AWS
re:Invent

DAT374-S

Using Redis beyond caching

Alvin Richards
Chief Product Officer
Redis Labs

Brandon Mahoney
Enterprise Architect
Alliance Data

Junaid Fakhruddin
Director of IT
Gap

Bhilhanan Jeyaram
SW Engineer
Gap

Agenda

History of Redis

Redis for Caching at Alliance Data

Redis for Inventory Management at Gap

Redis for Modern Data Models

redislabs
HOME OF REDIS

History of Redis

The Redis Project – 11 Year Retrospective

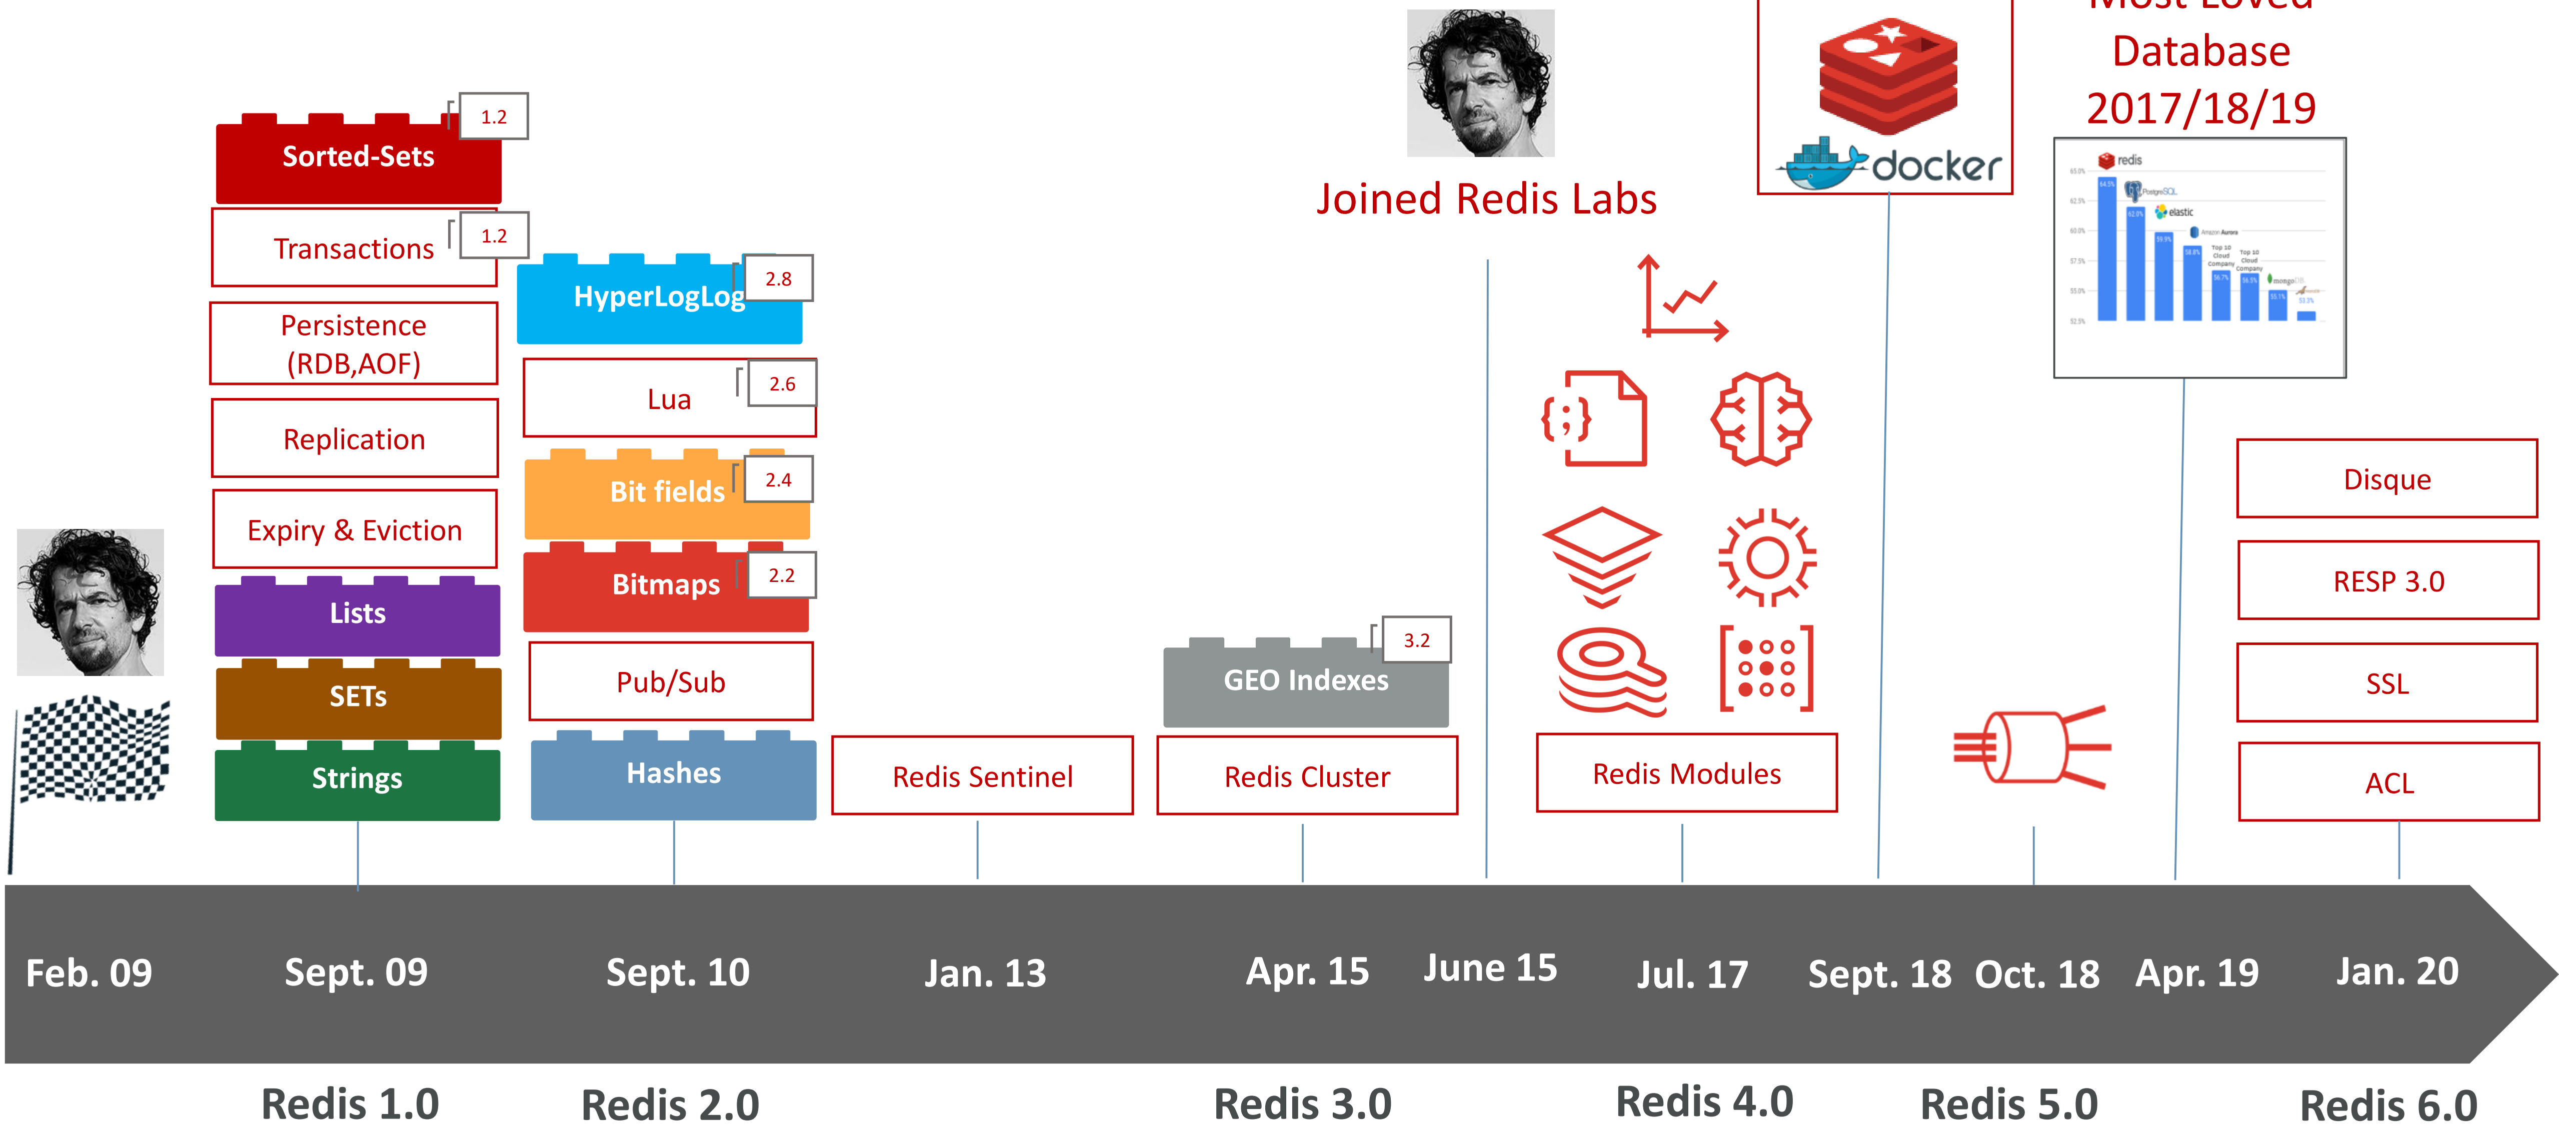

Sorted-Sets 1.2

Transactions 1.2

Persistence (RDB,AOF)

Replication

Expiry & Eviction

Lists

SETs

Strings

HyperLogLog 2.8

Lua 2.6

Bit fields 2.4

Bitmaps 2.2

Pub/Sub

Hashes

Redis Sentinel

GEO Indexes 3.2

Redis Cluster

Joined Redis Labs

Redis Modules

1B

Most Loved Database 2017/18/19

Disque

RESP 3.0

SSL

ACL

Feb. 09 | Sept. 09 | Sept. 10 | Jan. 13 | Apr. 15 | June 15 | Jul. 17 | Sept. 18 | Oct. 18 | Apr. 19 | Jan. 20

Redis 1.0

Redis 2.0

Redis 3.0

Redis 4.0

Redis 5.0

Redis 6.0

redislabs HOME OF REDIS

How Popular Is Redis?

