

redislabs

Home of Redis

Redis for Fast Data Ingest



Agenda


- Fast Data Ingest and its challenges
- Redis for Fast Data Ingest
 - Pub/Sub
 - List
 - Sorted Sets as a Time Series Database
- The Demo
- Scaling with Redis^e Flash




Fast Data Ingest Scenarios



IOT




wiselife^{uw}
Control your business
utilities from your phone

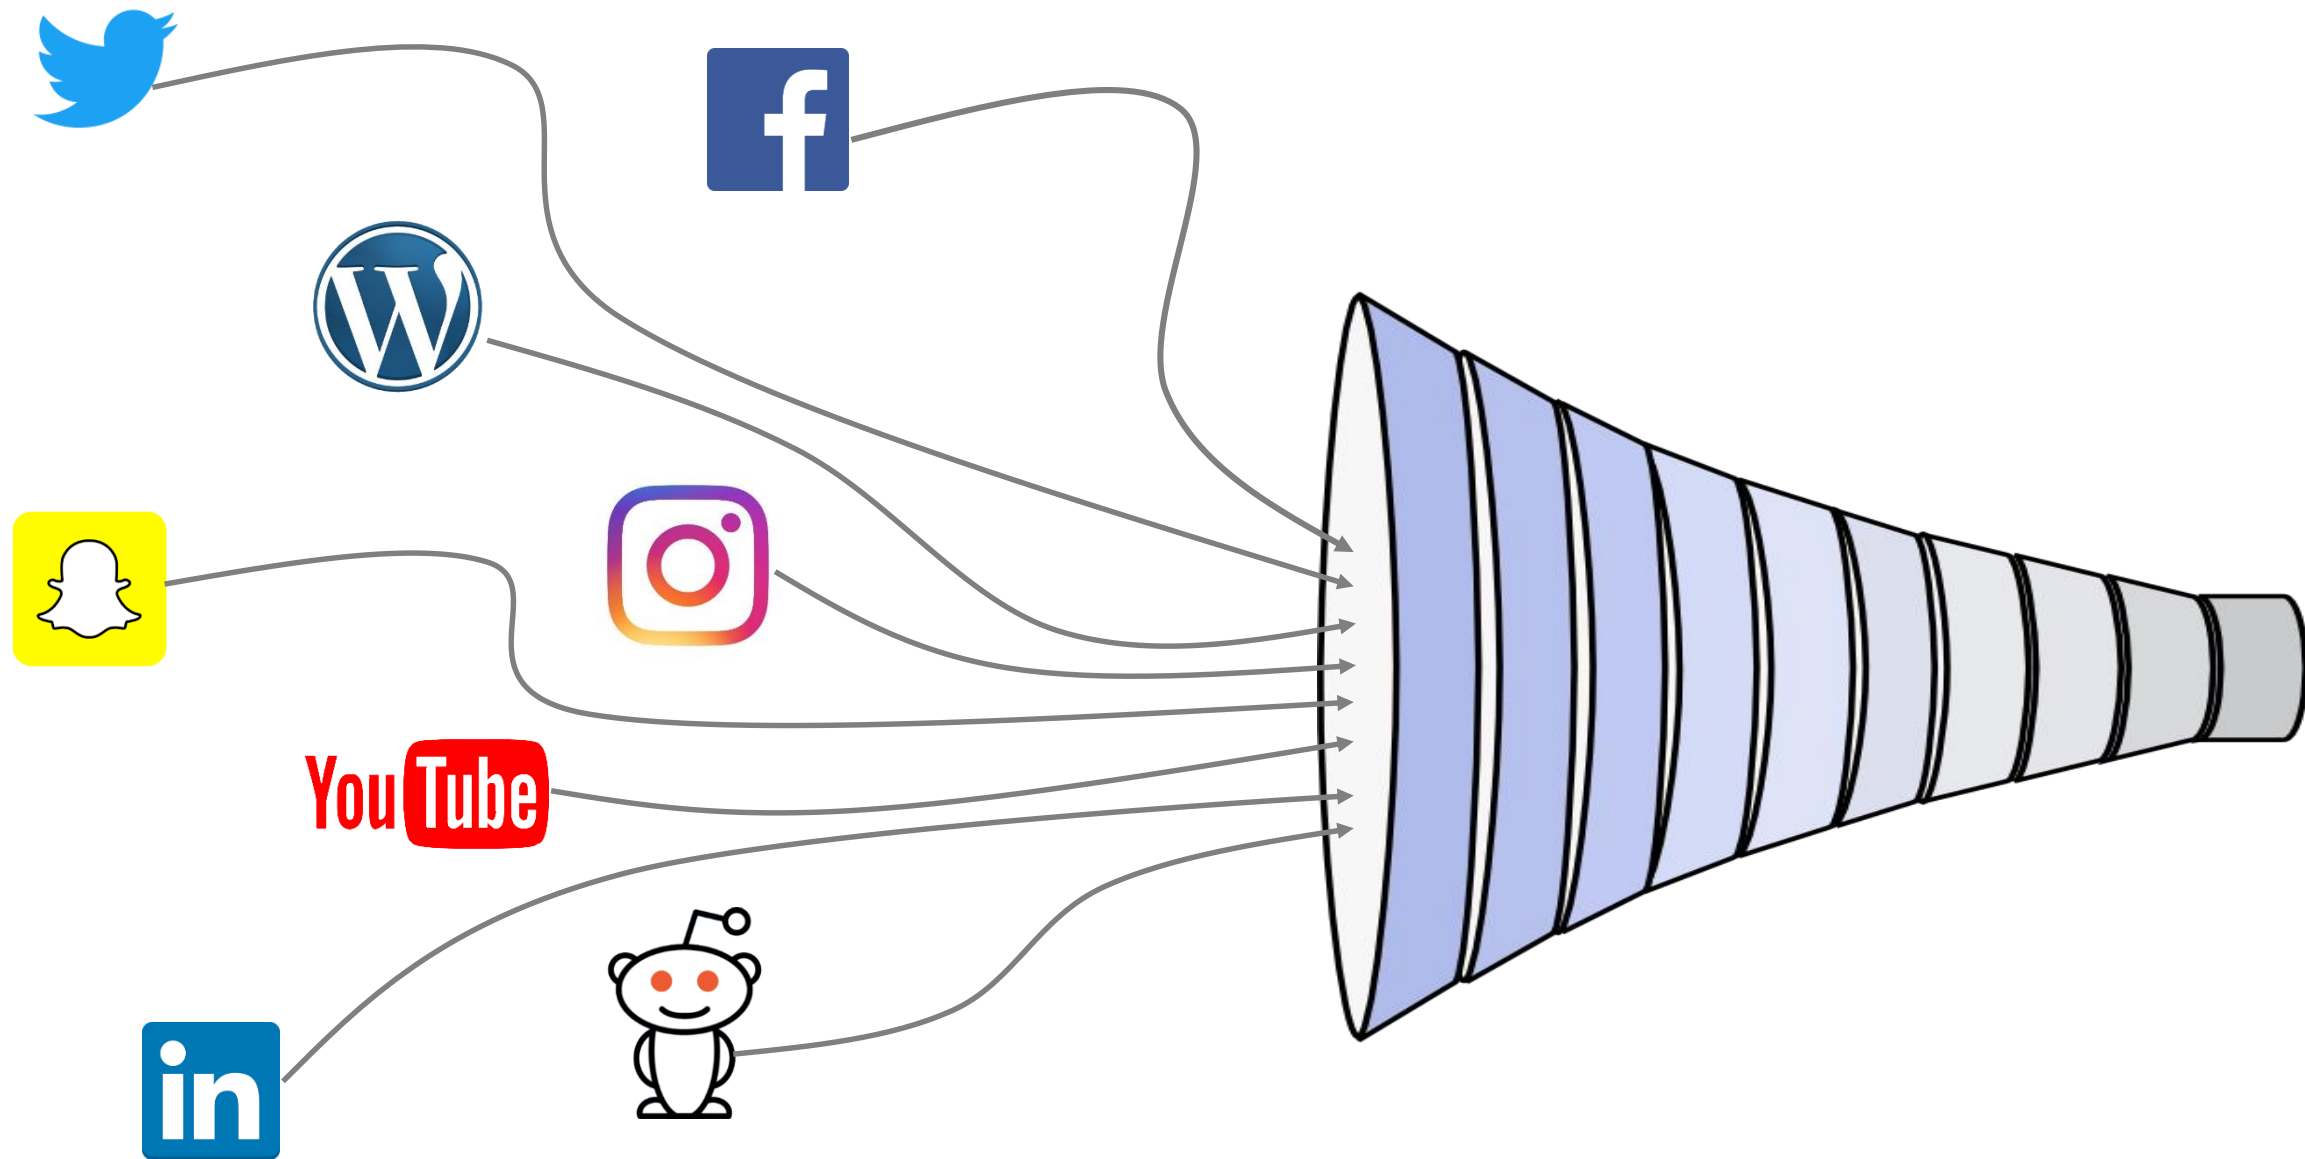


utilitywise^{uw}


Network Traffic Inspection



Social Media Analysis



More Scenarios



Log Collection



User Activity Tracking



Multi-player Gaming



Fintech

And more...

