **programs to handle text streams**, because that **is a universal interface**.

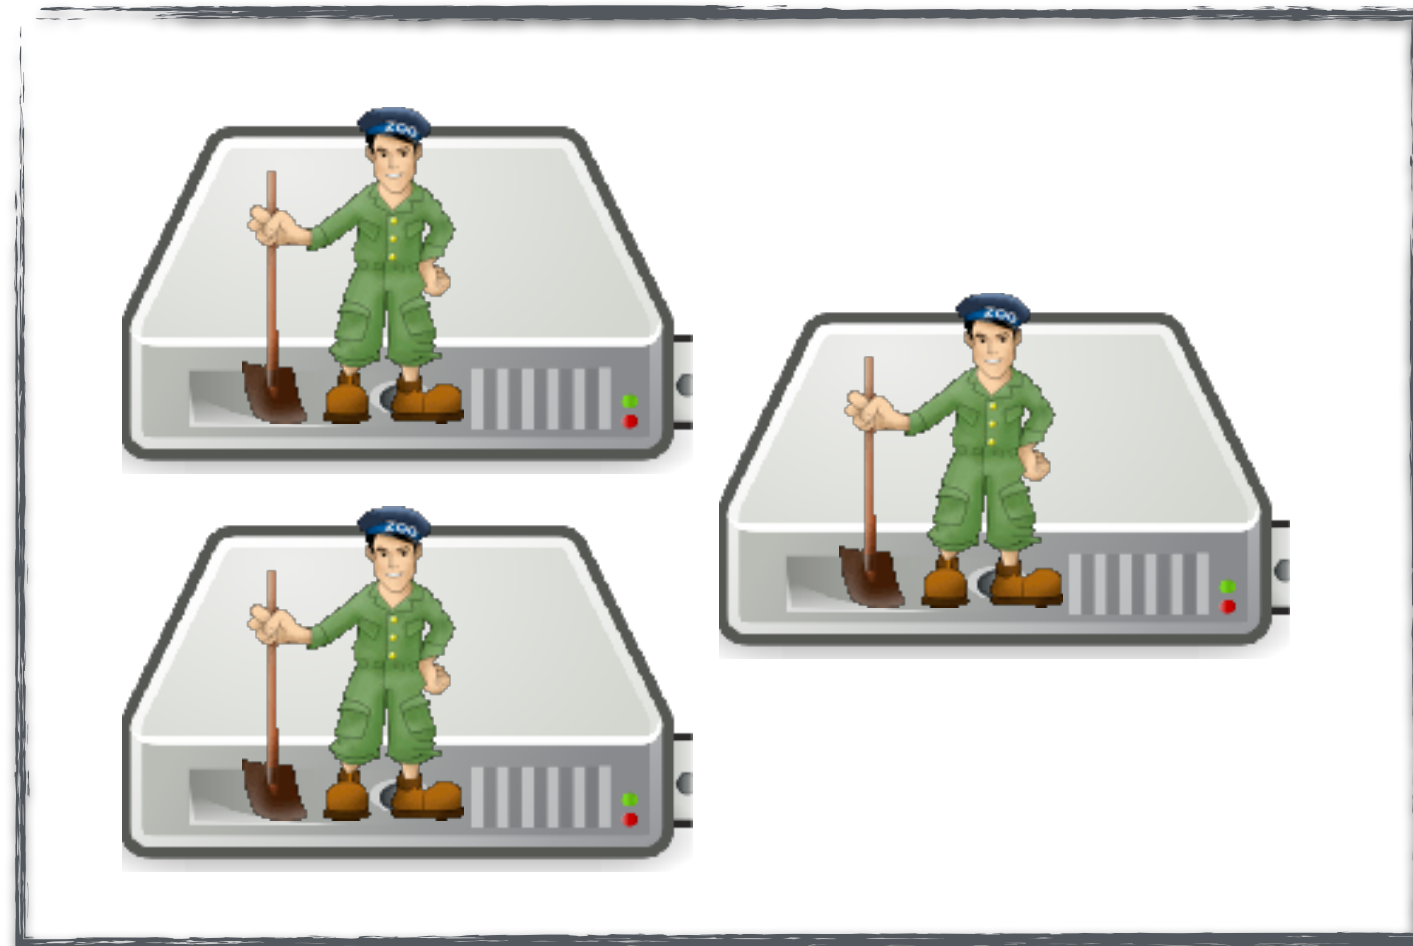


**KAFKA**

# Background




# Apache Kafka Cluster




Zookeeper


Kafka Broker




# Apache Kafka Topic




# Apache Kafka Partitions




# Apache Kafka Replications



# Apache Kafka Distributed Partitions




# Apache Kafka Producer

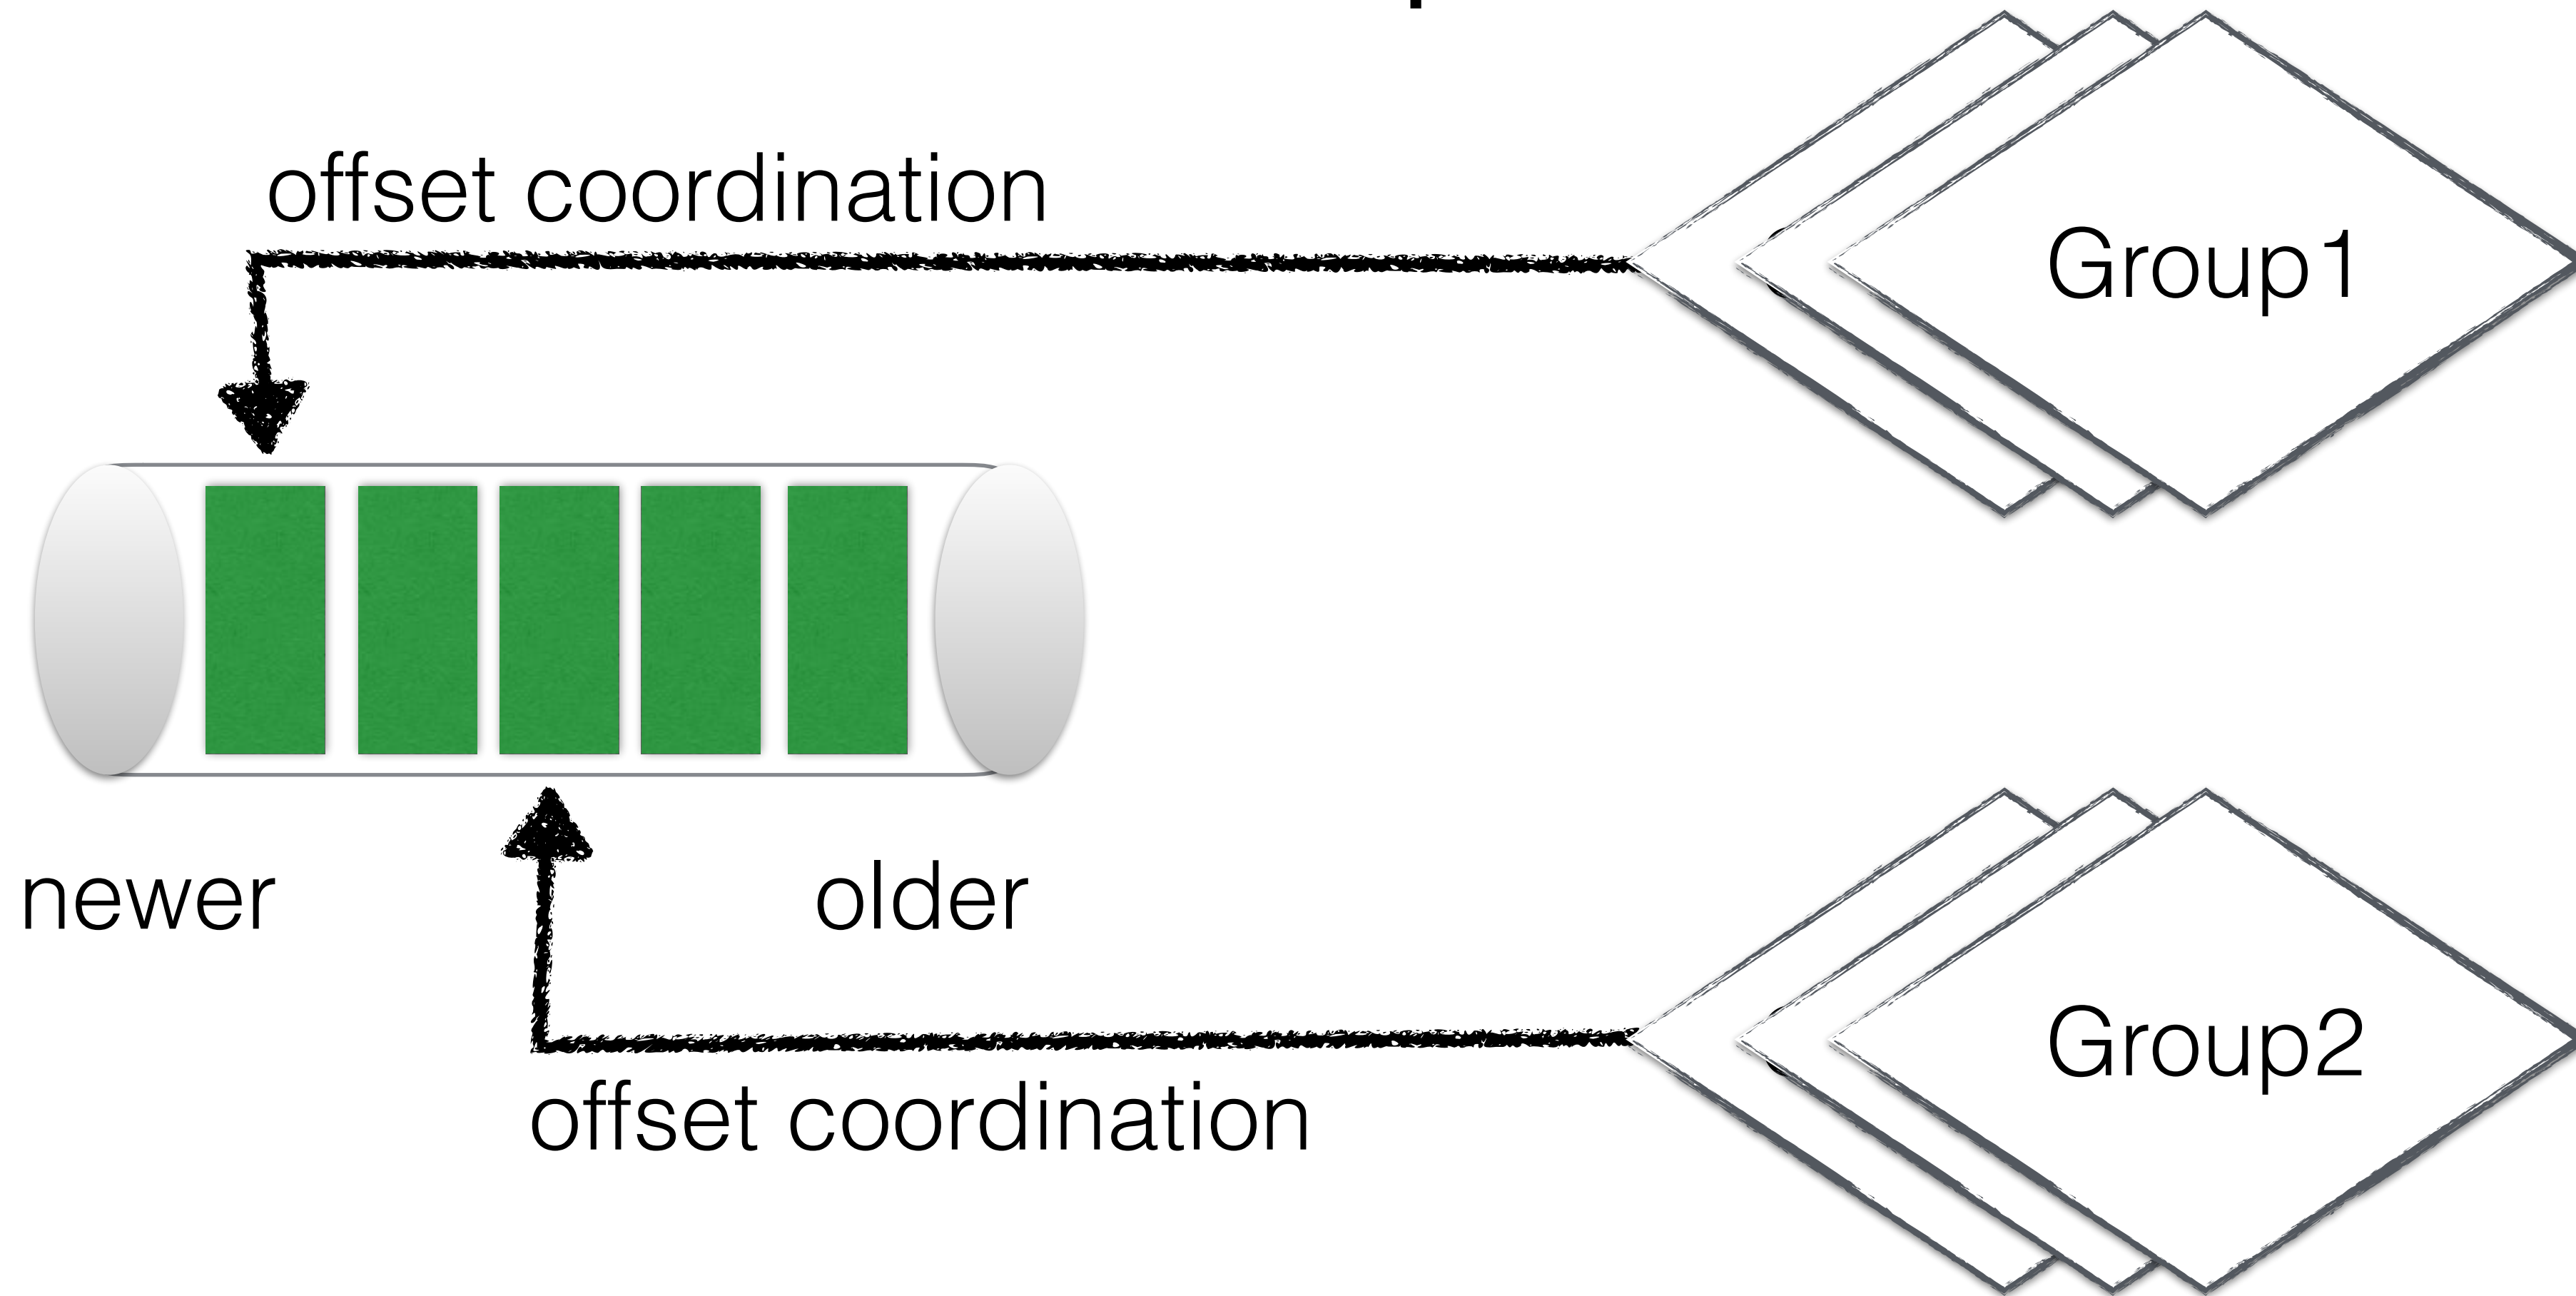