MOST LAUNCHED

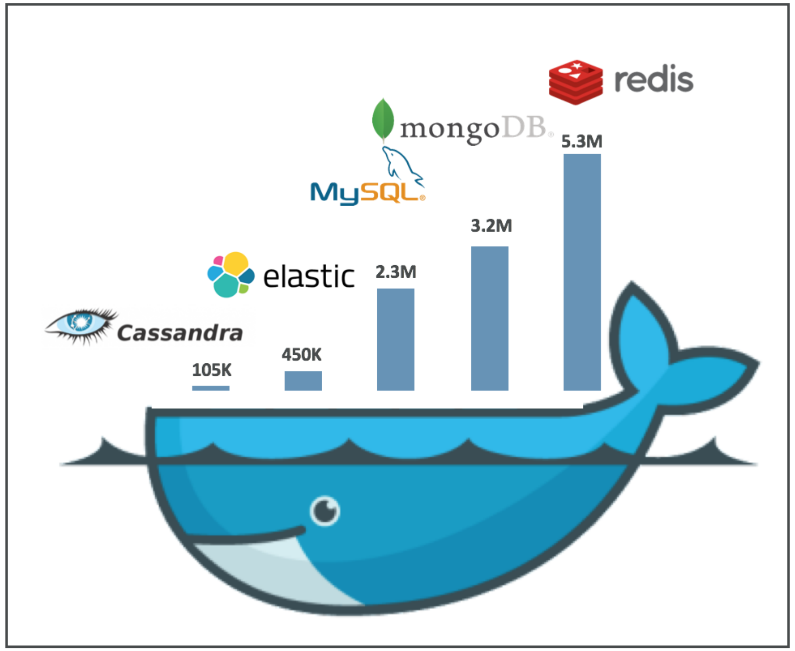

Launches/day: Docker Hub, Nov. 2019

MOST USED

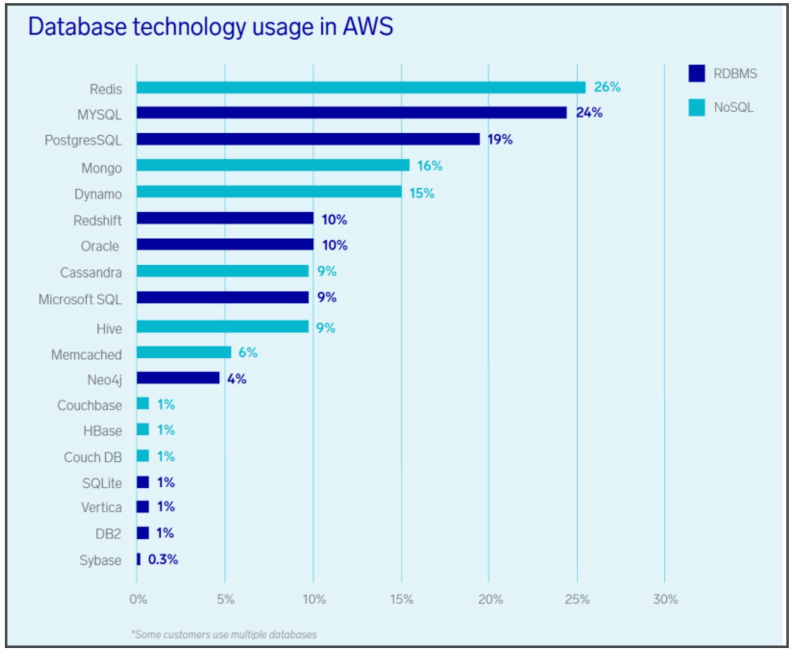

Sumo Logic, Sept. 2019

MOST LOVED

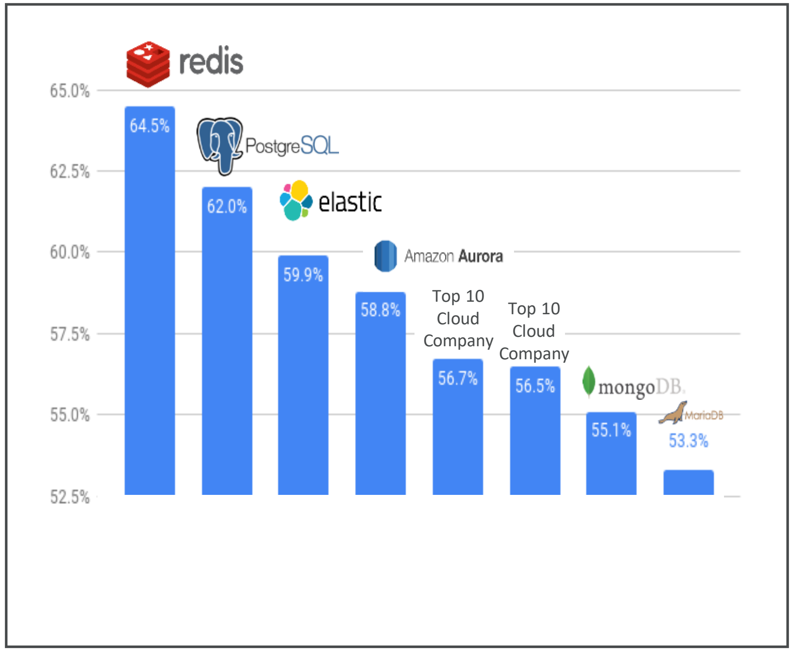

Stack Overflow, 2017, 2018, 2019

Why Redis?