Fast Data Ingest Challenges

- Keeping up with the pace of data arrival
- Data from multiple sources with no standard data format
- Filter, analyze, and transform data in real-time
- Managing data arriving from sources distributed geographically

Requirements for Fast Data Ingest

- Physical infrastructure – network, computational resources, etc.
 - Software stack to:
 - Filter
 - Aggregate
 - Transform
 - Distribute
- data in real-time with sub-millisecond latency

Fast Data Ingest with Redis



About Redis

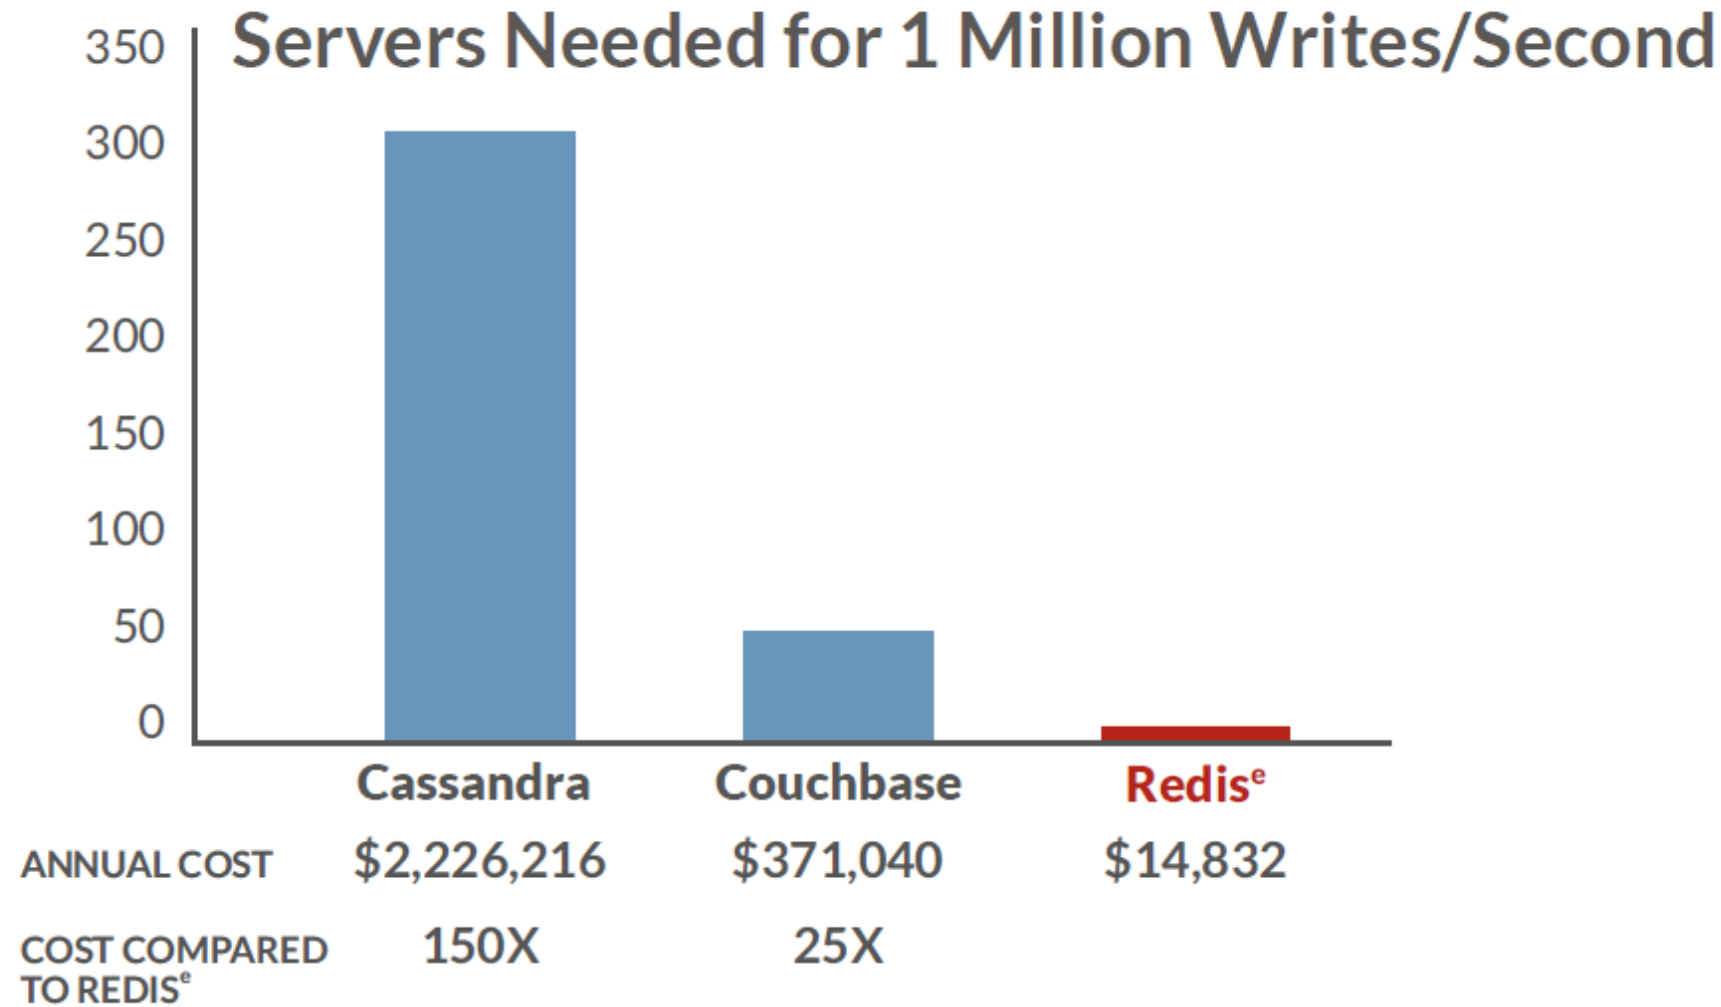


Open source. The leading **in-memory database platform**, supporting any high performance operational, analytics or hybrid use case.




The open source home and commercial provider of **Redis Enterprise (Redis[®])** technology, platform, products & services.

Redis for Fast Data Ingest




Redis^e delivers the maximum throughput with the lowest number of servers, slashing operational costs by up to 99% (Google Cloud Platform performance benchmarks)