# Apache Kafka Consumer



# Apache Kafka Consumer Groups



# Apache Kafka Streams

Low-Level-API	High-Level-API
---------------	----------------

Topology Builder

Stream and Table

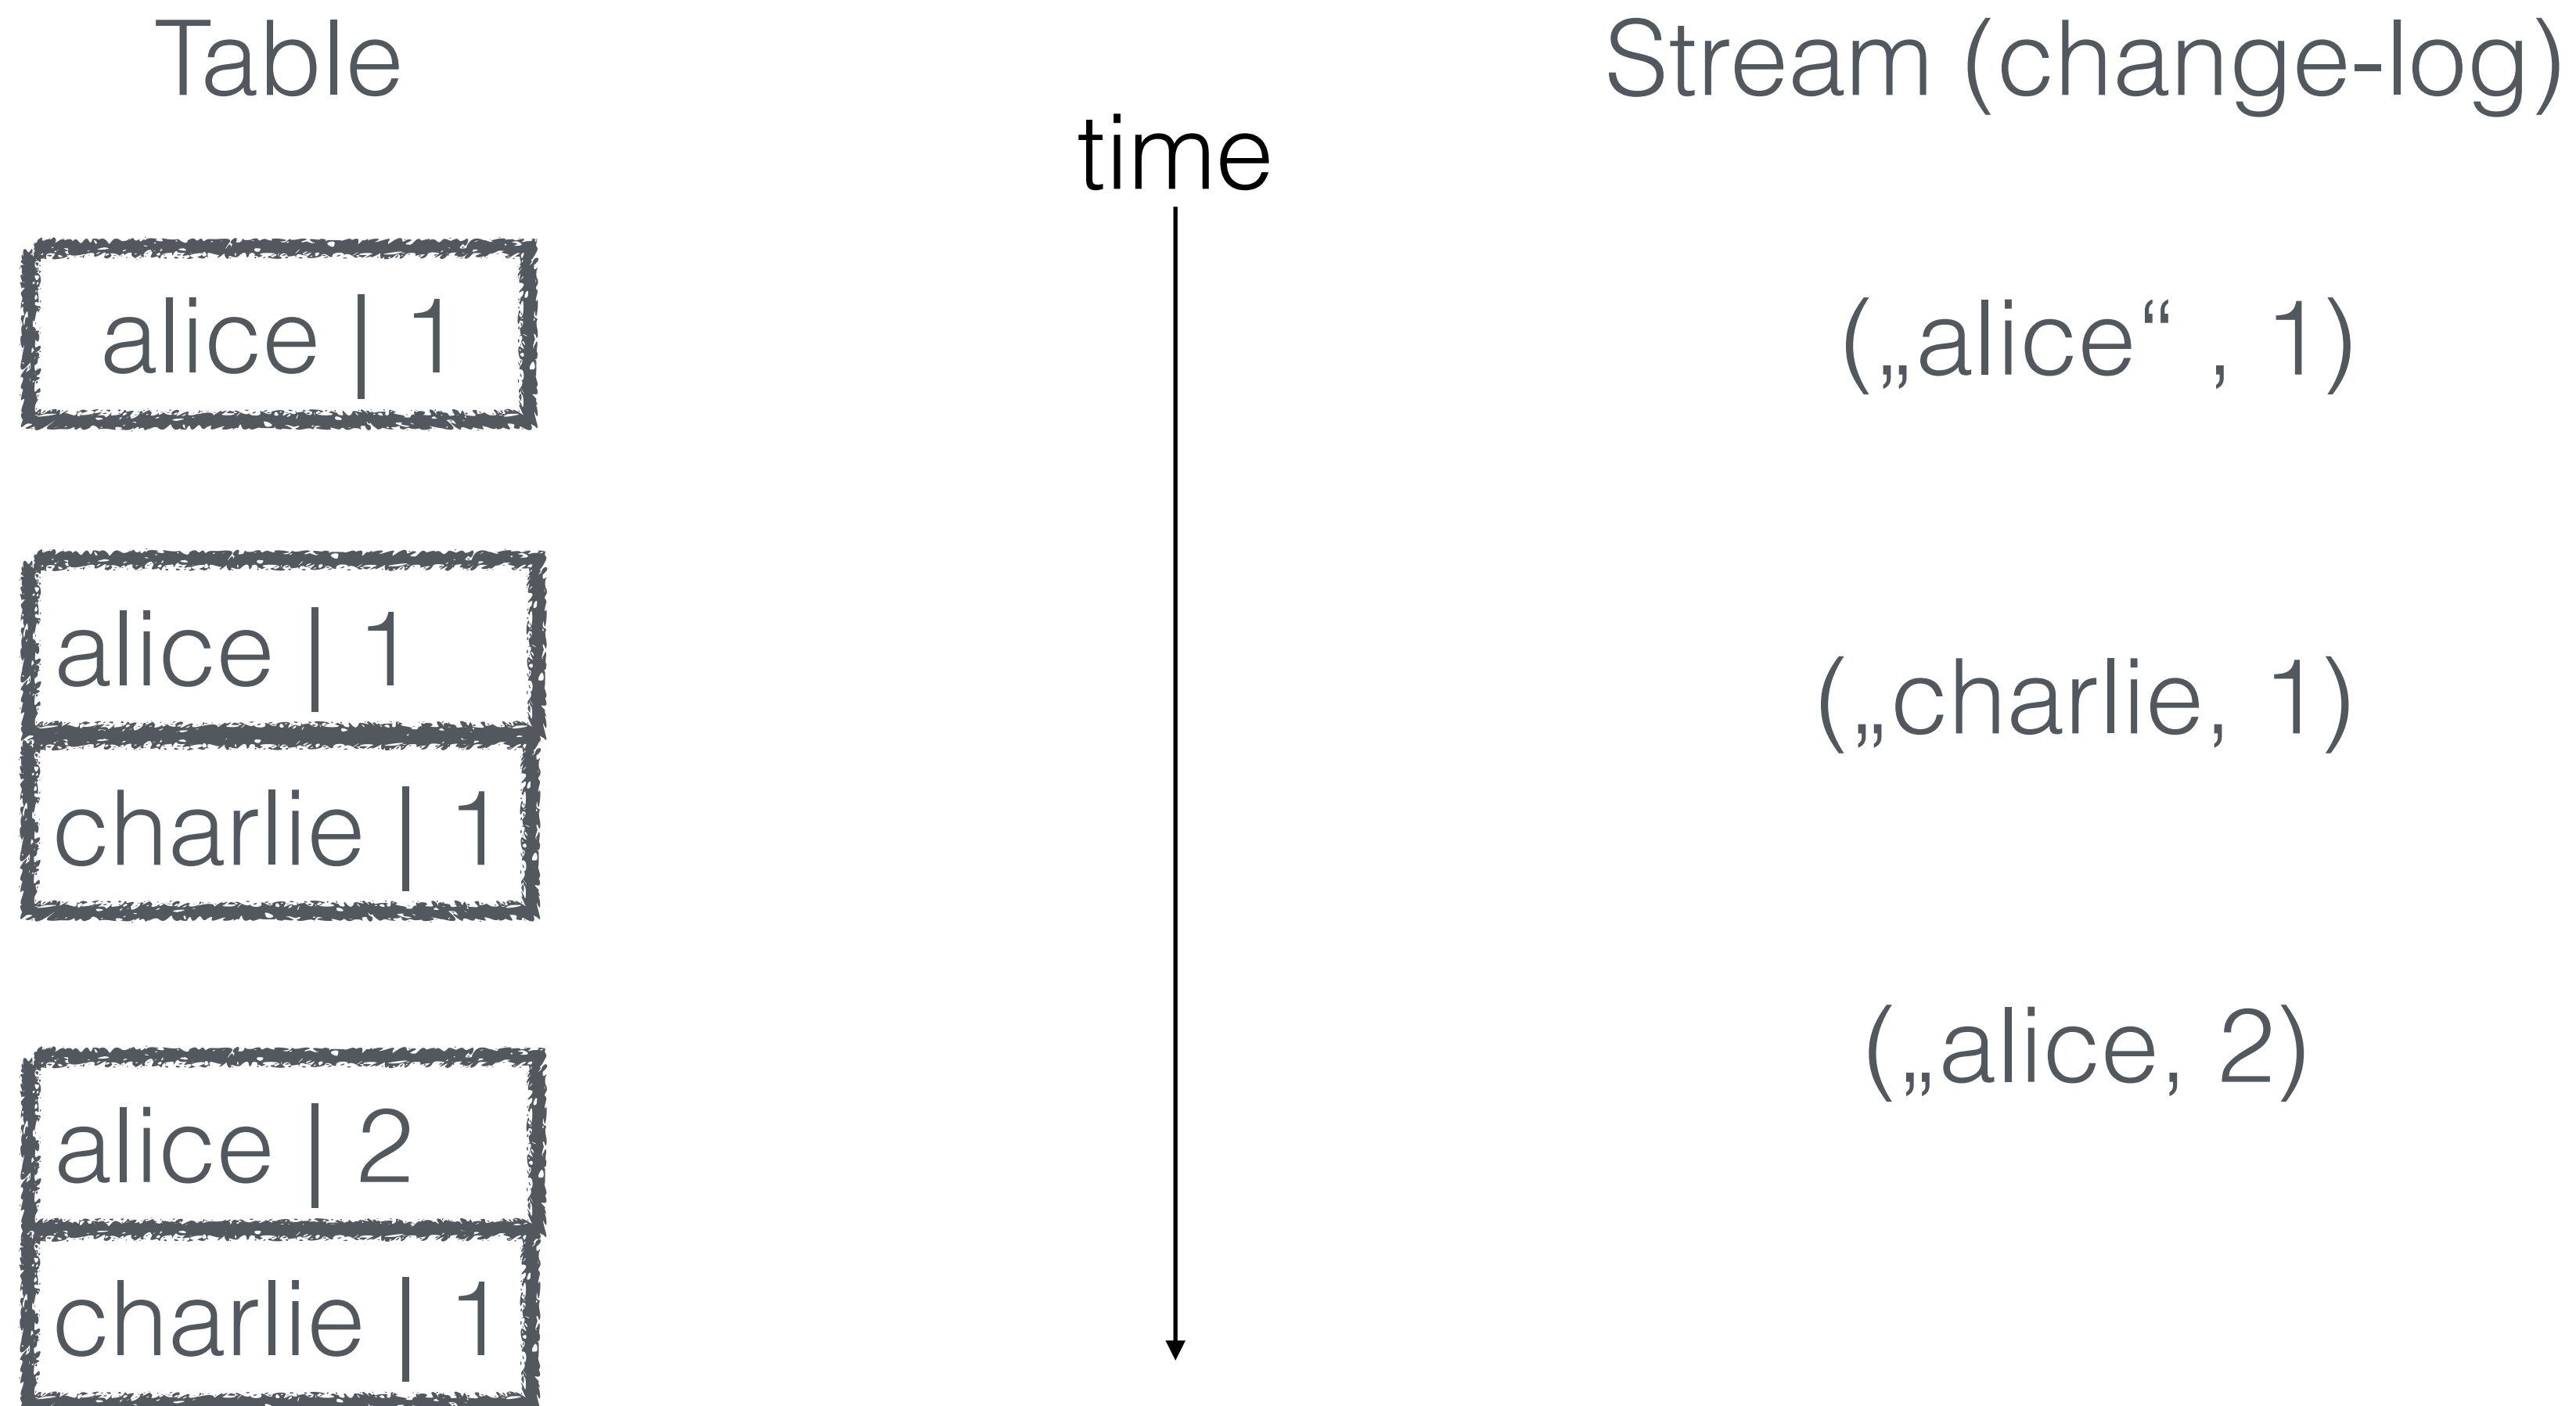
Custom Aggregators

Simple Transformation

Custom Processors

Simple Joins

# Kafka Streaming API



# STREAM PROCESSING

# Background

Capability



Simplicity

# Background



# DISTRIBUTED PIPE'ING



# Kafka featuring Unix

|

(split %1 "**s**")