Redis for Fast Data Ingest



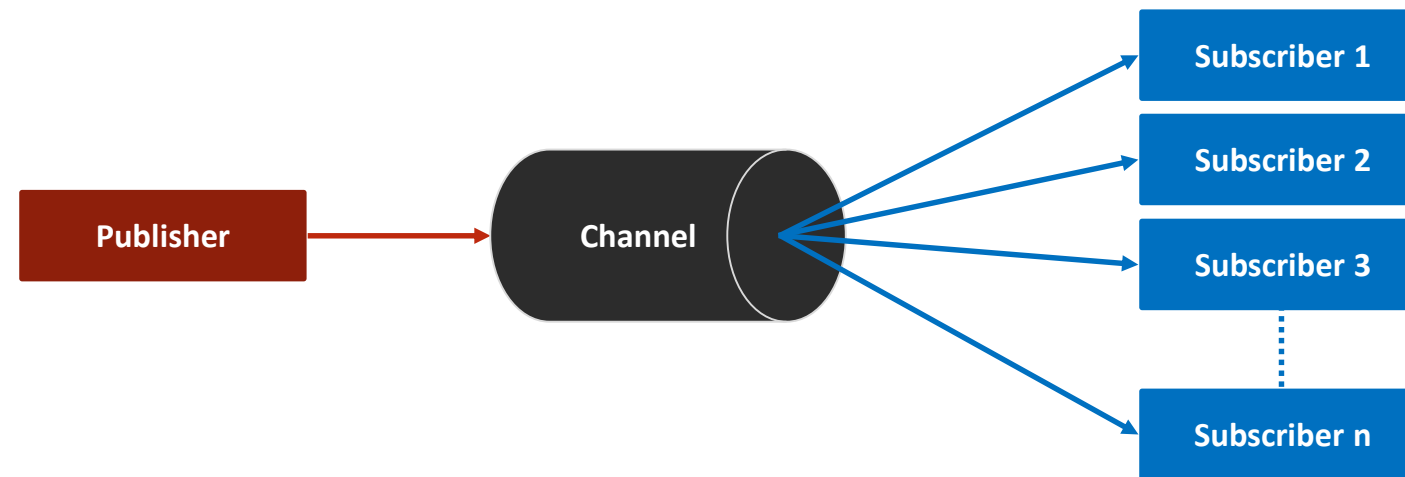
Redis Data Structures



Redis Pub/Sub

Common Ingest Techniques in Redis

Pub/Sub




Commands

Publisher: `publish <channel name> <message>`

Subscriber: `subscribe <channel name>`

List




Commands

Publisher: `lpush <list name> <message>`

Subscriber: `brpop <list name> <timeout>`

Sorted Set



Commands

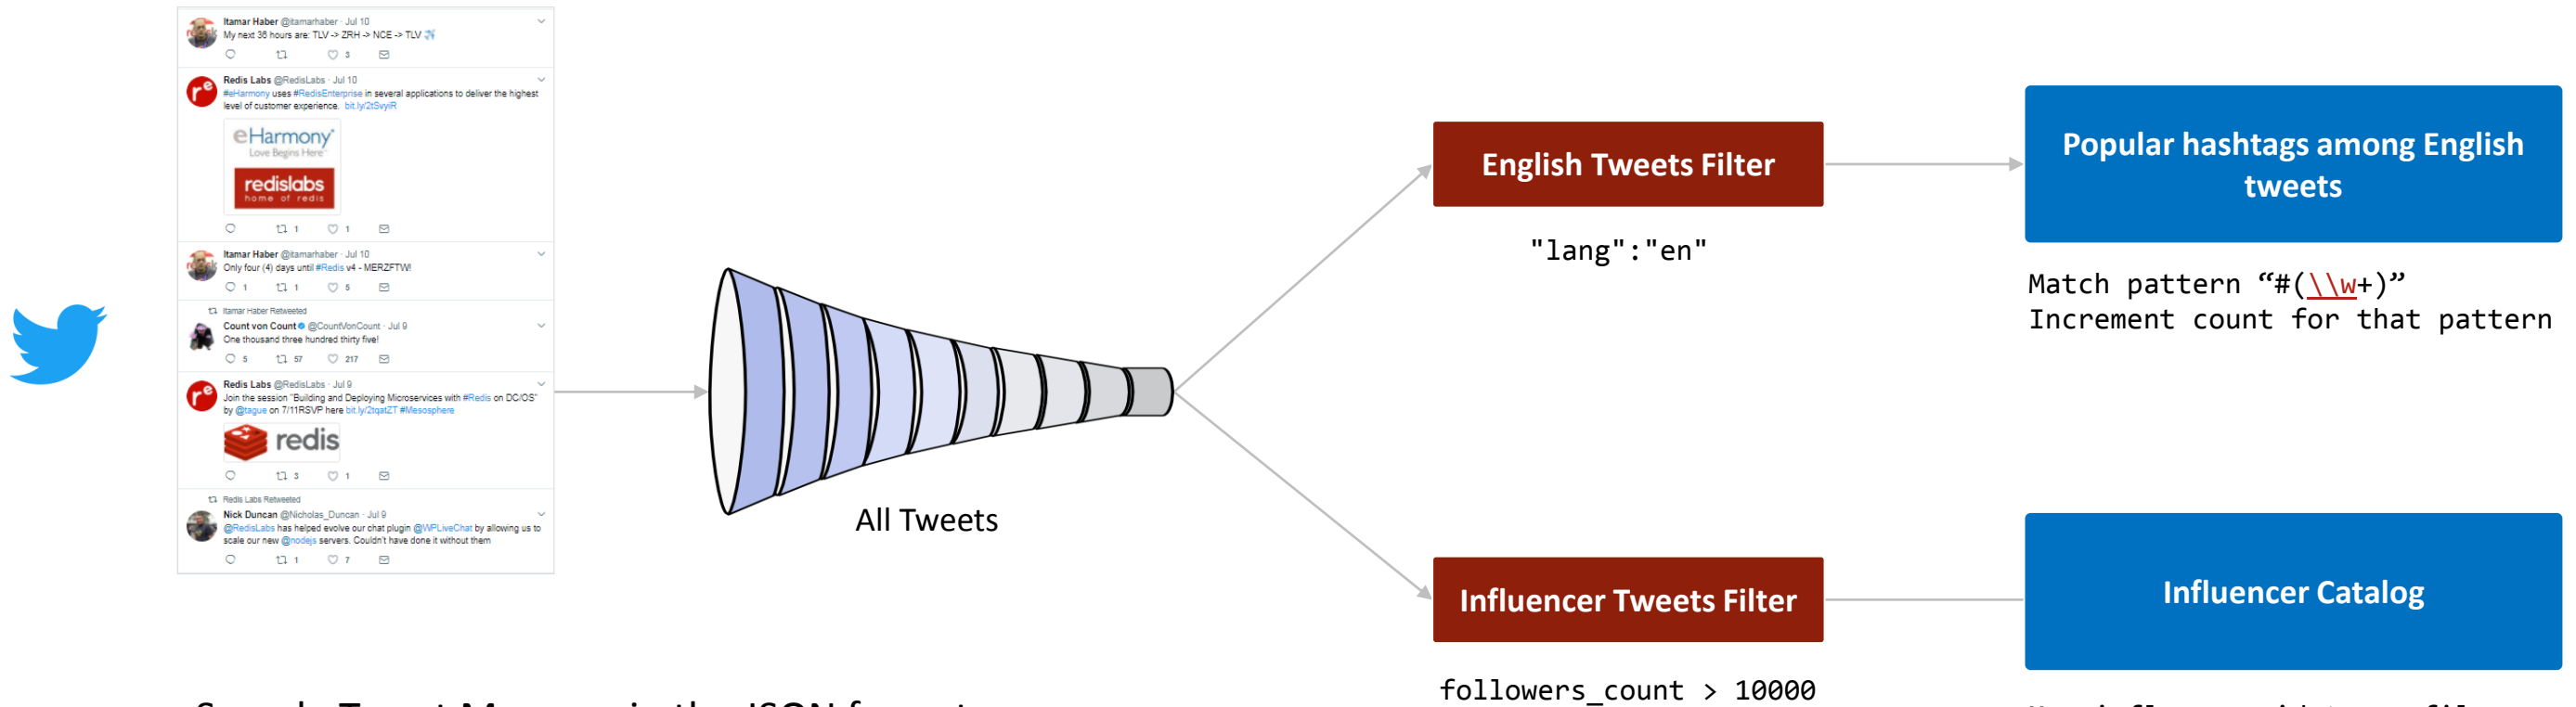
Publisher: `zadd <timeseries name> <timestamp> <message>`

Subscriber: `zrangebyscore <timeseries name> <last timestamp> <current timestamp> WITHSCORES`

The Demo



Demo: Problem Description



Sample Tweet Message in the JSON format:

```
{
  "created_at": "Tue Jul 11 17:06:03 +0000 2017",
  "id": 884821096440004600,
  "text": "USGS reports a M2 #earthquake 31km WSW of Enterprise, Utah on 7/11/17 @ 17:01:53 UTC https://t.co/xXQH2Mfy93 #quake",
  "user": {
    "id": 1414684496,
    "name": "Every Earthquake",
    "screen_name": "everyEarthquake",
    "location": "Earth",
    "followers_count": 18978,
    "friends_count": 17,
    "lang": "en"
  }
}
```