# Kafka featuring Unix



(split %1 "\s")



Message Broker




Stream processing job

# Kafka featuring Unix

(split %1 **"\s"**) | (sketch %1) | (agg-by-key %1) | (store %1)

# Distributed Tools

# Distributed Tools



<https://xkcd.com/297/>

**CUTTER**

# Distributed Tools

## **Input:**

Edmilson Alves 0 Edmilson Alves -LRB- born February 17 , 1976 -RRB- , is a Brazilian midfielder who currently plays for Roasso Kumamoto in the J. League Division 2 .

## **Output:**

[ Edmilson, Alves, 0, Edmilson, Alves, LRB, born ...]

# Distributed Tools

```
41 (defn split-string-value-of-dict [data selector]
42   "Select a value given by the selector path, and split the string at the white space."
43   (try
44     (let [field-value (get-in data (into [] (first selector)))]
45       (map #(clojure.string/trim %1)
46            (clojure.string/split field-value #"\s"))))
47     (catch Exception e
48       (error "Failed parsing field: " selector e)
49       (list))))
```



# Distributed Tools

```
51 (defn- stream-mapper
52   "Main stream processor takes a configuration
53   [conf ]
54   (let [streamBuilder (KStreamBuilder.)
55         ^KStream log-stream (.stream
56                               streamBuilder
57                               stringSerde
58                               stringSerde
59                               (into-array String [(:input-topic conf)]))]
60     (-> log-stream
61       (.flatMapValues (reify ValueMapper
62                        (apply [this value]
63                              (try
64                                (let [value-as-dict (json/read-str value :key-fn keyword)]
65                                  (split-string-value-of-dict value-as-dict (:selector conf) ))
66                                (catch Exception e
67                                  (error "Failed parsing .json" e)
68                                  (list))))))
69       (.map (reify KeyValueMapper
70            (apply [this k v]
71                  (KeyValue. v v))))
72       (.through stringSerde stringSerde (:output-topic conf)))
73     (.start (KafkaStreams. streamBuilder (get-props conf))))))
```


# HEAVY HITTER

# clj-kstream-hh

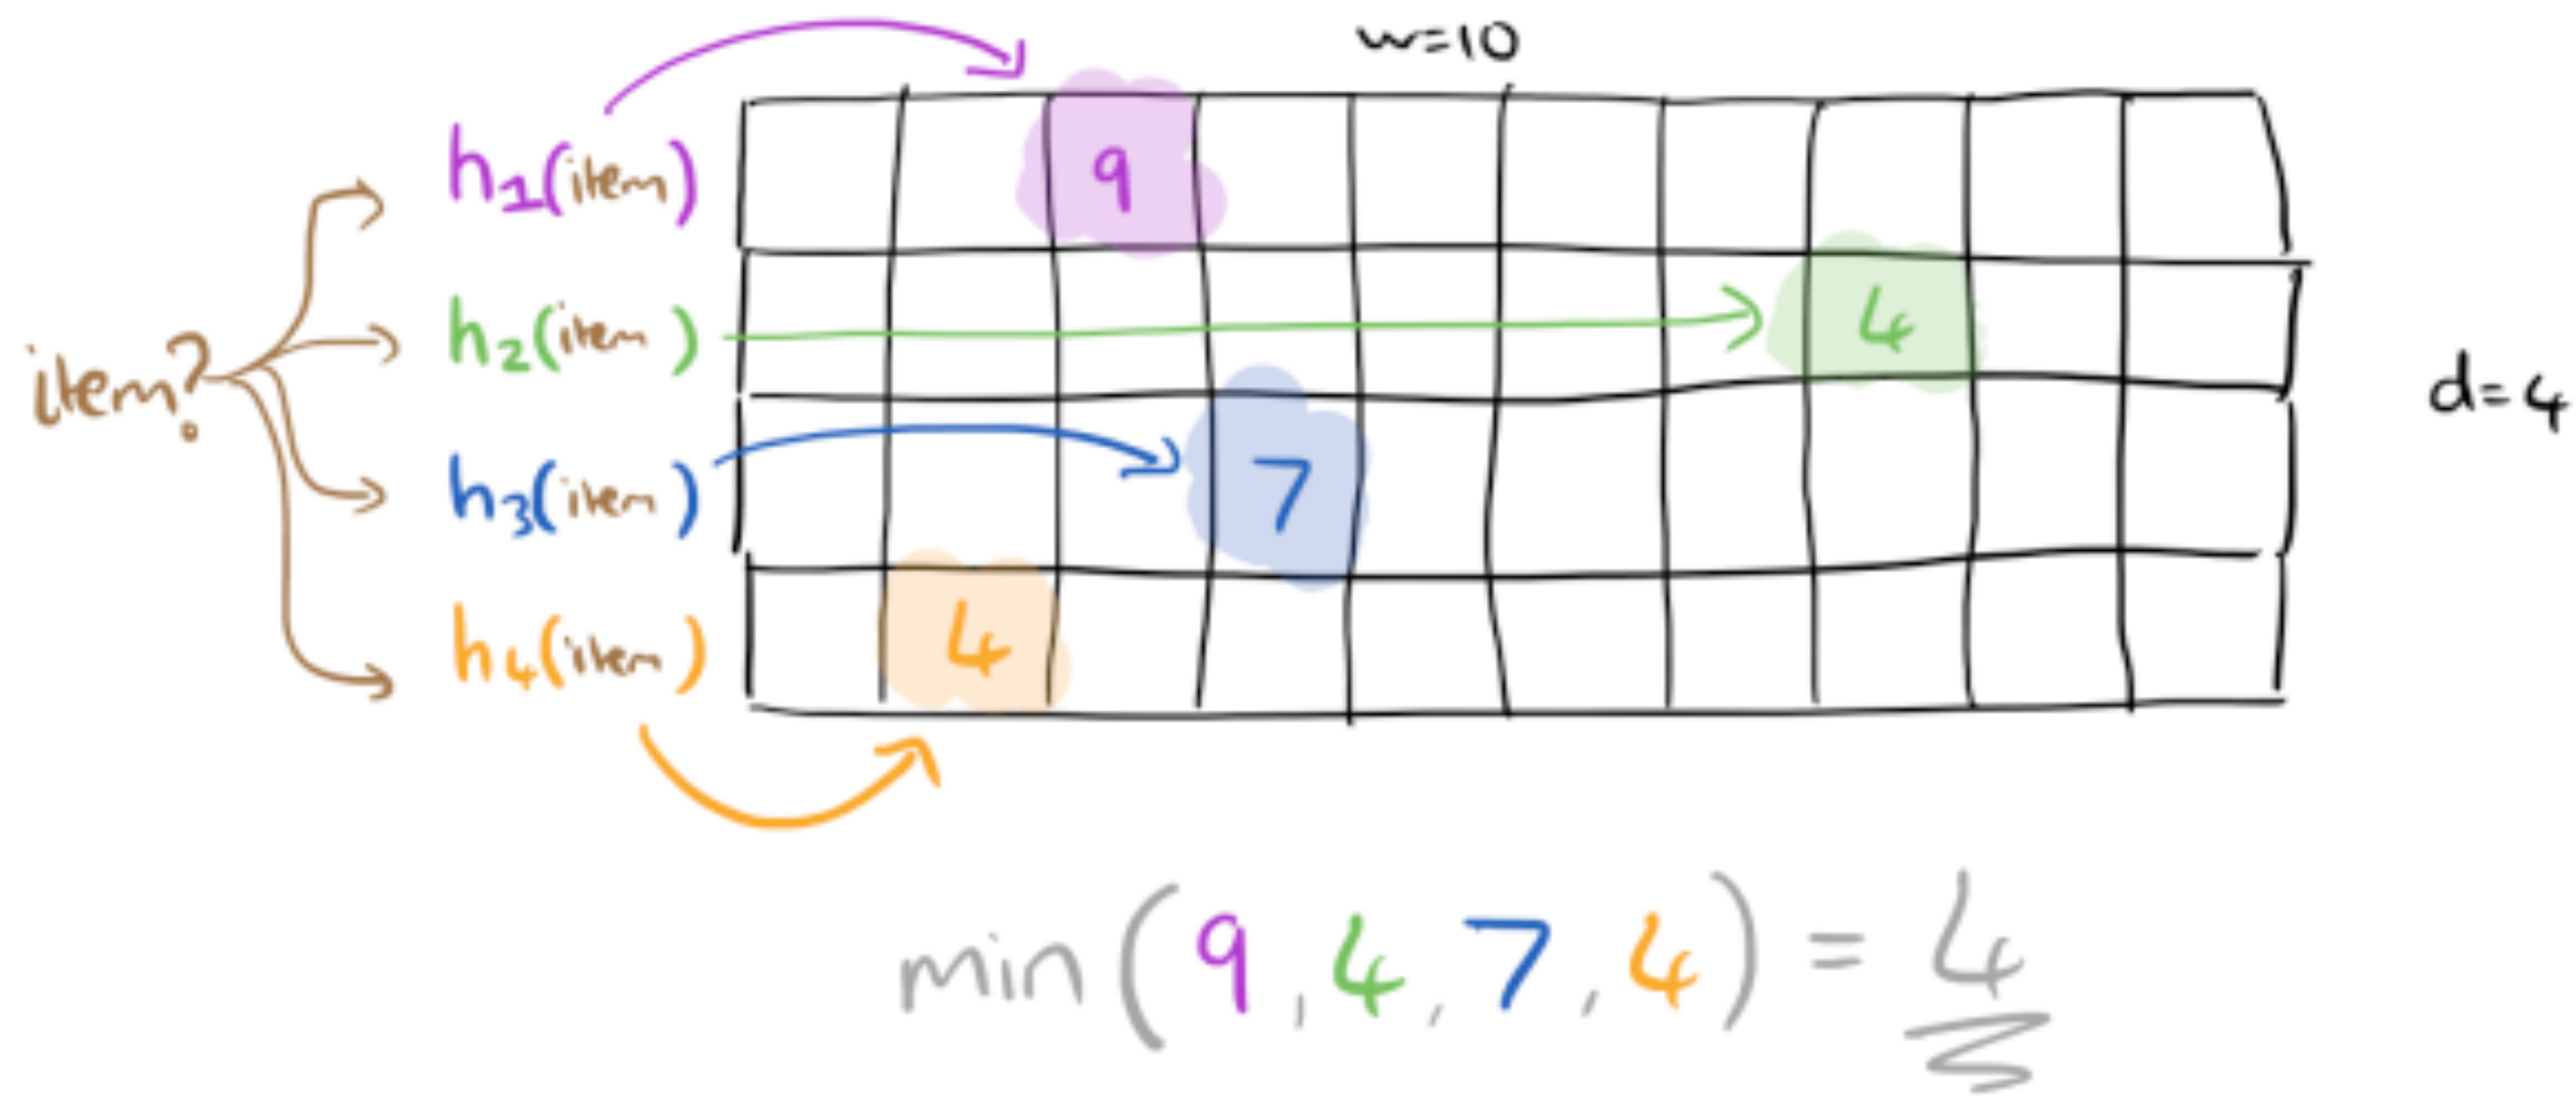
**Input:** [ Edmilson, Alves, 0, Edmilson, Alves, LRB, born, ...]

**Output:** Edmilson ~10 Alves ~8


# Count Min (CM) sketch



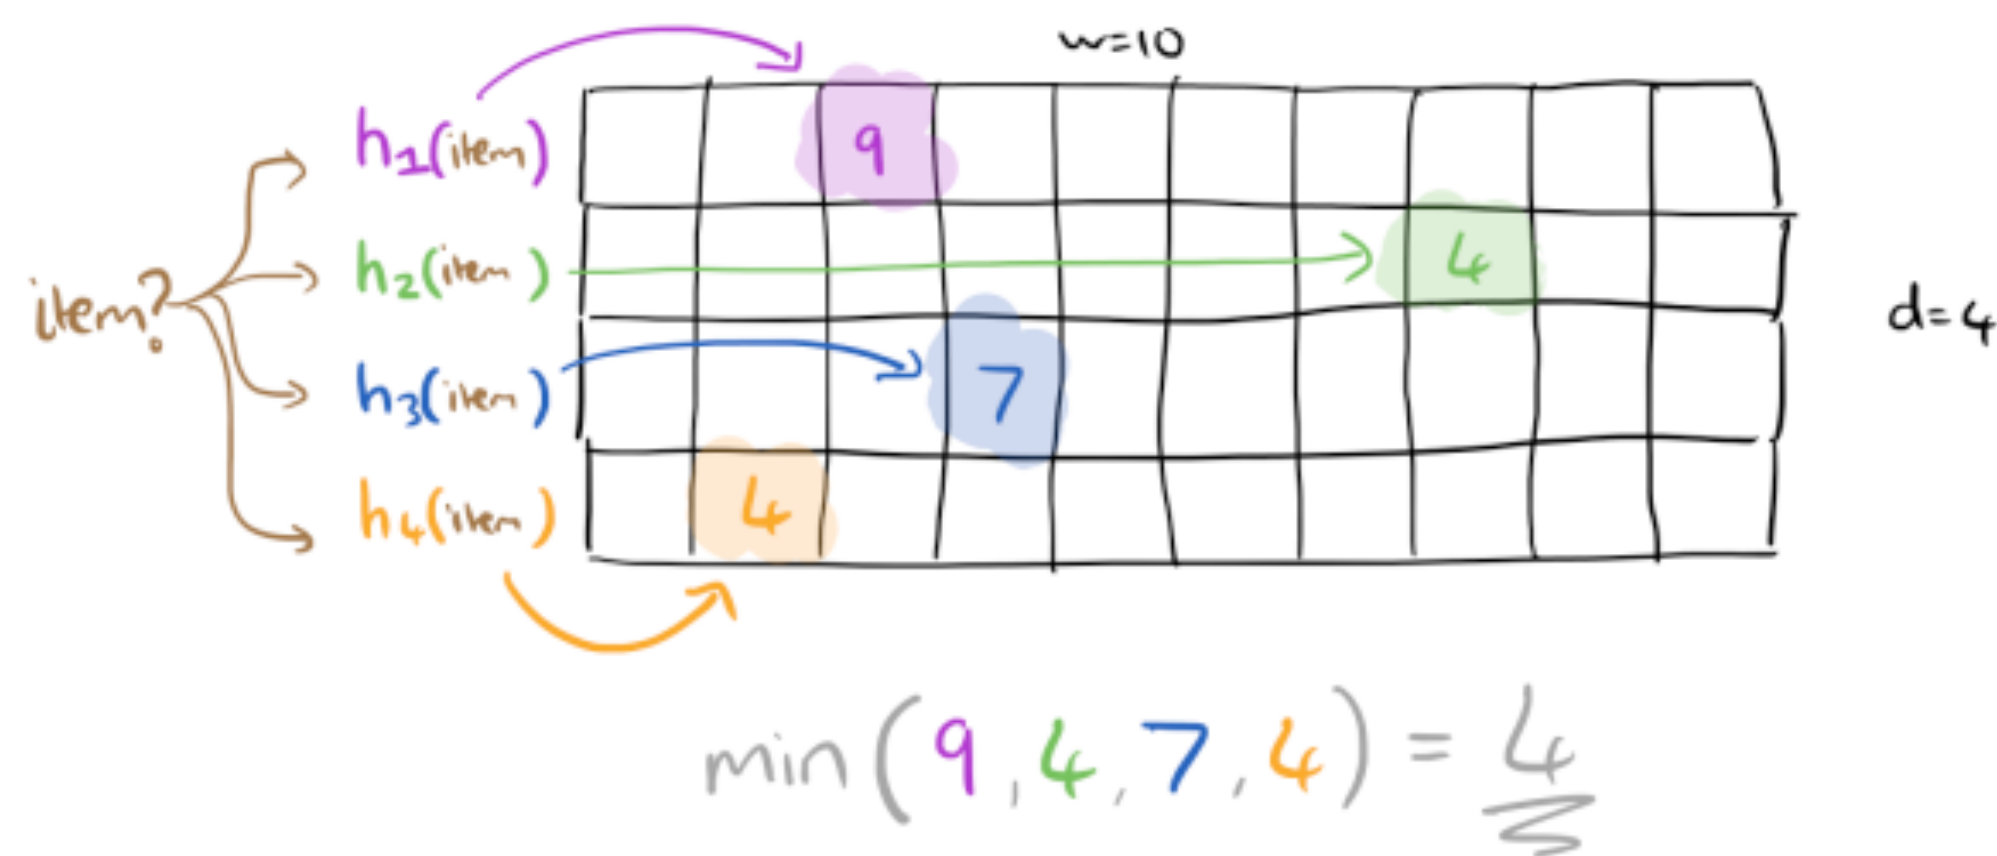
# CM sketch retrieval



# Distributed Tools



retrieve sketched value



take top\_n

## Heavy Hitter for t\_1

And ~10
Bob ~7
Alice ~5
Foo ~3
Bar ~2



# Topology

```

(defn- heavy-hitter-processor
  "Main stream processor takes a configuration and a mapper function to apply."
  [conf]
  (let [streamBuilder (-> (new TopologyBuilder)
    (.addSource (-name conf) string_dser string_dser (into-array [(:input-topic conf)]))
    (.addProcessor "HeavyHitter"
      (reify ProcessorSupplier
        (get [this]
          (get-processor)))
        (into-array [(:name conf)]))
    (.addStateStore
      (->> (Stores/create storeName)
        (.withStringKeys)
        (.withLongValues)
        (.inMemory)
        (.build))
      (into-array ["HeavyHitter"]))
    (.addSink
      "Sink"
      (:output-topic conf)
      string_ser
      string_ser
      (into-array ["HeavyHitter"]))))]
    (.start
      (KafkaStreams.
        streamBuilder
        (get-props conf))))))

```

# Processor

```
(defn ^Processor get-processor []  
  (reify org.apache.kafka.streams.processor.Processor  
    (init [this context]  
      (.schedule (:context @application-state) (:time-window @application-state))  
      (swap! hh/state assoc  
        :top-n 5  
        :number-of-hashfn 10N  
        :bucket-size 1000N)  
      (reset! hh/hitter ^(priority-map))  
      (reset! hh/min-sketch (make-array Integer/TYPE 10N 1000N)) ...)  
    (process [this key value]  
      (debug "Process (k,v)::" key value)  
      (hh/sketch-value value)  
      (hh/add-to-hitter value) ...)  
    (punctuate [this timestamp ...])  
    (close [this]  
      (.close (:store @application-state))))))
```



# WINDOW AGGREGATE

# Distributed Tools

**(key, value)**

**Input:**

Alves ~8 Alves ~10 Edmilson ~5 Edmilson ~3

**Output:**

Alves ~18 Edmilson ~8

# Distributed Tools

```
79 (defn- stream-mapper
80   "Main stream processor takes a configuration and a mapper function to apply."
81   [conf ]
82   (let [streamBuilder (KStreamBuilder.)
83         ^KStream a-stream (.stream
84                             streamBuilder
85                             stringSerde
86                             stringSerde
87                             (into-array String [(:input-topic conf)]))]
88     (-> a-stream
89       (.aggregateByKey (reify Initializer
90                         (apply [this] 0))
91                          (reify Aggregator
92                            (apply [this key value aggregate]
93                                  (+ aggregate (Long/parseLong value))))
94                          (.until (TimeWindows/of "counts" (:window-size conf)) (:window-size conf))
95                          stringSerde
96                          longSerde)
97       (.toStream)
98       (.map (reify KeyValueMapper
99             (apply [this key value]
100                   (toJsonBlob key value))))
101       (.to stringSerde stringSerde (:output-topic conf)))
102     (.start (KafkaStreams. streamBuilder (get-props conf))))))
```

# ELASTICSEARCH SINK

**Input:**

Alves ~18

**Output:**

```
{  
  "name": "Alves",  
  "count": 18,  
  "time": "January 26th 2017, 17:03:00.000"  
}
```

# Distributed Tools

```
20 (defn index [msg conf]
21   (try
22     (debug "Try indexing" (:key msg) (:value msg))
23     (esd/put (:es_connection conf)
24              (:es_index conf)
25              (:es_type conf)
26              (if (nil? (:key msg)) (str (UUID/randomUUID)) (:key msg))
27              (json/read-str (:value msg)))
28     (catch Exception e
29       (error "Failed indexing: " e))))
```

```
31 (defn start-indexing [conf]
32   (go-loop []
33     (let [response (index (<! (:from_kafa conf)) conf)]
34       (when (not-empty response)
35         (debug "Response: " response)))
36     (recur)))
```

# ALMOST KISS

# Distributed Tools

Edmilson Alves 0 Edmilson Alves -LRB- born February 17 , 1976 -RRB- , is a Brazilian midfielder who currently plays for Roasso Kumamoto in the J. League Division 2 .

[ Edmilson, Alves, 0, Edmilson, Alves, LRB, born ...]