Demo Setup

PubNub

Service Provider for Messages



Programming Language for the demo



IDE



Redis container on Docker

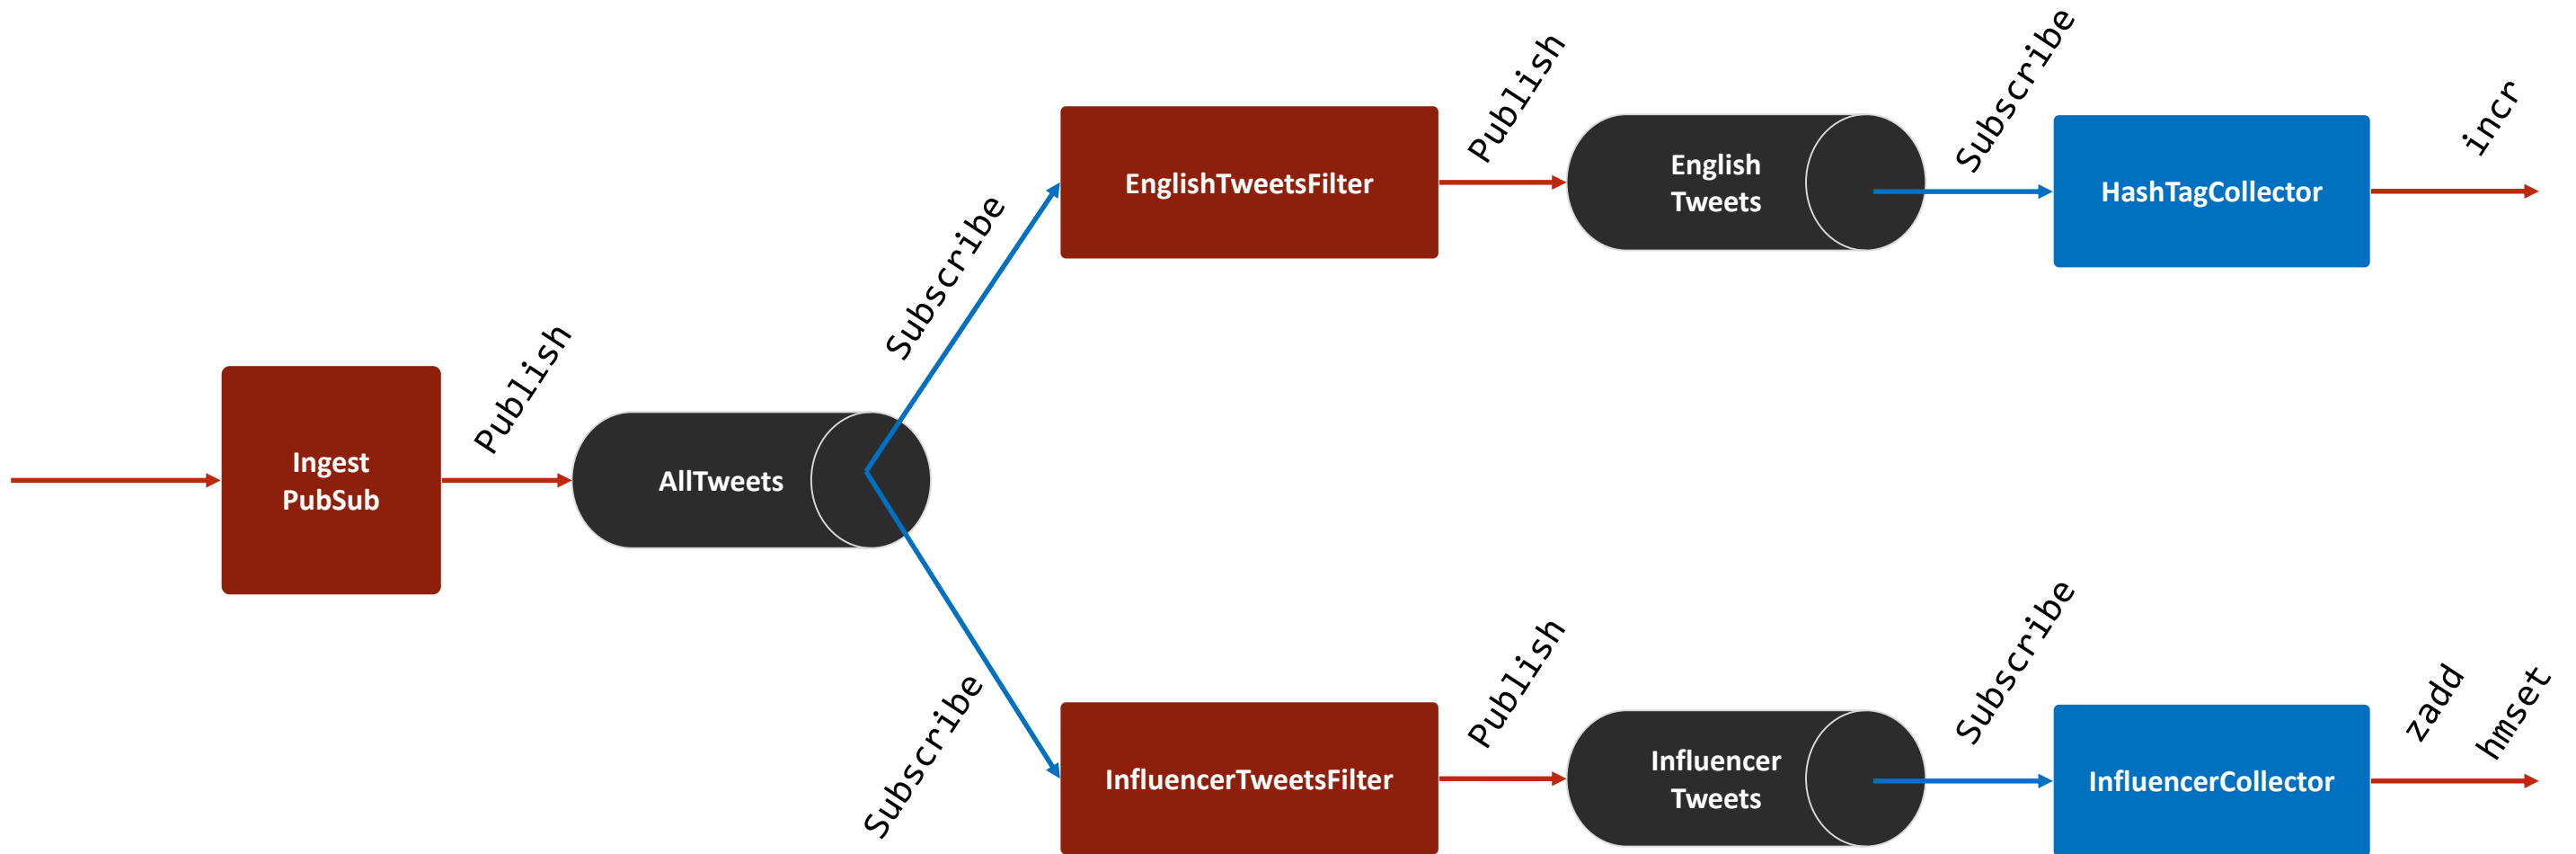


The Three Data Ingest Techniques

Fast Data Ingest Technique	Pros	Cons
Pub/Sub	<ul style="list-style-type: none">• Easy• Decoupled setup• Good for geographically distributed setup	<ul style="list-style-type: none">• Not resilient to connection loss• Requires many connections
Lists	<ul style="list-style-type: none">• Easy• Resilient to connection loss	<ul style="list-style-type: none">• Tightly coupled producers and consumers• Data duplication
Sorted Sets	<ul style="list-style-type: none">• Resilient to connection loss• Least chance of losing data• Access to historical data• Loosely coupled producers and consumers	<ul style="list-style-type: none">• Consumes space• Complex logic