Alves ~10 Alves ~8

Alves 18

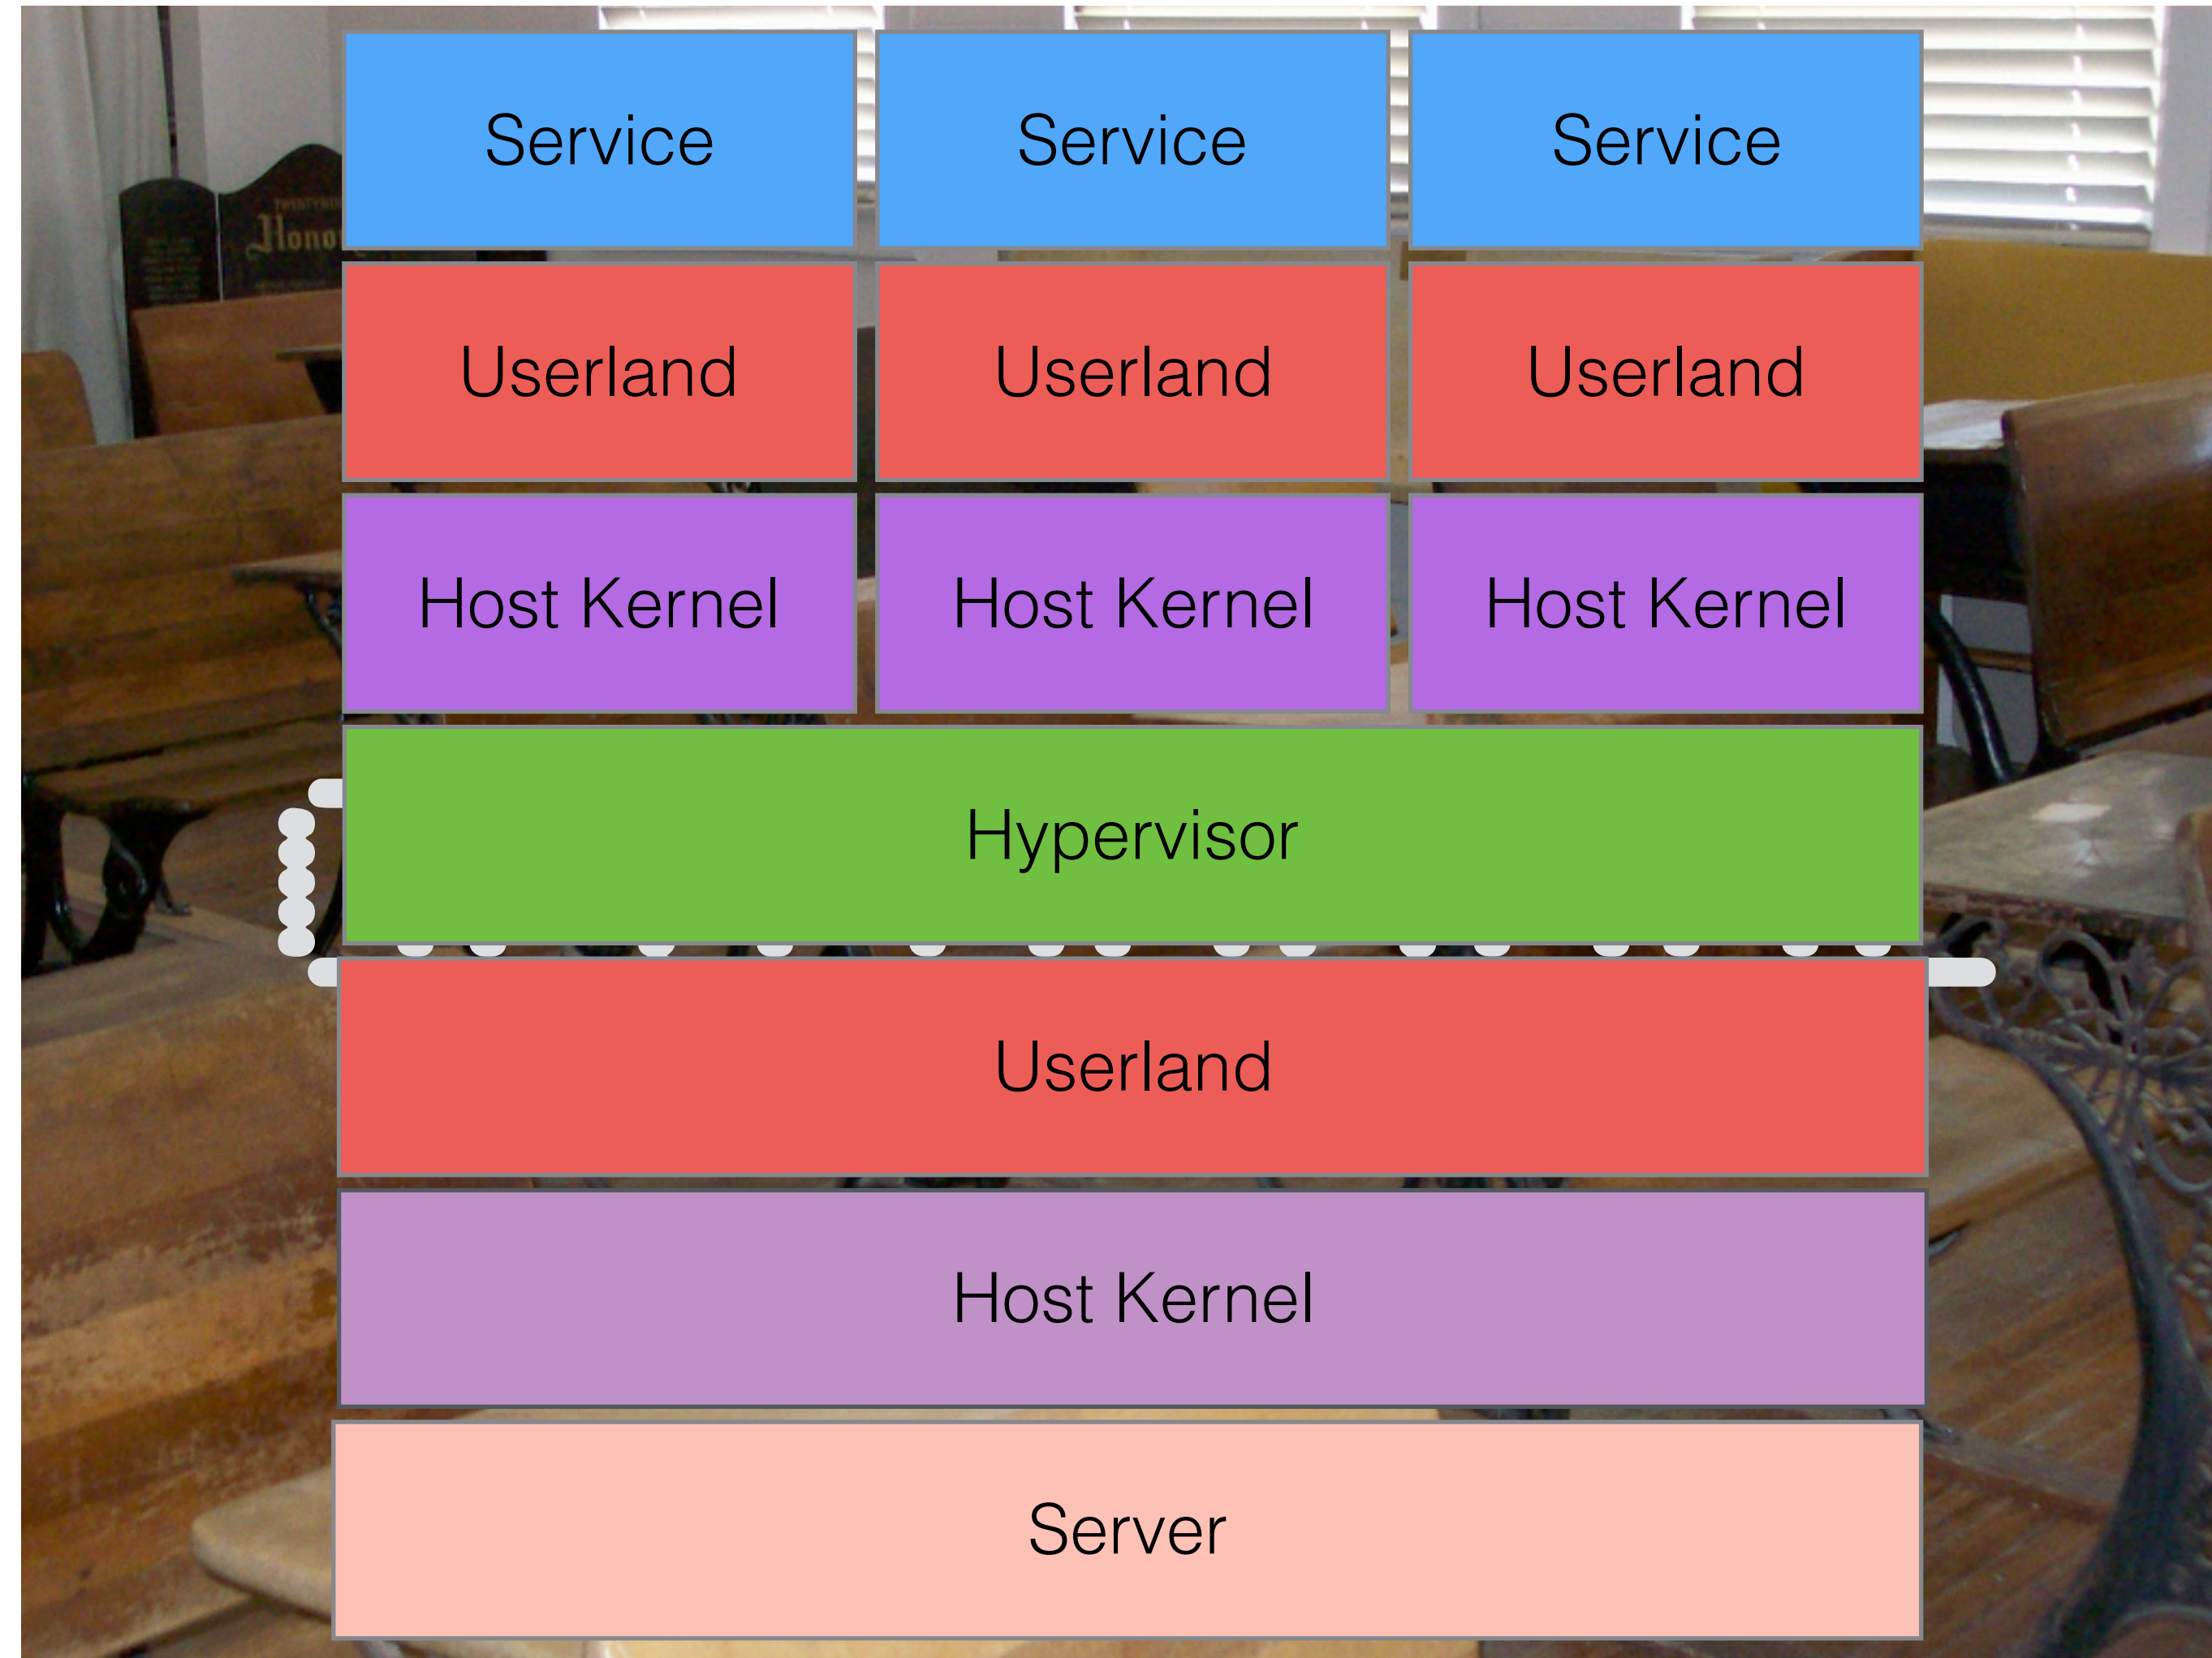
```
{"name": "Alves", "count": 18, "time": "January 26th 2017, 17:03:00.000"}
```