Technique 1: Fast Data Ingest with Pub/Sub

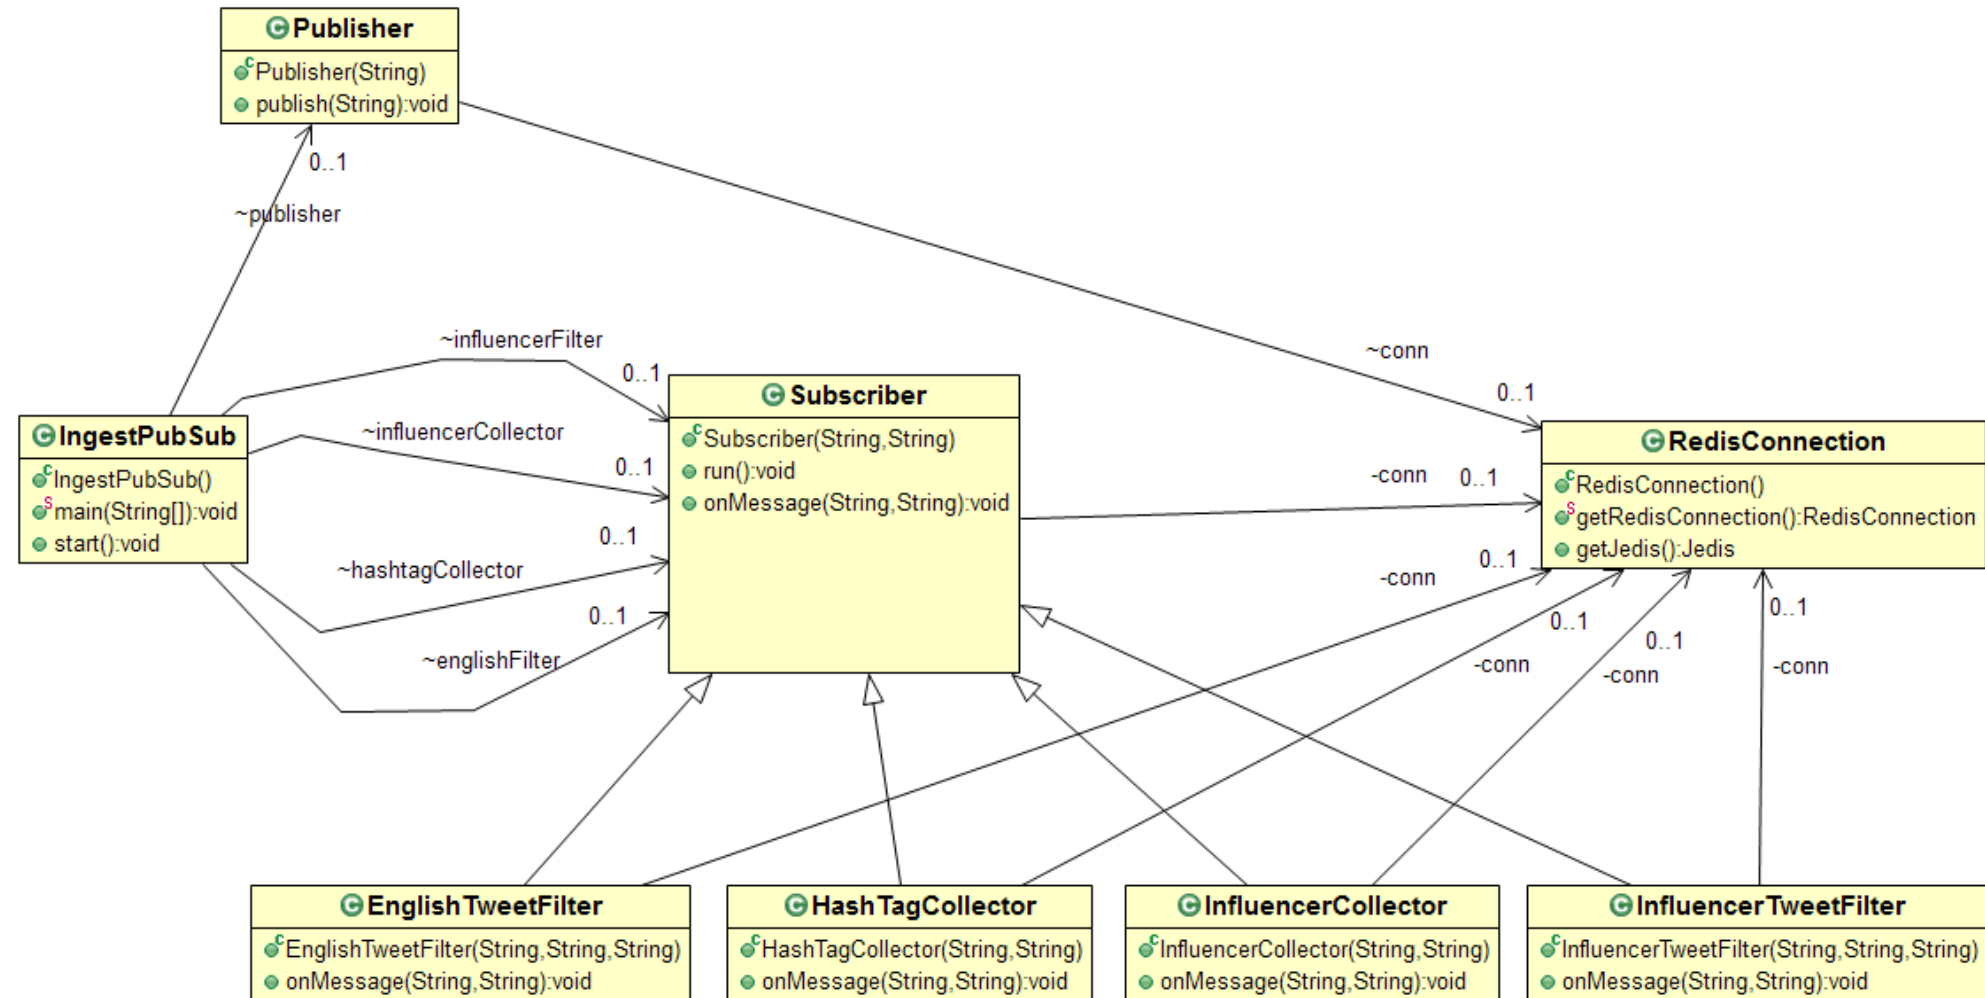
Fast Data Ingest with Pub/Sub



Advantages

- Easy
- Decoupled setup
- Good for geographically distributed setup

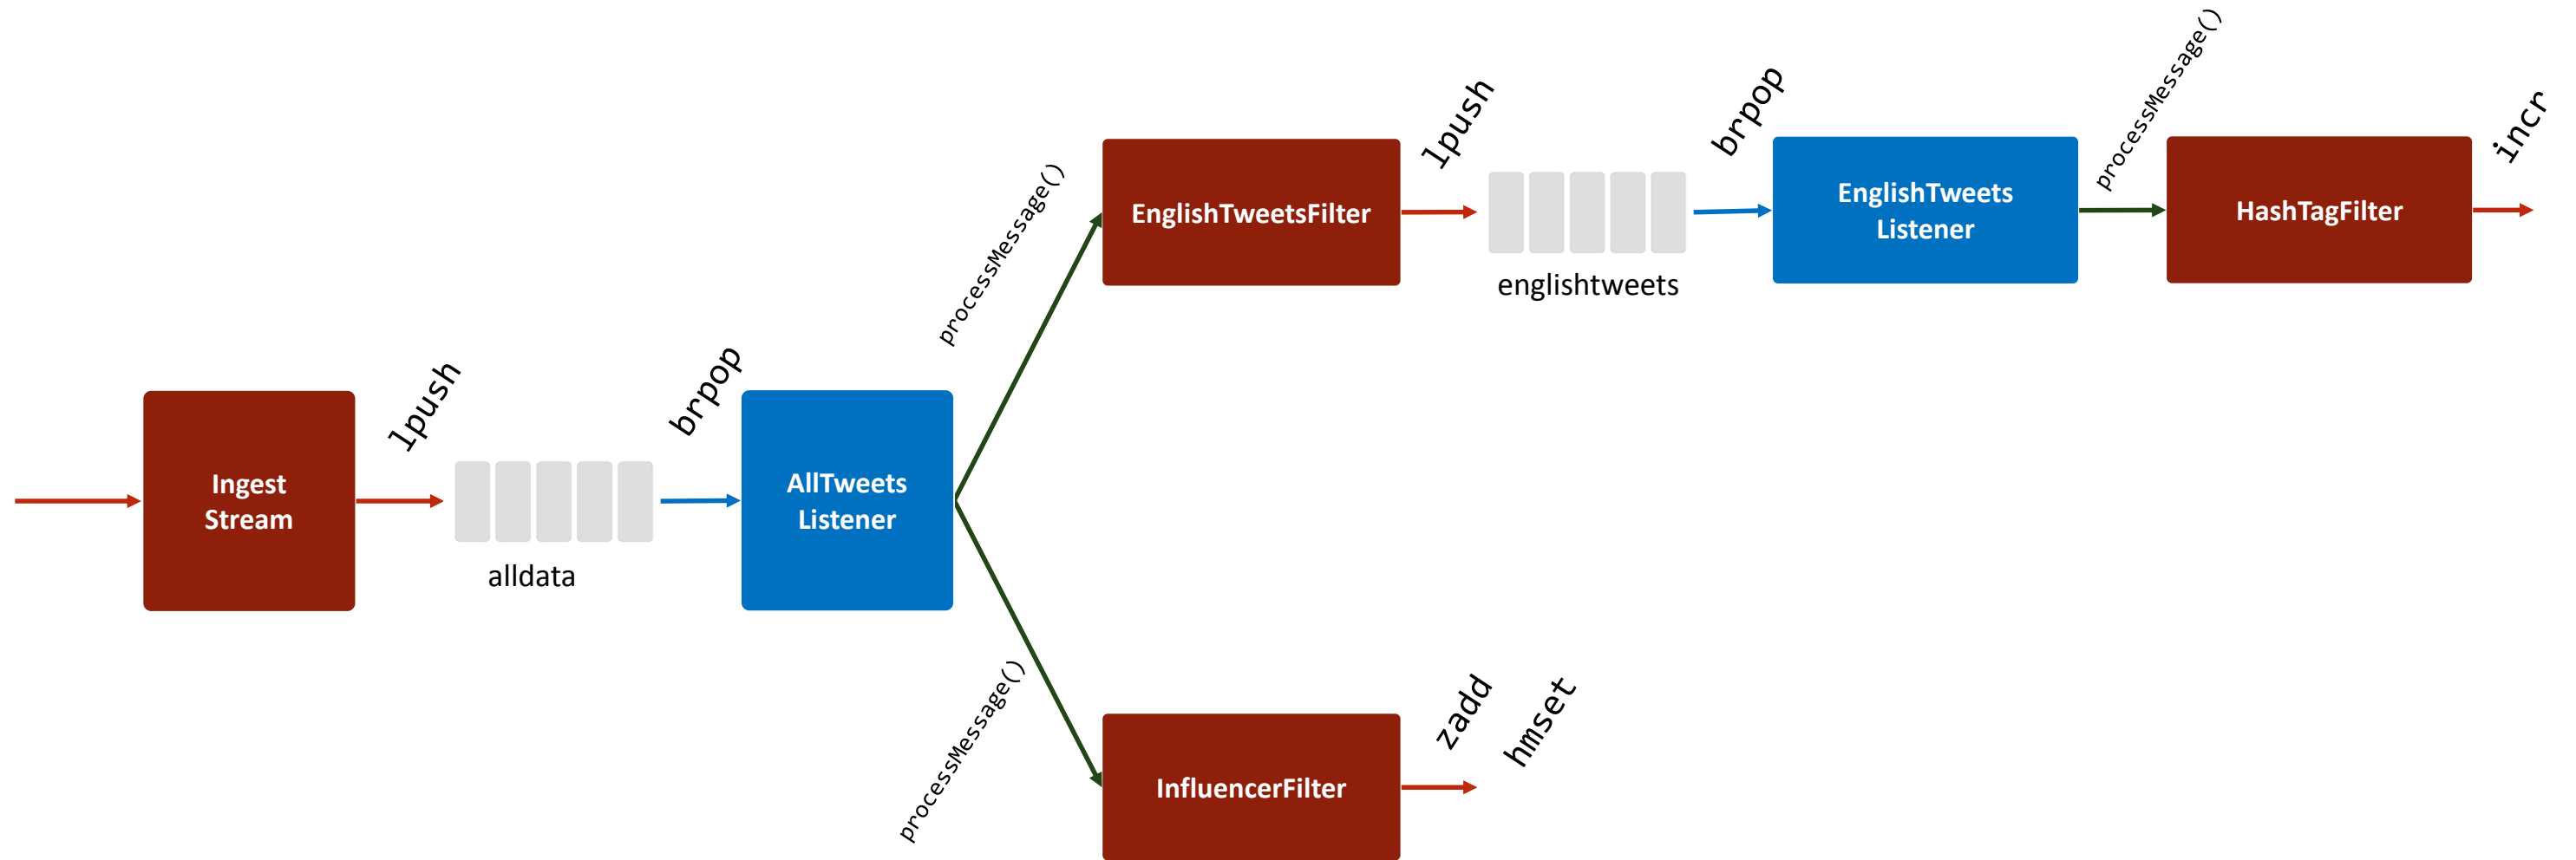
Class Diagrams and Sample Code



<https://github.com/redislabsdemo/IngestPubSub>

Technique 2: Fast Data Ingest with Lists

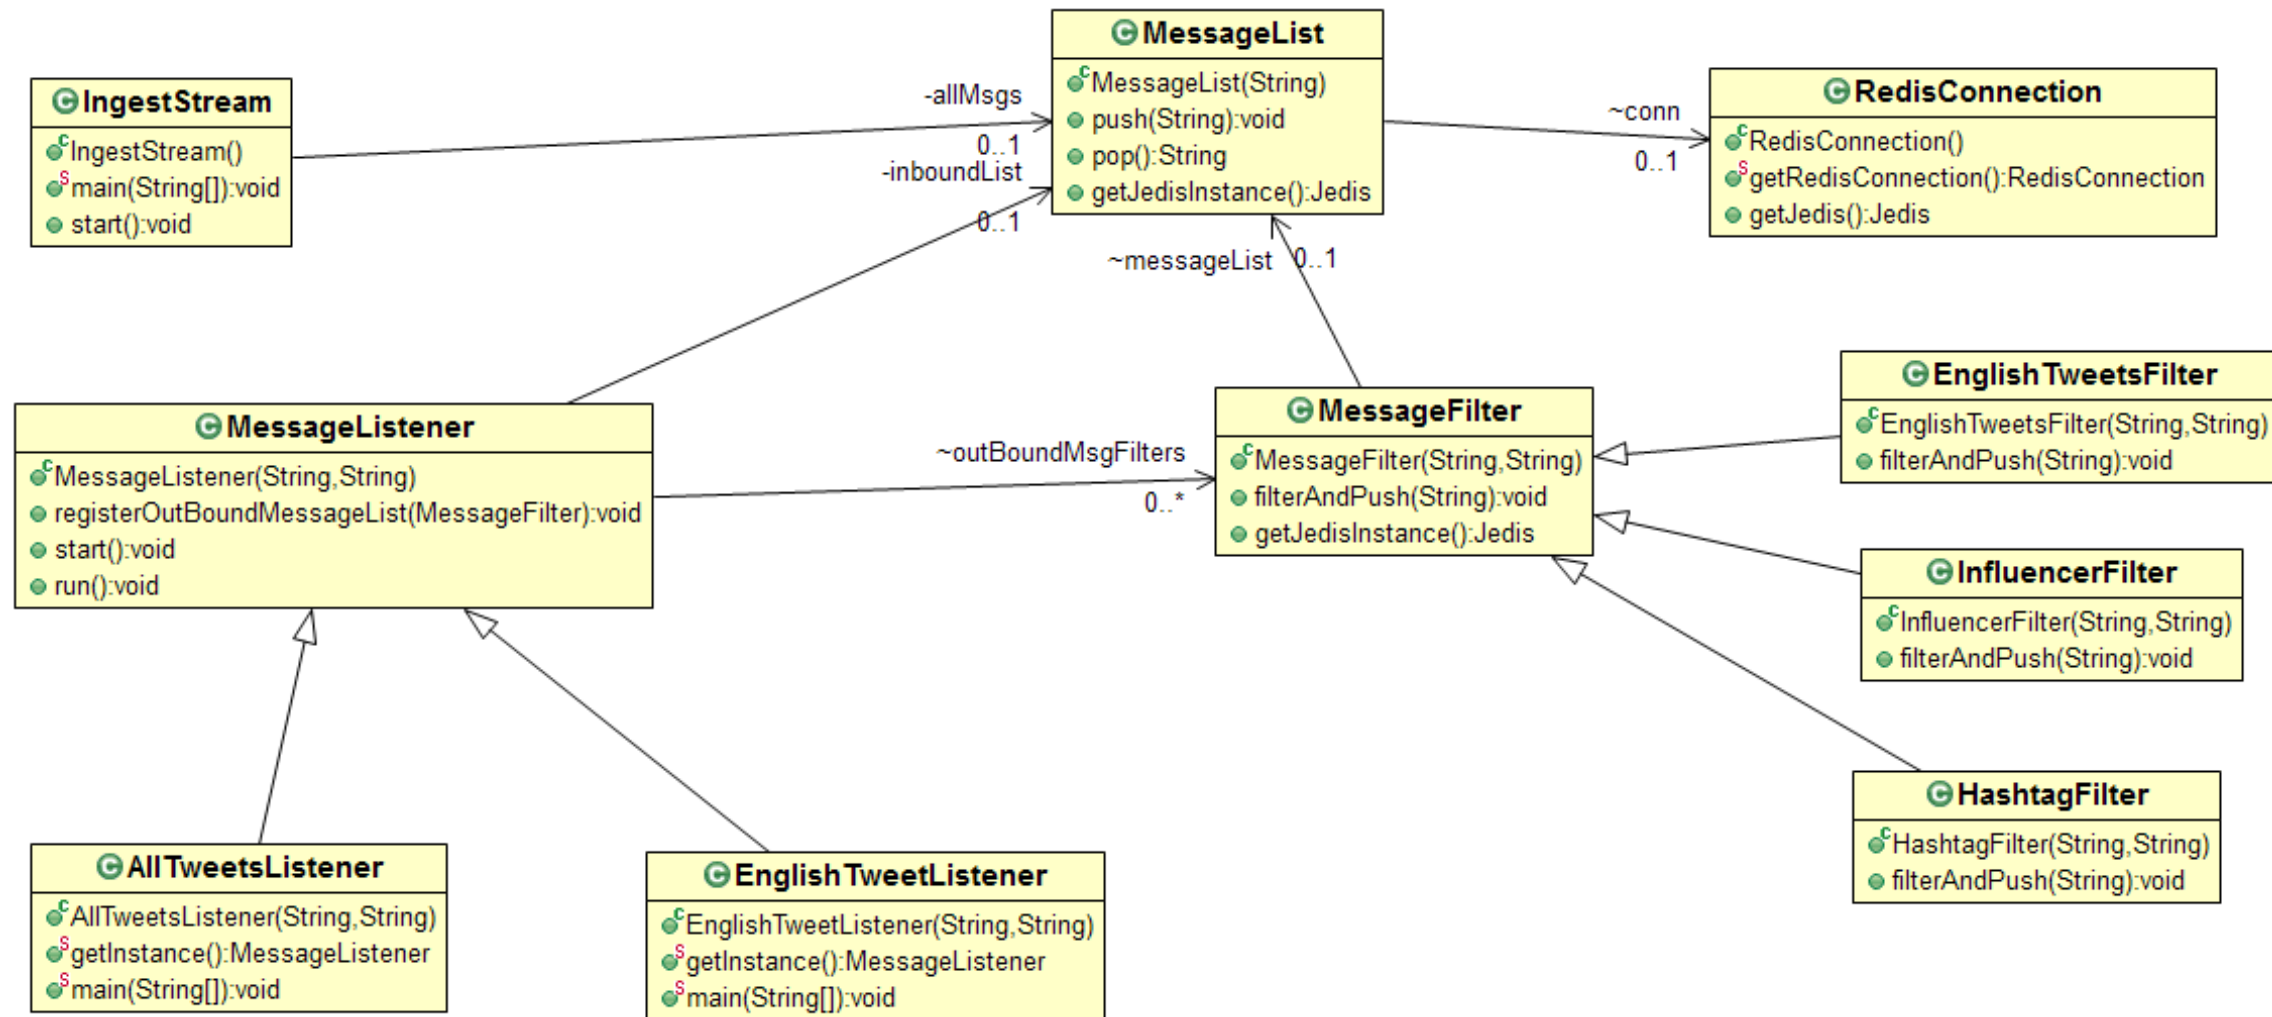
Fast Data Ingest with Lists



Advantages

- Easy
- Resilient to connection loss

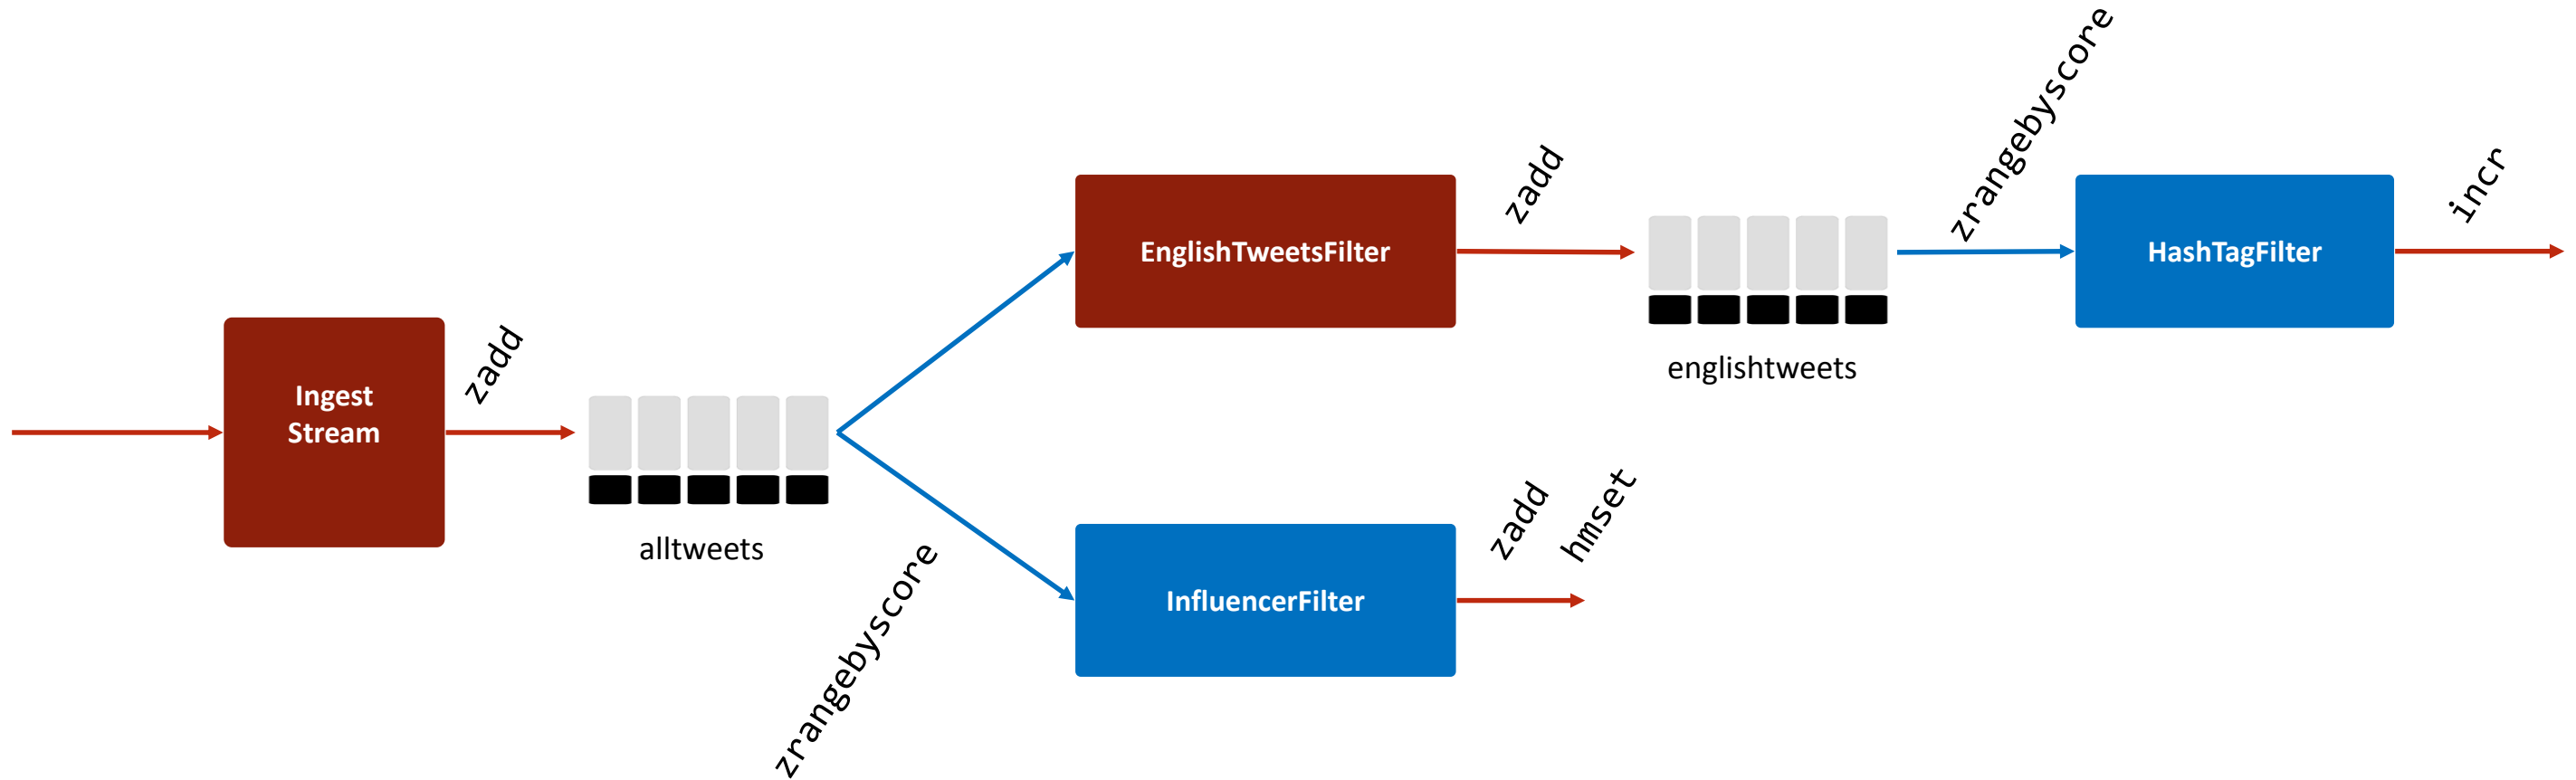
Class Diagrams and Sample Code