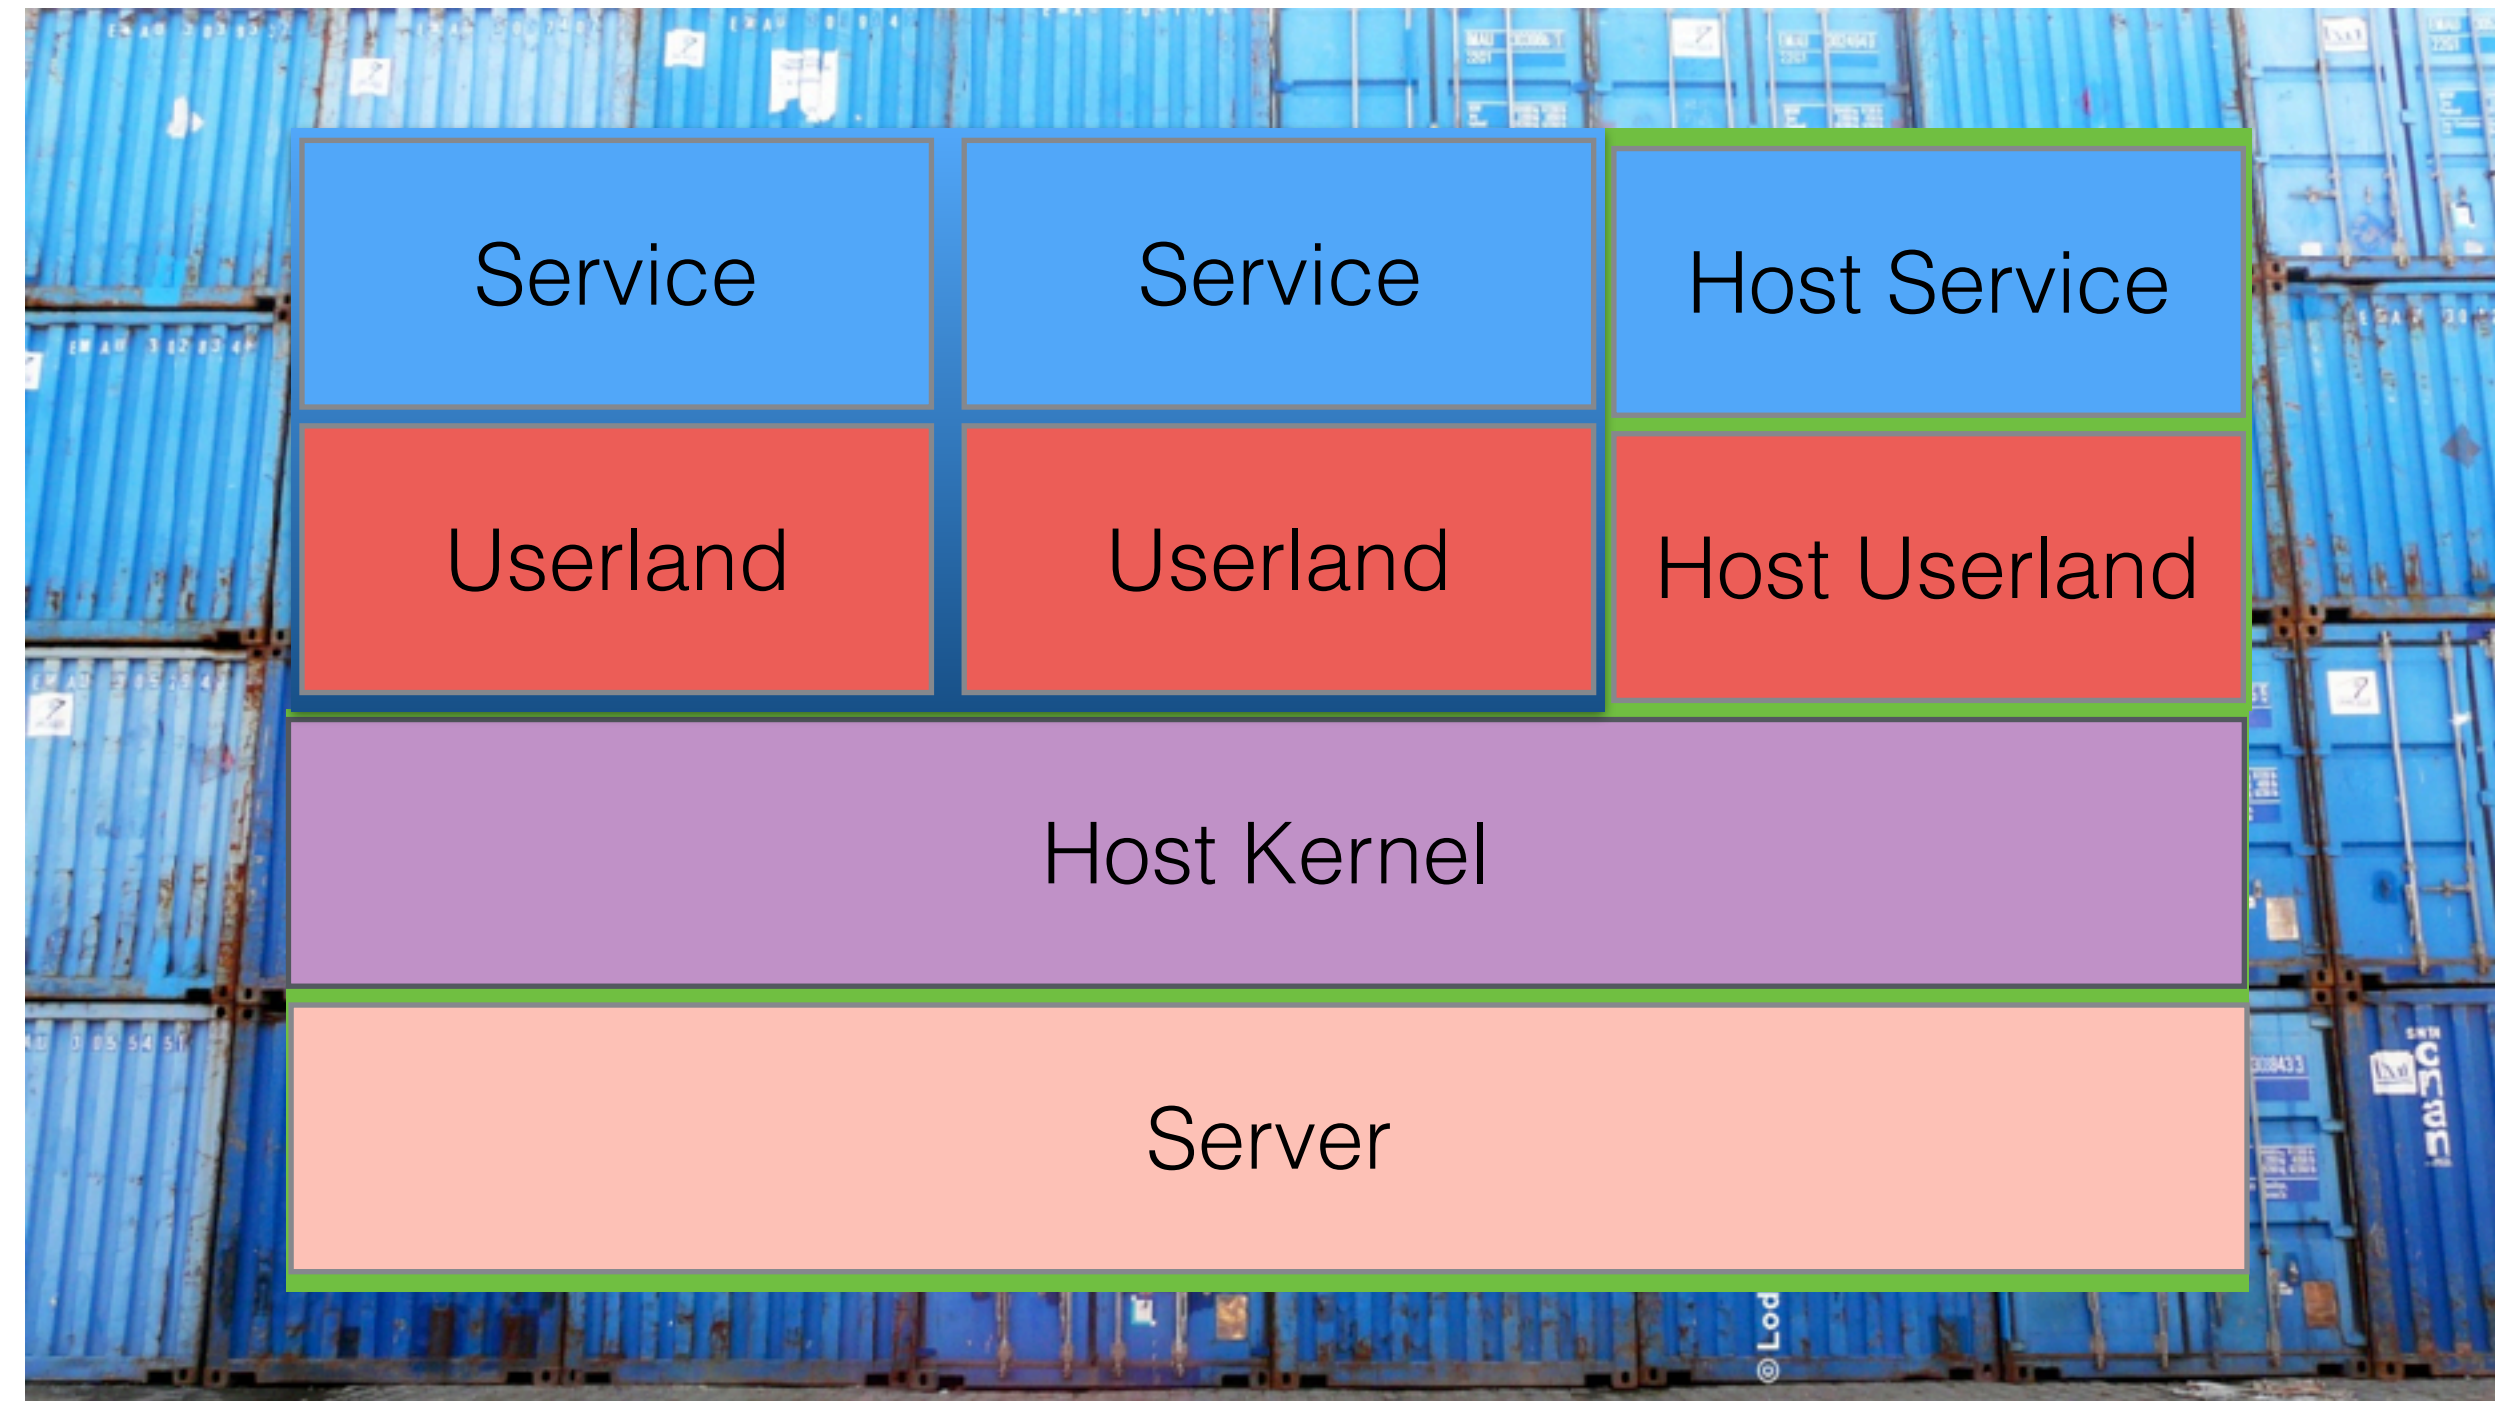
# Container

- [...] **Operating-system-level virtualization** is a server virtualization method in which the kernel of an operating system allows the existence of **multiple isolated user-space** instances

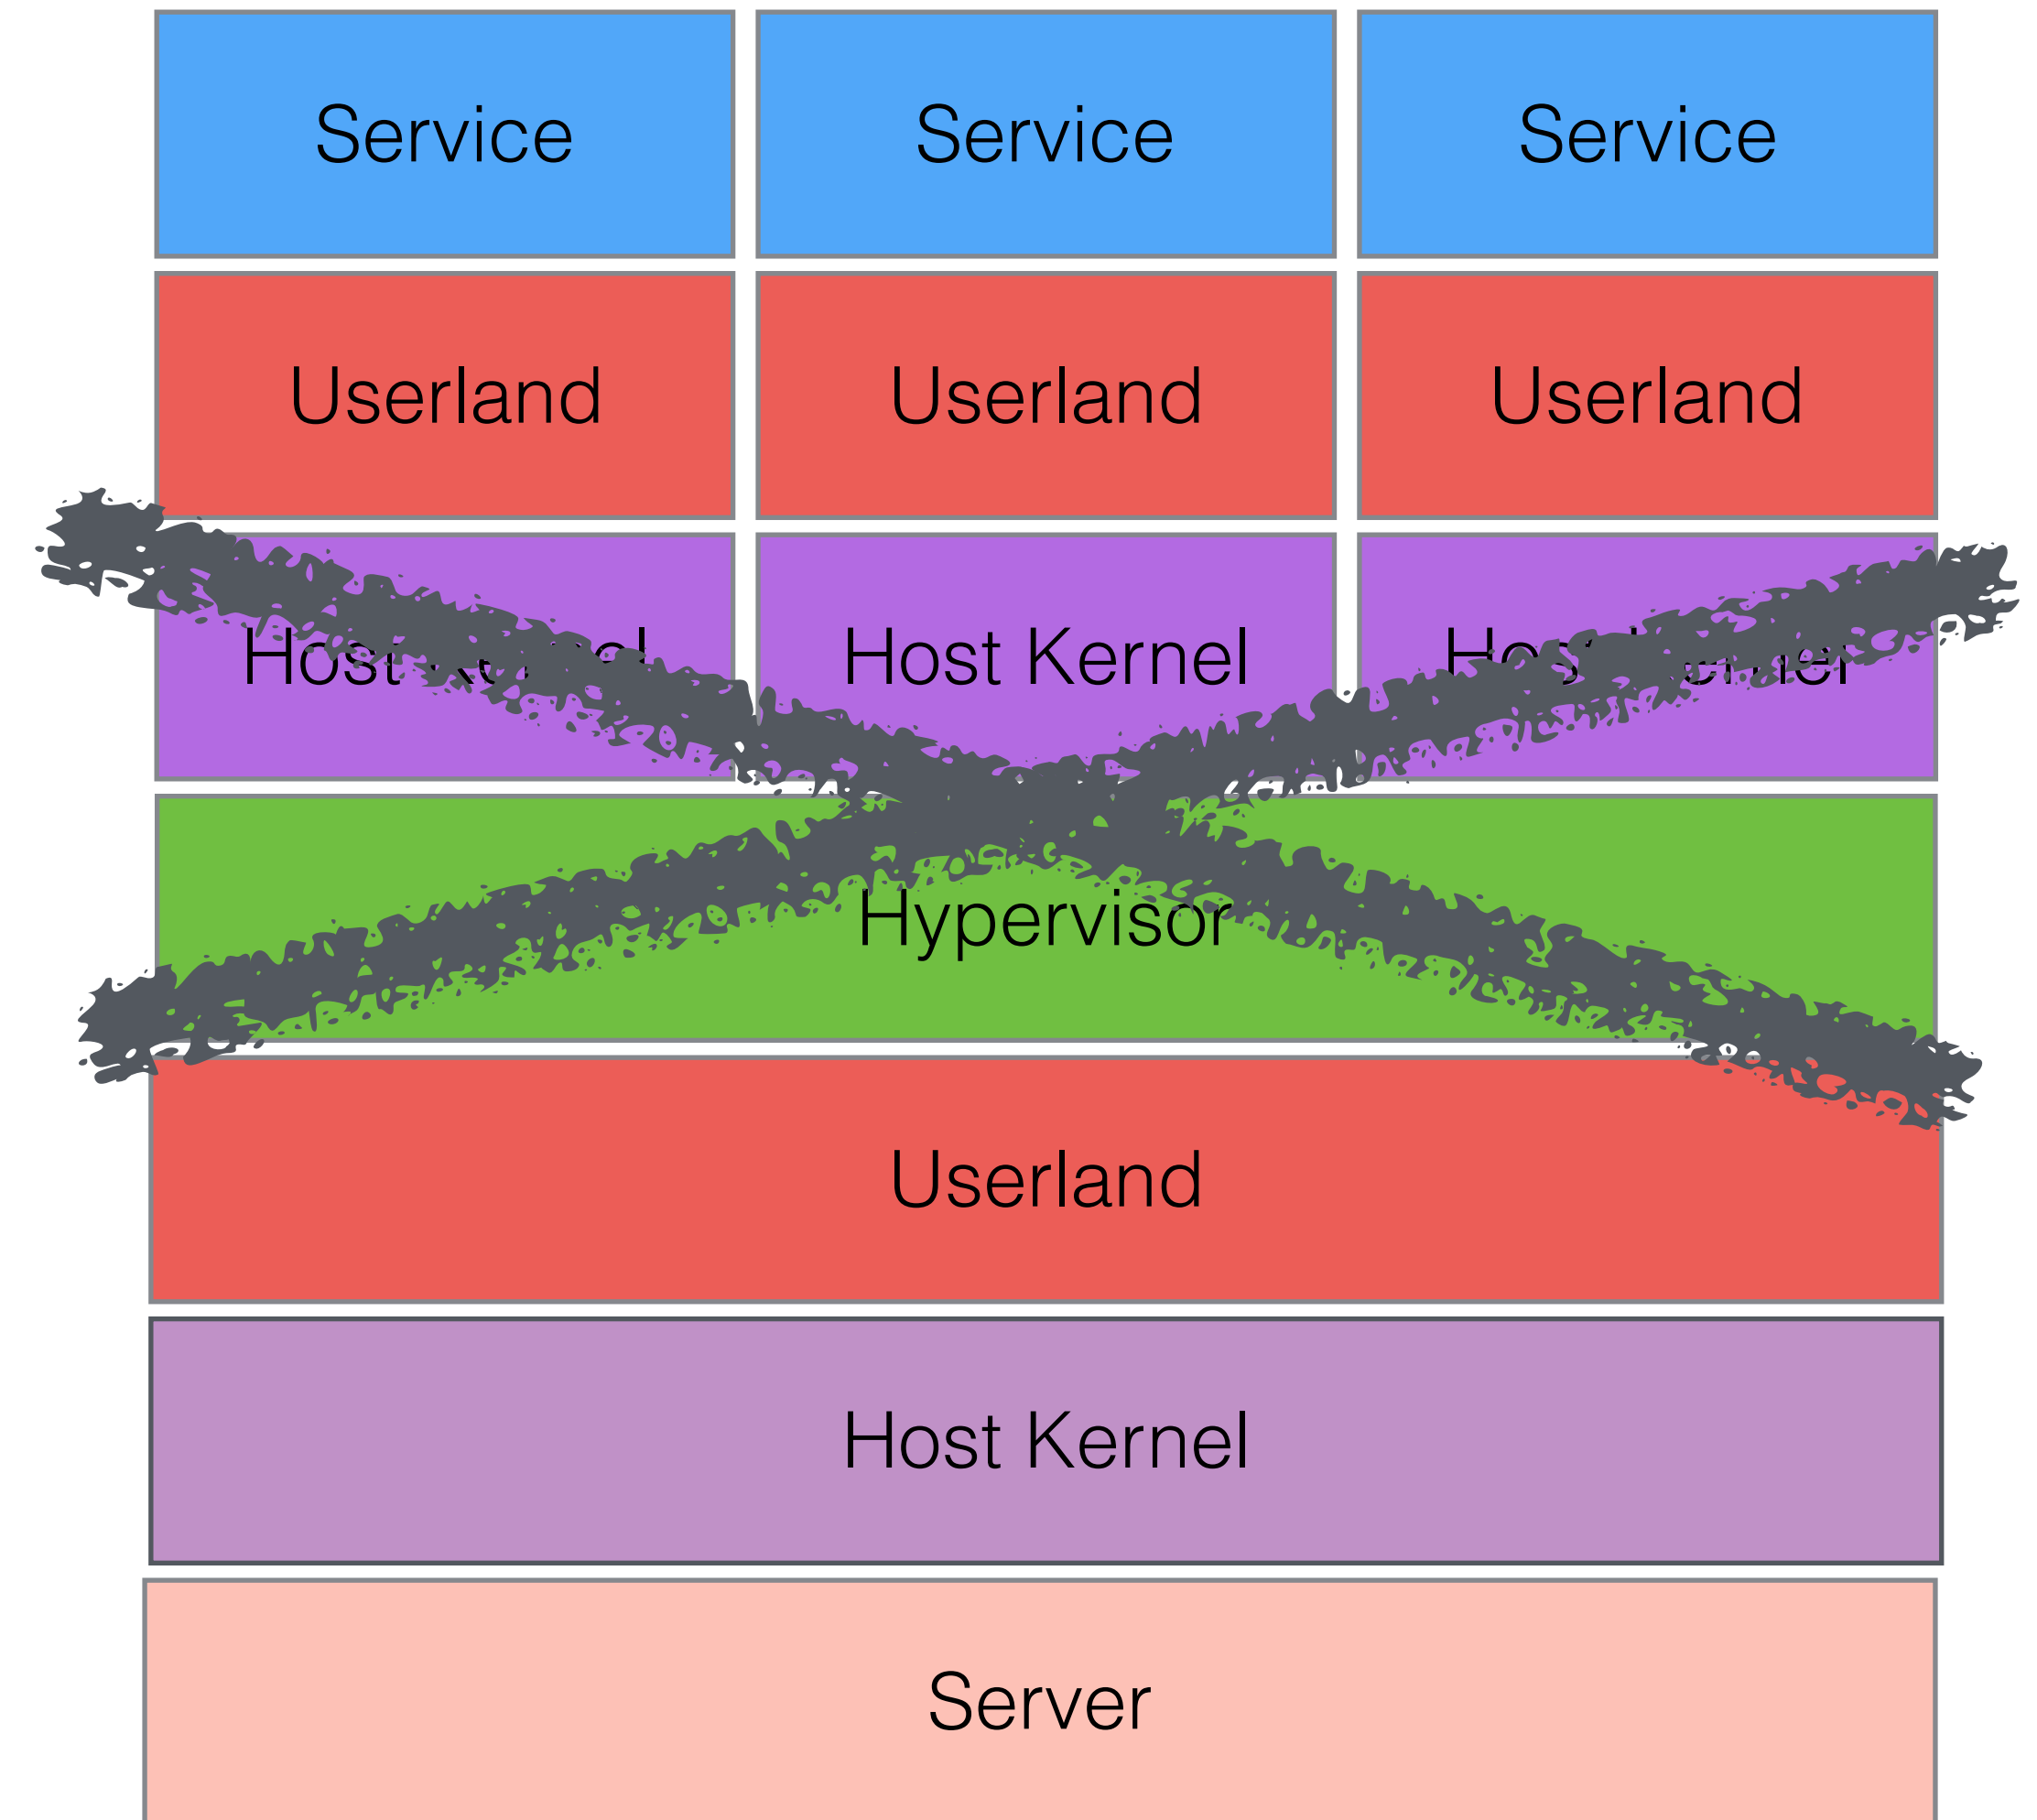
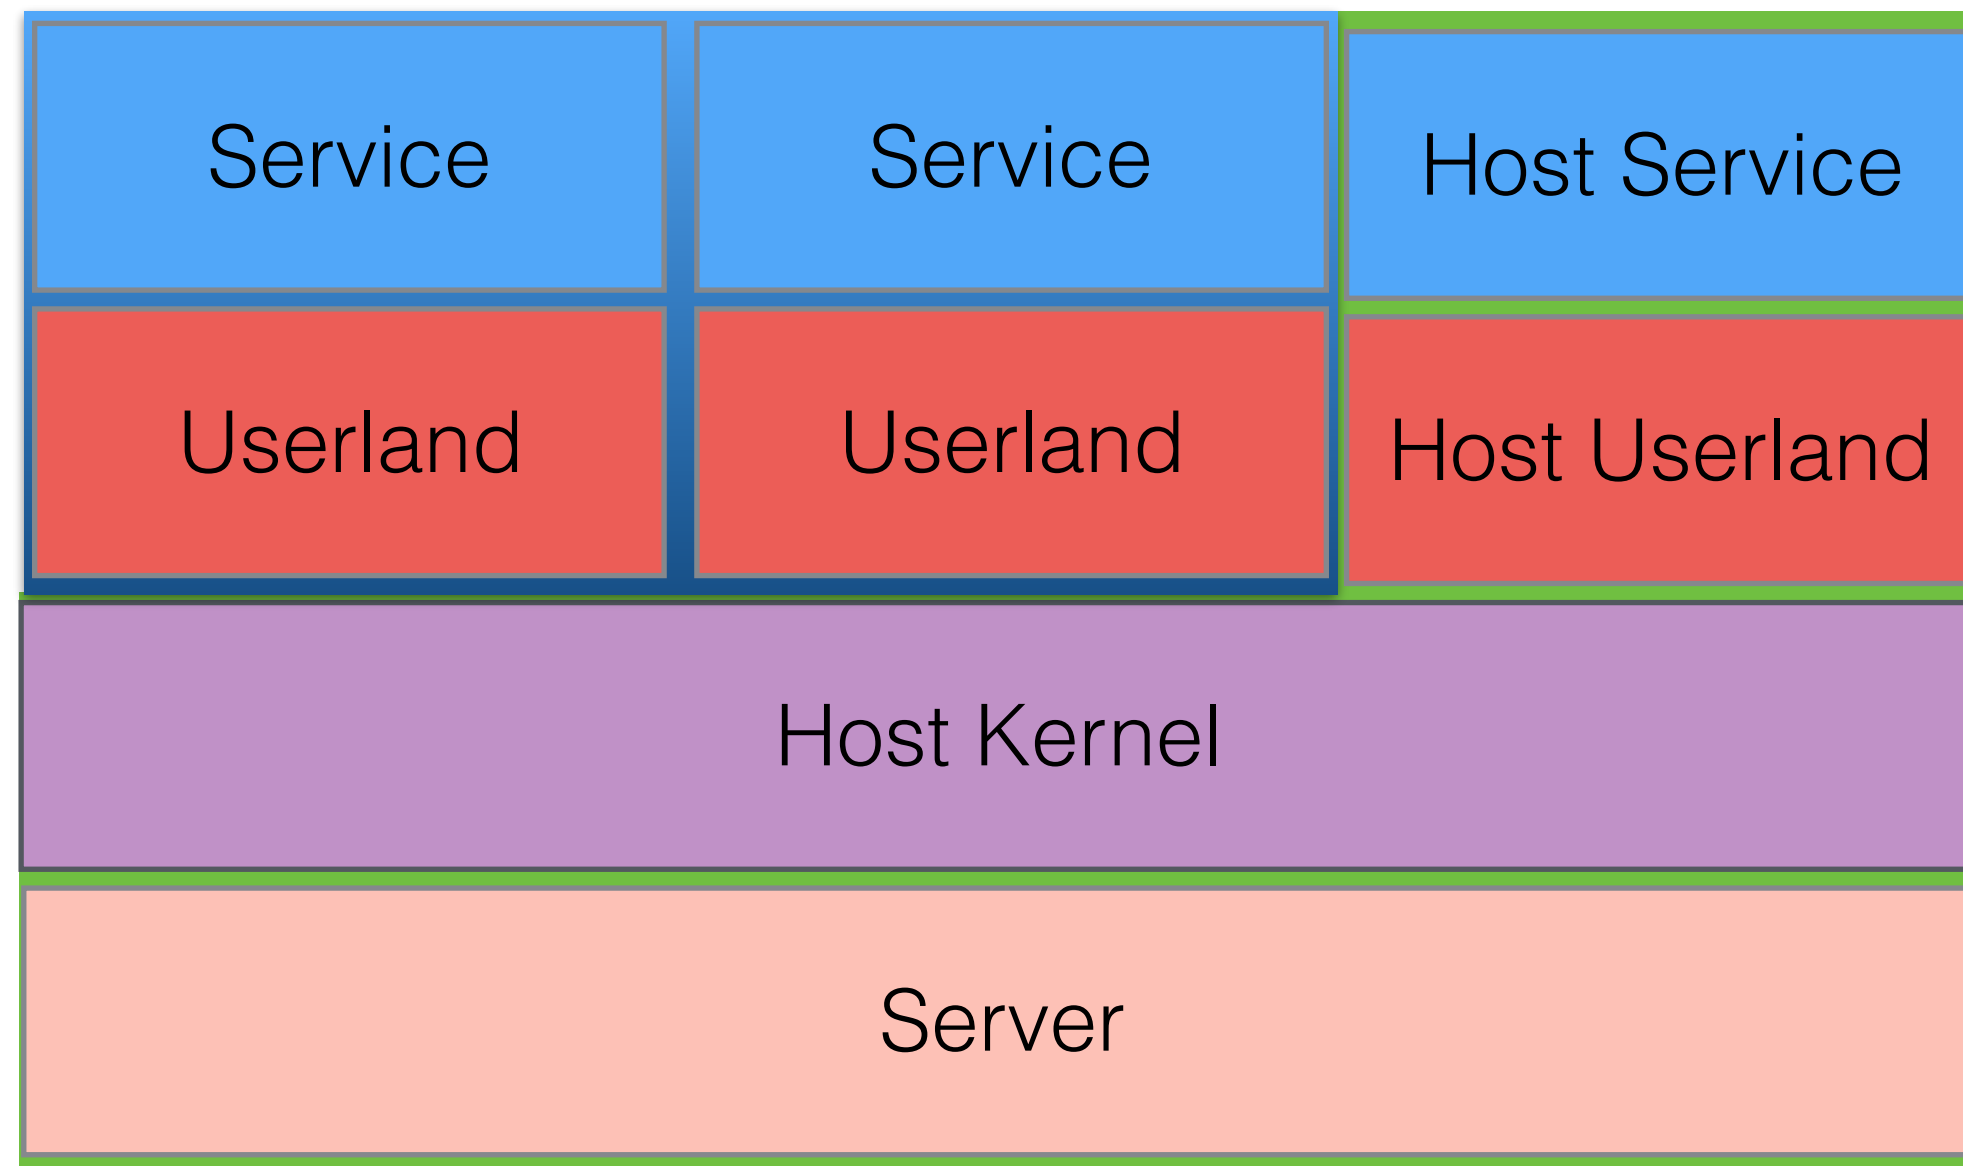
# Containers