<https://github.com/redislabsdemo/IngestList>

Technique 3: Fast Data Ingest with Sorted Sets

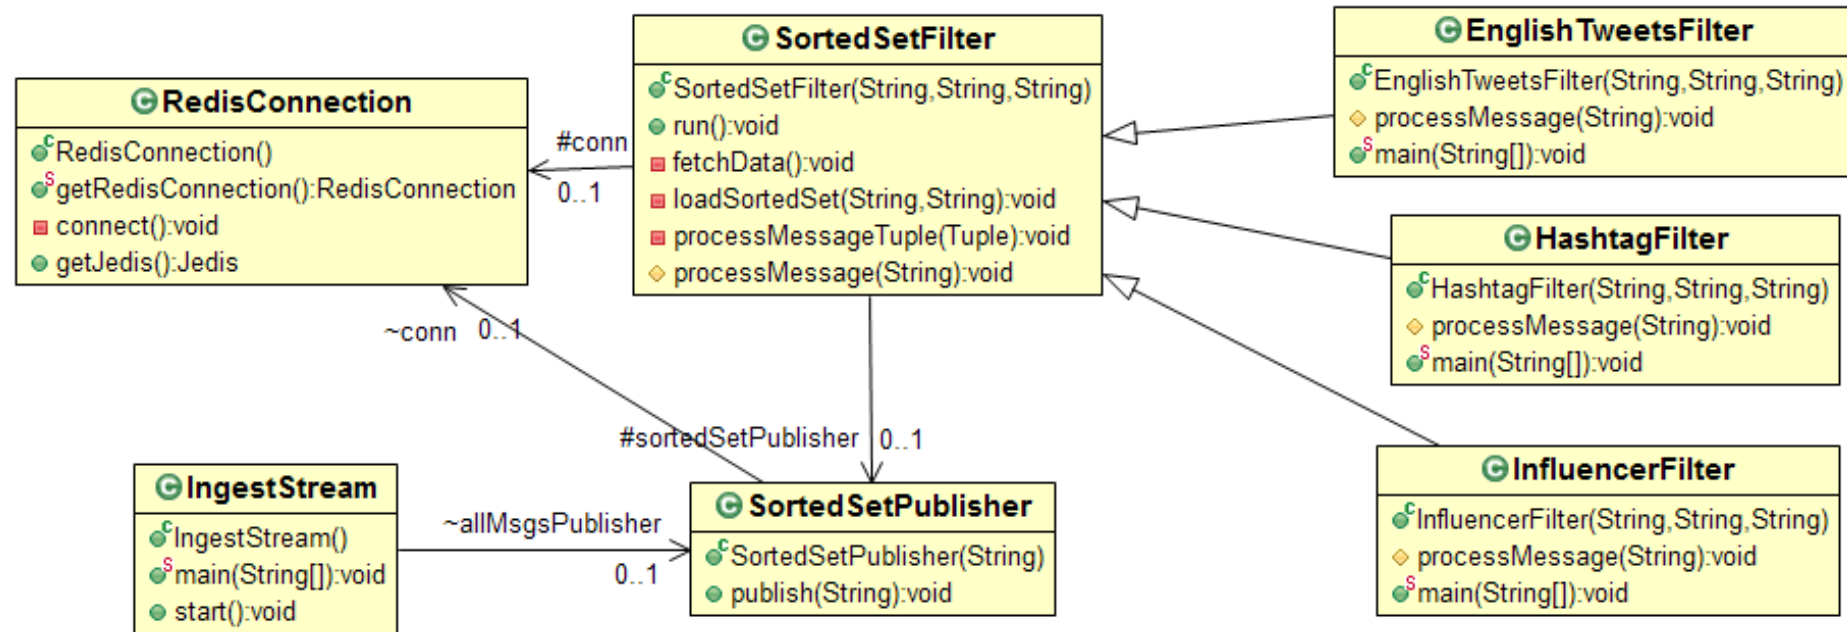
Fast Data Ingest with Sorted Sets



Advantages

- Resilient to connection loss
- Least chance of losing data
- Access to historical data
- Loosely coupled producers and consumers