# Containers



# Containers






# Containers

java -jar ...args

docker run -t -i CONTAINER-NAME ...args

```
121 clj-kstream-hh:  
122   image: sojoner/clj-kstream-hh:0.1.0  
123   hostname: clj-kstream-hh  
124   container_name: clj-kstream-hh  
125   extends:  
126     file: base.yml  
127     service: sojoner  
128   command: "--broker kafka-broker:9092 --input-topic mapped-test-json --output-topic heavy-hitters --window-size 1 --name stream-hh"
```

# The development setup...



```
$ export DOCKER_HOST=tcp://my.desktop.de:2576
```

# THE PIPE IN CONTAINERS



# Containers



***Discovering and configuring services in your infrastructure.***



ZooKeeper

***enables highly reliable distributed coordination.***



kafka

***Distributed append log a.k.a Message Broker***




elastic

***Provides a distributed full-text search engine***

# Containers



```
3 consul:
4   image: qnib/alpn-consul
5   hostname: consul
6   container_name: consul
7   networks:
8     - network
9   environment:
10    - DC_NAME=es
11    - RUN_SERVER=true
12    - BOOTSTRAP_CONSUL=true
13    - DNS_RECURSOR=8.8.8.8
14   ports:
15    - "8500:8500"
```




# Containers



## ZooKeeper


```
17  zookeeper:  
18  image: qnib/zookeeper  
19  hostname: zookeeper  
20  container_name: zookeeper  
21  extends:  
22  | file: base.yml  
23  | service: sojoner  
24  ports:  
25  - "2181:2181"
```



# Containers



```
27 kafka-broker:  
28 image: qnib/kafka:0.10.0.1  
29 hostname: kafka-broker  
30 container_name: kafka-broker  
31 extends:  
32   file: base.yml  
33   service: sojoner  
34 volumes:  
35   - /tmp/kafka-logs  
36 ports:  
37   - "9092:9092"
```







# Containers



```
52 elasticsearch1:  
53 image: docker.elastic.co/elasticsearch/elasticsearch:5.3.0  
54 container_name: elasticsearch1  
55 environment:  
56   - xpack.security.enabled=false  
57   - cluster.name=clj-kstream-es-docker-cluster  
58   - bootstrap.memory_lock=true  
59   - "ES_JAVA_OPTS=-Xms512m -Xmx512m"  
60 ulimits:  
61   memlock:  
62     soft: -1  
63     hard: -1  
64   nofile:  
65     soft: 65536  
66     hard: 65536  
67 mem_limit: 1g  
68 cap_add:  
69   - IPC_LOCK  
70 volumes:  
71   - esdata1:/usr/share/elasticsearch/data  
72 ports:  
73   - 9200:9200  
74 networks:  
75   - network
```




# Containers




# THE TOOLS IN CONTAINERS

# Containers



```
1 FROM qnib/alpn-jre8
2 ADD clj-kstream-cutter.jar /usr/share/clj-kstream/clj-kstream-cutter.jar
3 ENTRYPOINT ["java", "-jar", "/usr/share/clj-kstream/clj-kstream-cutter.jar"]
4 CMD []
```



```
1 #!/bin/bash
2 mv ../target/clj-kstream-cutter.jar .
3 docker build --tag "sojoner/clj-kstream-cutter:0.2.1" .
4 docker tag <HASH> sojoner/clj-kstream-cutter:0.2.1
5 docker login
6 docker push sojoner/clj-kstream-cutter
```




# Containers

```
112 clj-kstream-cutter:  
113   image: sojoner/clj-kstream-cutter:0.2.0  
114   hostname: clj-kstream-cutter  
115   container_name: clj-kstream-cutter  
116   extends:  
117     file: base.yml  
118     service: sojoner  
119   command: "--broker kafka-broker:9092 --zookeeper zookeeper:2181 --input-topic logs-replay --output-topic mapped-test-json --selector msg --name stream-cut-json-field"  
120  
121 clj-kstream-hh:  
122   image: sojoner/clj-kstream-hh:0.1.0  
123   hostname: clj-kstream-hh  
124   container_name: clj-kstream-hh  
125   extends:  
126     file: base.yml  
127     service: sojoner  
128   command: "--broker kafka-broker:9092 --input-topic mapped-test-json --output-topic heavy-hitters --window-size 1 --name stream-hh"  
129  
130 clj-kstream-string-long-window-aggregate:  
131   image: sojoner/clj-kstream-string-long-window-aggregate:0.2.2  
132   hostname: clj-kstream-string-long-window-aggregate  
133   container_name: clj-kstream-string-long-window-aggregate  
134   extends:  
135     file: base.yml  
136     service: sojoner  
137   command: "--broker kafka-broker:9092 --input-topic heavy-hitters --window-size 1 --output-topic agg-result --name stream-agg"  
138  
139 clj-kstream-elasticsearch-sink:  
140   image: sojoner/clj-kstream-elasticsearch-sink:0.0.1  
141   hostname: clj-elasticsearch-sink  
142   container_name: clj-elasticsearch-sink  
143   extends:  
144     file: base.yml  
145     service: sojoner  
146   command: "--broker kafka-broker:9092 --topic agg-result --elasticsearch http://elasticsearch1:9200 --index heavy-hitters-test-idx --index-type hh-struct"
```

# A datacenter setup

Build a Docker Swarm



```
$ export DOCKER_HOST=tcp://my.datacenter.de:2576
```

# Containers

```
1  version: '3'
2  services:
3  zookeeper: <4 keys>
21 zkui: <5 keys>
41 broker:
42   image: qnib/plain-kafka:0.10.0.1
43   networks:
44     - backend_services
45   ports:
46     - "9092:9092"
47   deploy:
48     replicas: 1
49     resources:
50       limits:
51         cpus: '1'
52         memory: 768M
53       update_config: <2 keys>
56       restart_policy: <1 key>
58       environment: <2 keys>
61     kafkamanager: <5 keys>
77     esmaster: <5 keys>
101    esdata: <4 keys>
122    kibana: <5 keys>
139 # a network for our stack
140 networks:
141   backend_services:
142     external: true
```



# Containers

```
1  version: '3'
2  services:
3  - clj-kstream-lf-producer:
4    image: sojoner/clj-kstream-lf-producer:0.1.0
5    hostname: clj-kstream-lf-producer
6    networks:
7      - backend_services
8    command: "--broker backend_broker:9092 --topic logs-replay"
9    deploy:
10     replicas: 1
11     resources: <1 key>
15     update_config:
16       parallelism: 1
17       delay: 15s
18     restart_policy:
19       condition: on-failure
20
21     clj-kstream-cutter: <5 keys>
38
39     clj-kstream-hh: <5 keys>
56
57     clj-kstre...-aggregate: <5 keys>
74
75     clj-kstre...earch-sink: <5 keys>
92
93     # a network for our stack
94     networks:
95     - backend_services:
96       external: true
```

# Containers

```
$ docker network create --driver overlay --attachable=true backend_services
```

```
$ docker stack deploy --compose-file backend.yml backend
```

```
$ docker stack deploy --compose-file streamprocessors.yml kstream
```

# Recap

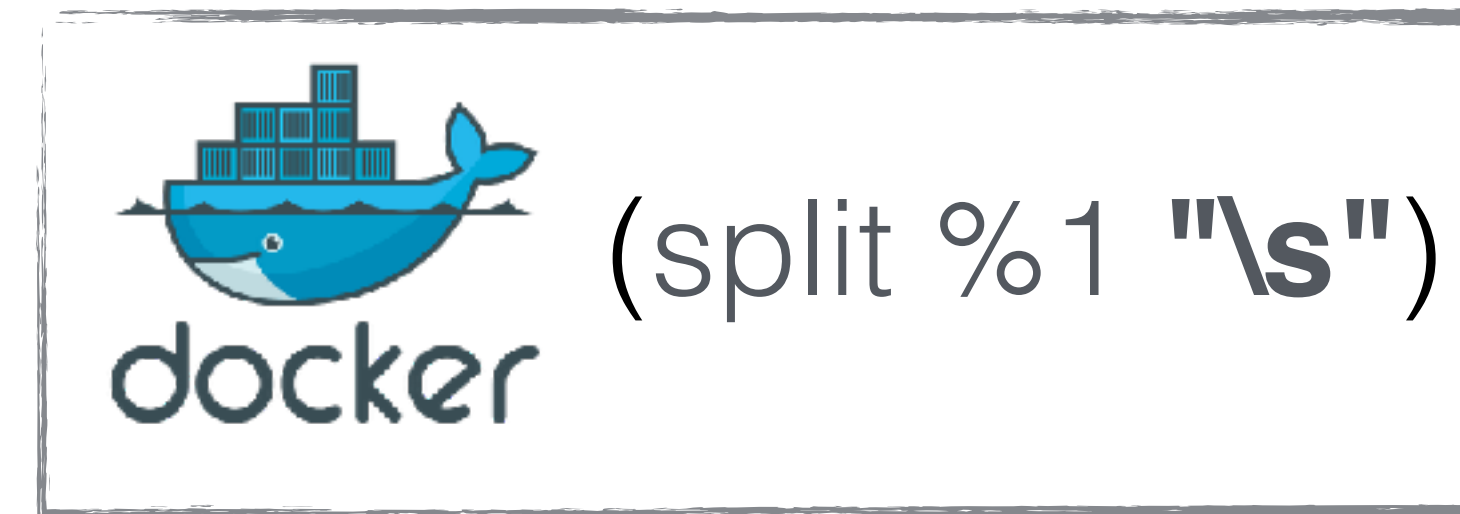

**TL;DR**



(split %1 "\s")

AS MESSAGE BROKER

AS STREAM PROCESSOR





**GOTCHAS ?**

## Recap

- Kafka Streams still **at least once**
  - but **exactly ones** is coming
- **Still need capacity planning**
- **Testing / Debugging is still a challenge**
  - **Consistency** of the **state storage**
  - **Processing Time vs. Event Time**
- What about **Amdahl's law** ?
- How to manage **Docker Volumes** nicely @scale

# ORIENTATION

Human Kind:

*„Take what you need“*

Albert Einstein:

*„Make things as simple as possible, but no simpler“*

William of Ockham:

*„Among competing hypotheses, the one with the fewest assumptions should be selected“*

# Containerizing Distributed Pipes

(thanks (listening [this]))

- <http://kafka.apache.org/>
- <https://martin.kleppmann.com/2015/05/06/data-agility-at-strata.html>
- <https://www.confluent.io/blog/apache-kafka-samza-and-the-unix-philosophy-of-distributed-data/>
- <https://speakerdeck.com/ept/kafka-and-samza-distributed-stream-processing-in-practice>
- <https://github.com/mhausenblas/dnpipes>
- [https://en.wikipedia.org/wiki/Pipeline\\_%28Unix%29](https://en.wikipedia.org/wiki/Pipeline_%28Unix%29)
- <https://zookeeper.apache.org/doc/trunk/zookeeperOver.html>
- <https://github.com/sojoner/container-stacks/tree/master/kafkaelasticsearch>
- [https://kafka.apache.org/documentation/streams#streams\\_processor](https://kafka.apache.org/documentation/streams#streams_processor)
- [https://kafka.apache.org/documentation/streams#streams\\_dsl](https://kafka.apache.org/documentation/streams#streams_dsl)
- <https://hub.docker.com/r/sojoner/clj-kstream-elasticsearch-sink/>
- <https://hub.docker.com/r/sojoner/clj-kstream-cutter/>
- <https://hub.docker.com/r/sojoner/clj-kstream-hh/>
- <https://hub.docker.com/r/sojoner/clj-kstream-string-long-window-aggregate/>
- <https://blog.acolyer.org/2016/07/21/time-adaptive-sketches-ada-sketches-for-summarizing-data-streams/>
- <https://github.com/mhausenblas/dnpipes>