Class Diagrams and Sample Code



<https://github.com/redislabsdemo/IngestSortedSet>

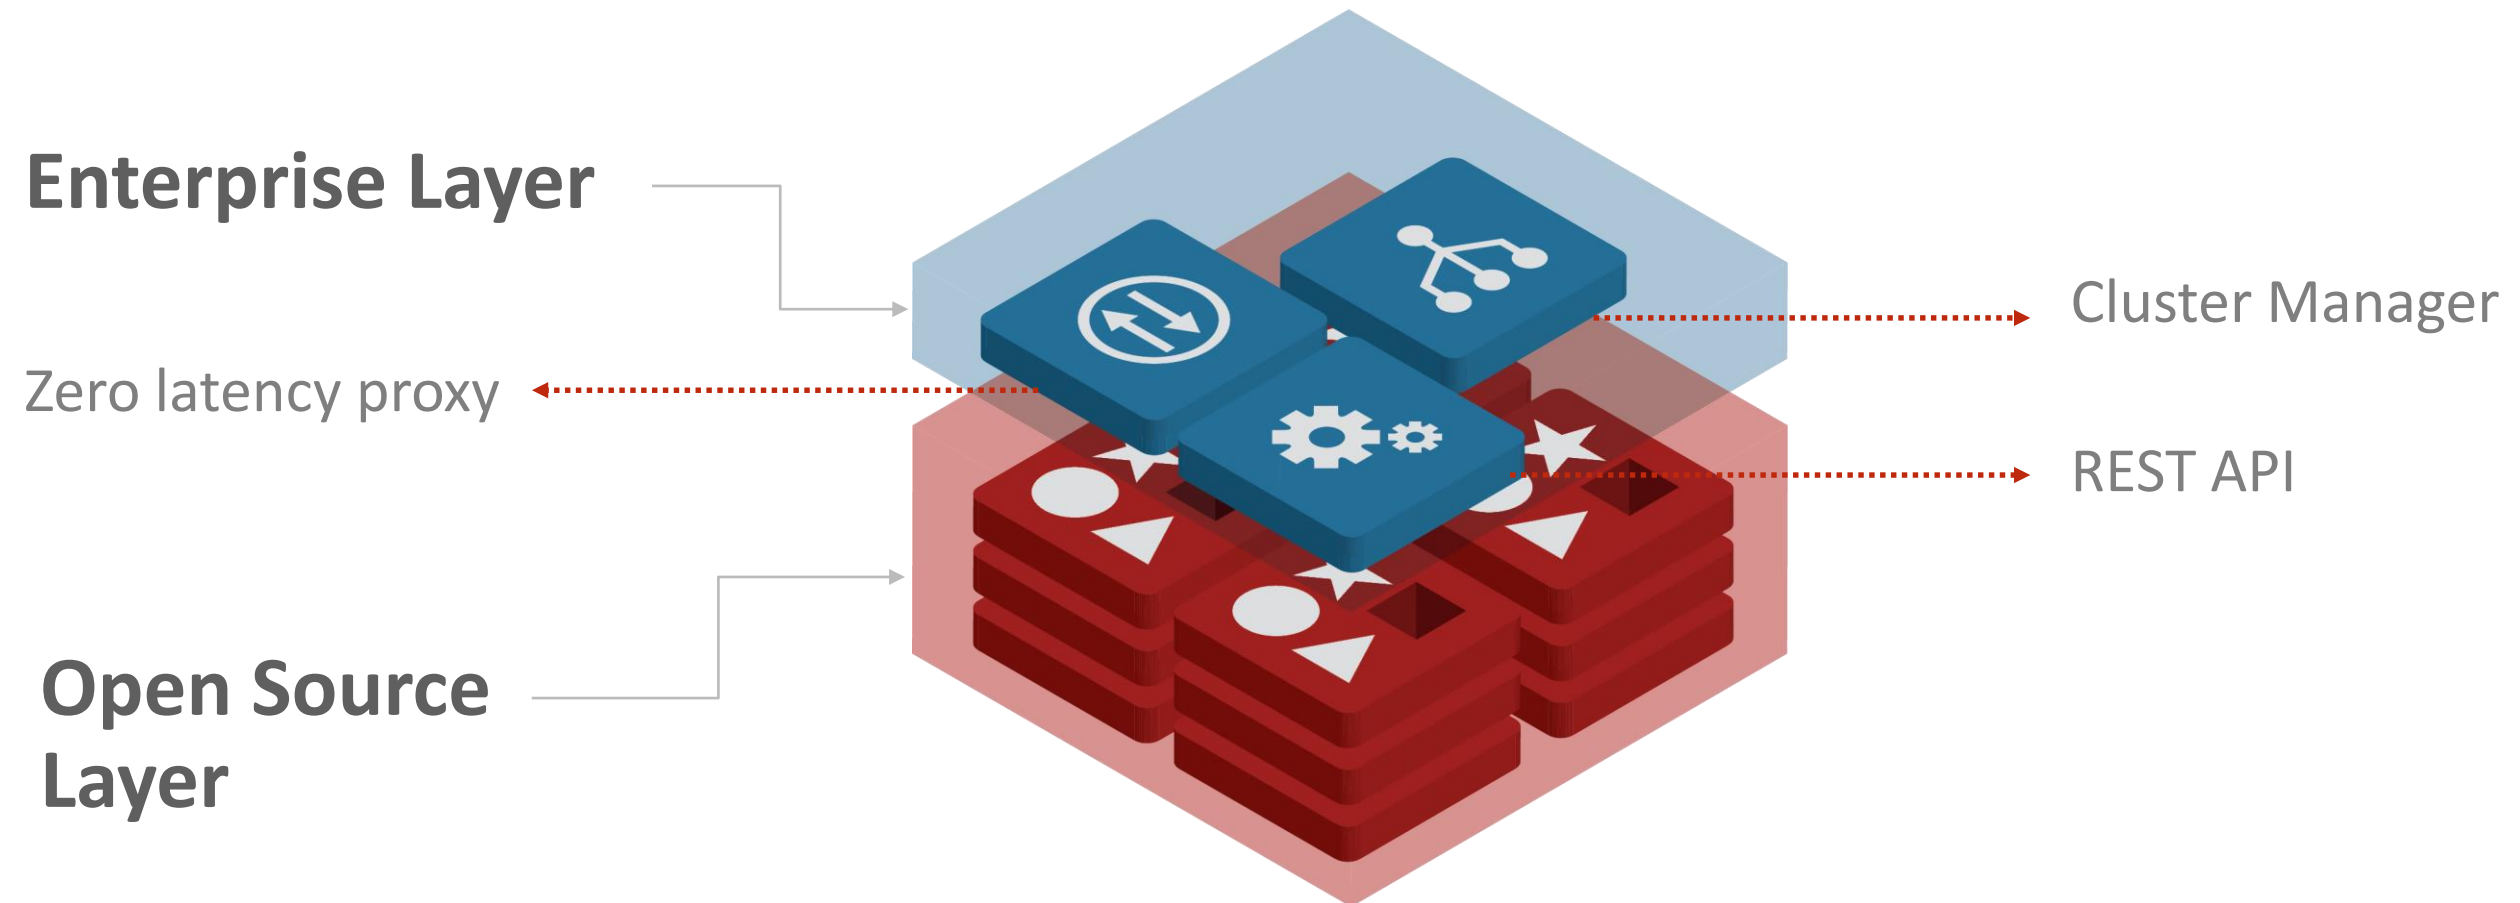
Redis^e for Fast Data Ingest




Redis Database Instances



Redis^e Technology



Redis^e Technology



Redis^e Technology


Redis^e Cluster

- Shared nothing cluster architecture
- Fully compatible with open source commands & data structures



Redis^e - Shared Nothing Symmetric Architecture

Distributed Proxies
Single or Multiple Endpoints



Unique multi-tenant “Docker” like architecture enables running hundreds of databases over a single, average cloud instance without performance degradation and with maximum security provisions

Redis^e Benefits for Data Ingest

Effortless Scaling



Simple, Seamless Clustering. Linear scaling



ACID Compliance in Cluster Architecture

Always On Availability




Instant Failure Recovery, No Data loss



Stable and Predictable High Performance

Substantially Lower Costs



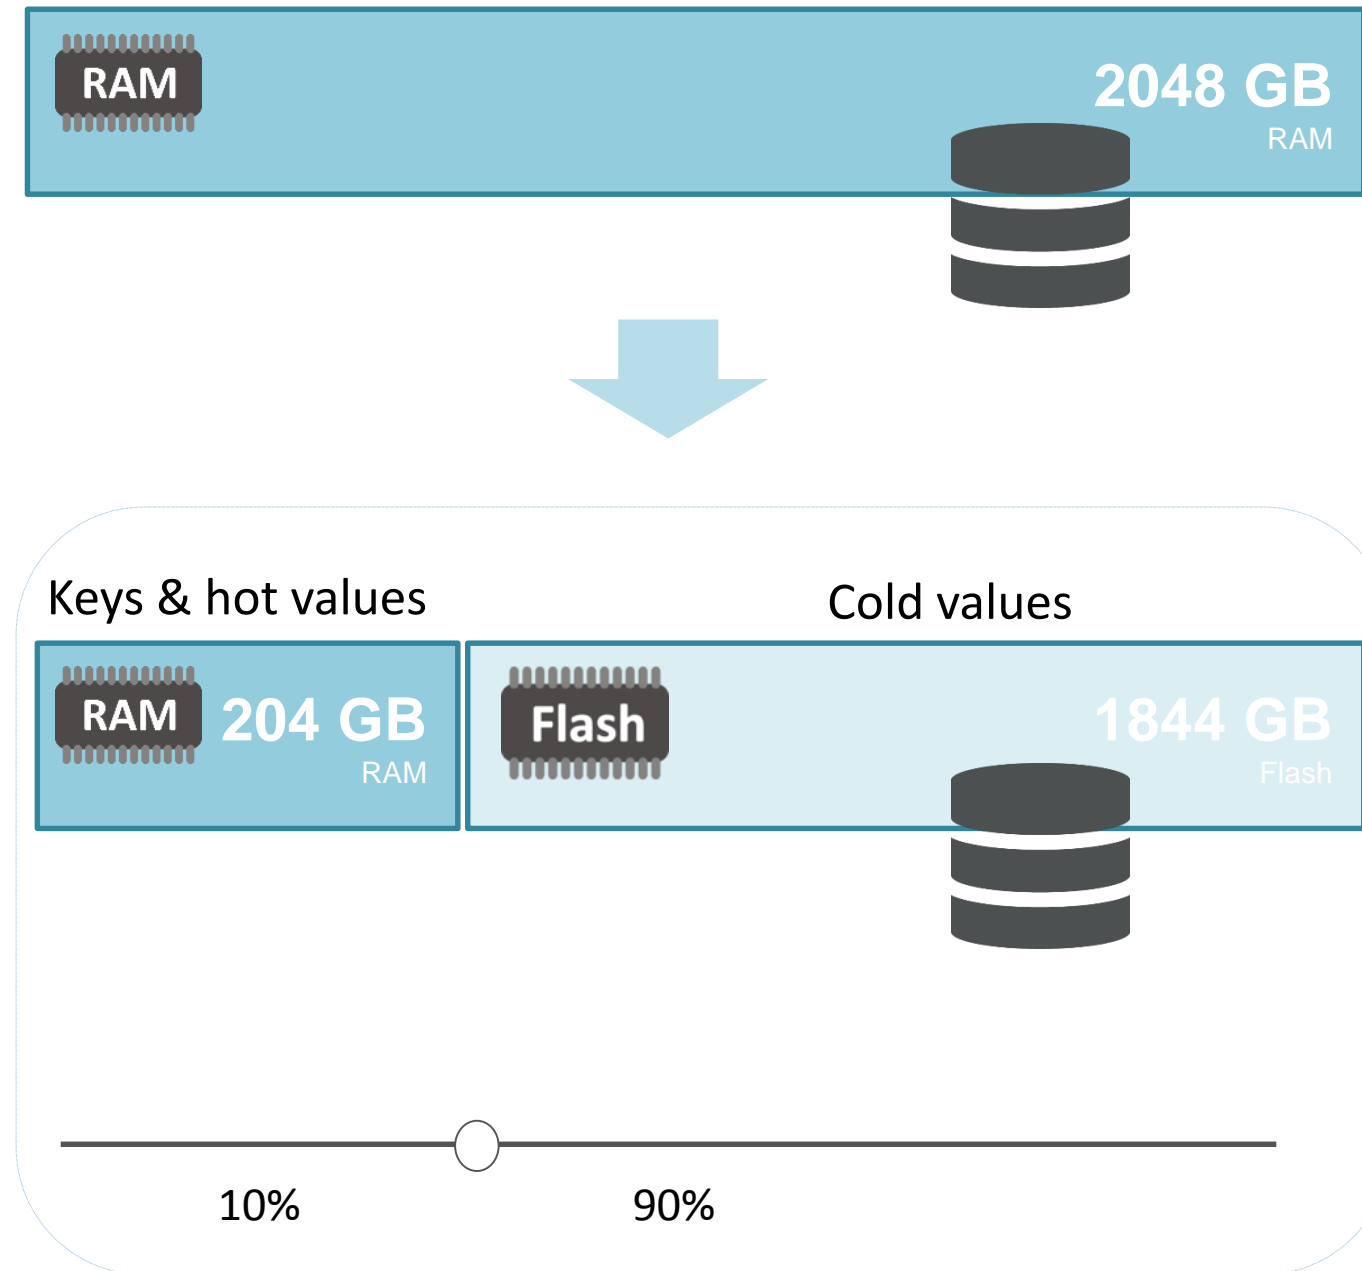
Run on Flash as a RAM extension



Top notch 24x7 expert support

Redis^e Flash

- Near-RAM performance at 70%+ lower costs
- Technology treats Flash as a RAM replacement (or extension)
- RAM/Flash ratio can be easily configured
- Pluggable storage engine
- Available on SATA-based SSD, NVMe-based SSD, NVDIMM like 3D XPoint/SCM on x86 and P8 platforms




Redis^e Flash - 10TB Redis Deployment on EC2

	Redis on RAM	Redis ^e Flash
Dataset size	10 TB	10 TB
Database size with replication	30 TB	20 TB*
AWS instance type	x1.32xlarge	i3.16xlarge
Actual instance size (RAM, and RAM+Flash)	1.46 TB	3.66 TB
# of instances needed	21	6 + 1 (for quorum)
Persistent Storage (EBS)	154 TB	110 TB
1 year cost (reserved instances)	\$1,595,643	\$298,896
Savings	-	81.27%

* Redis Enterprise only needs 1 copy of the data because quorum issues are solved at the node level

Questions



One more thing....

redis.conf setting:

```
client-output-buffer-limit pubsub 32mb 8mb 60
```

With this setting, Redis will force the clients to disconnect under two situations:

- If the output buffer grows more than 32mb
- If the output buffer holds 8mb of data consistently for 60 seconds

Thank You

Roshan Kumar

 roshan@redislabs.com
 @roshankumar

Redis Labs

 expert@redislabs.com
 @redislabs