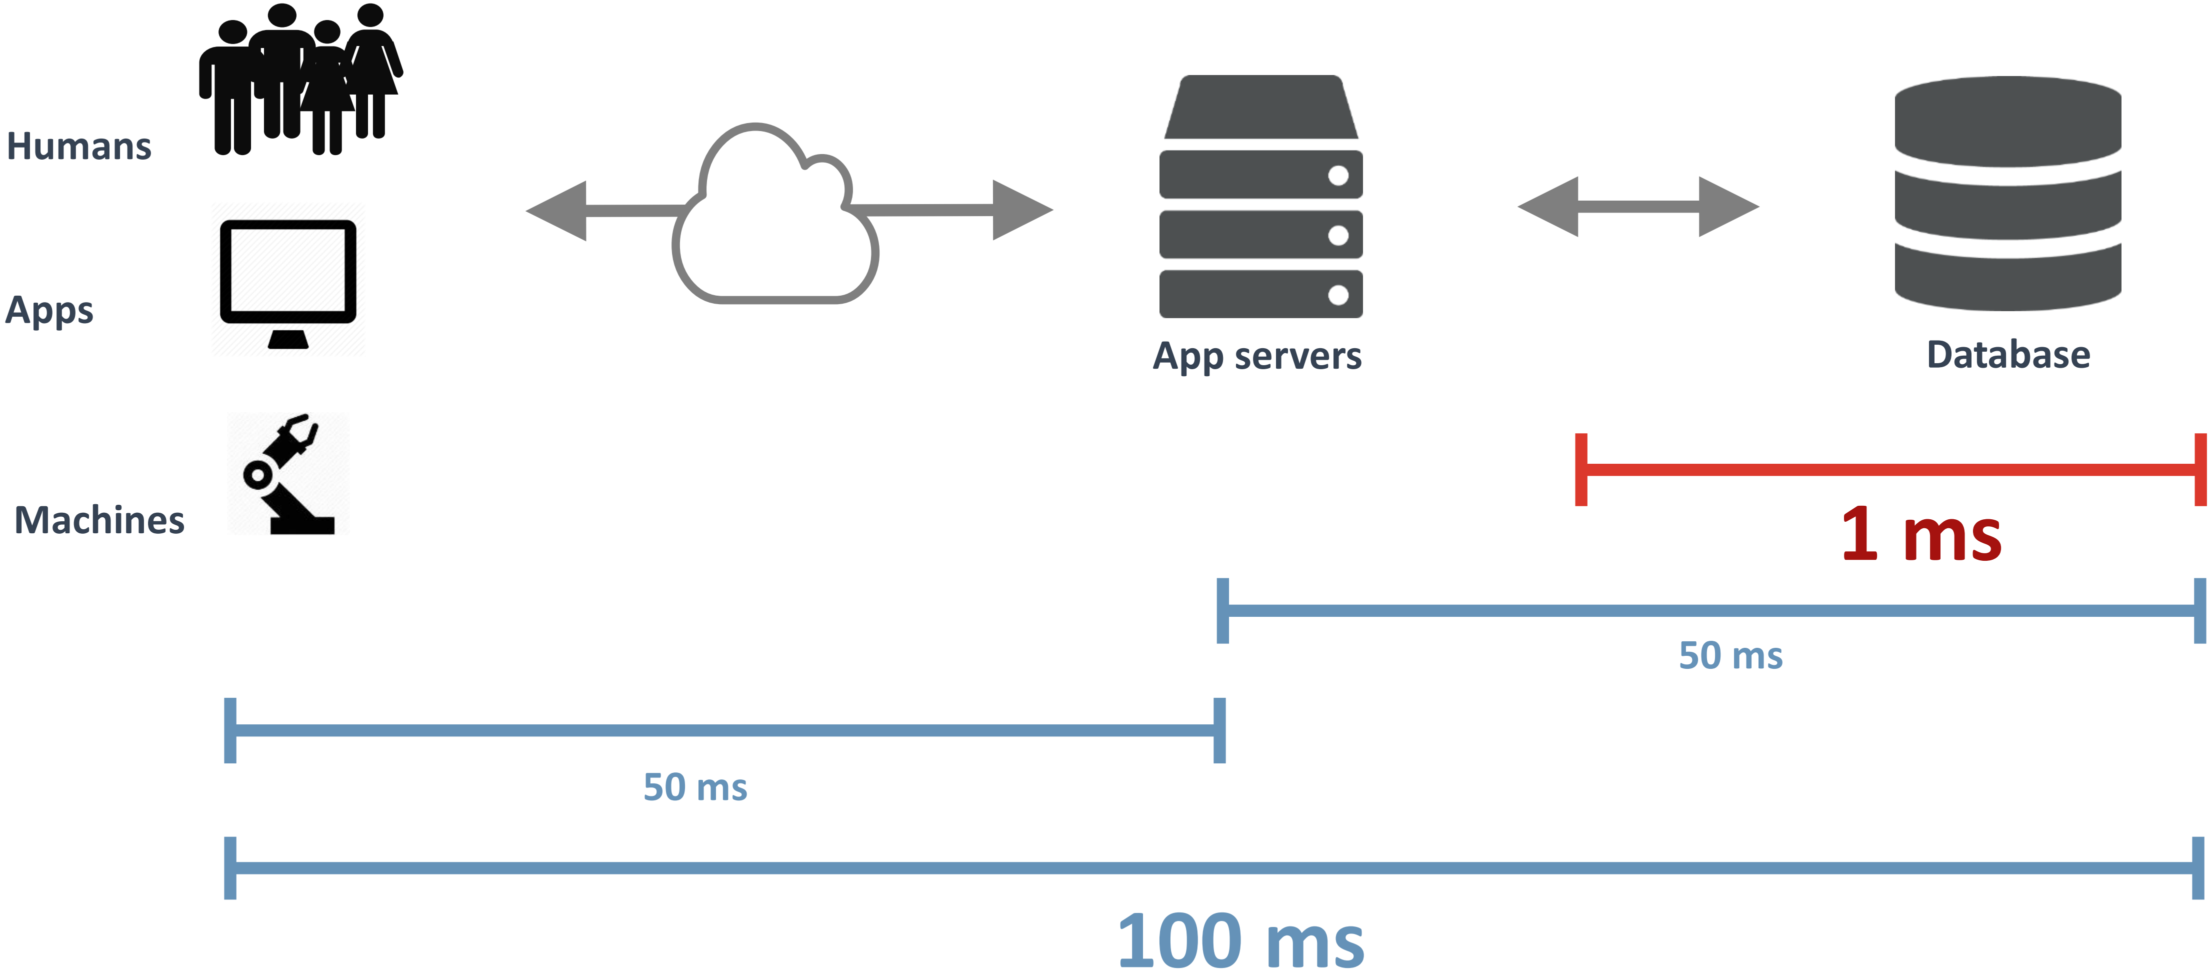

redislabs
HOME OF REDIS

Why Redis?

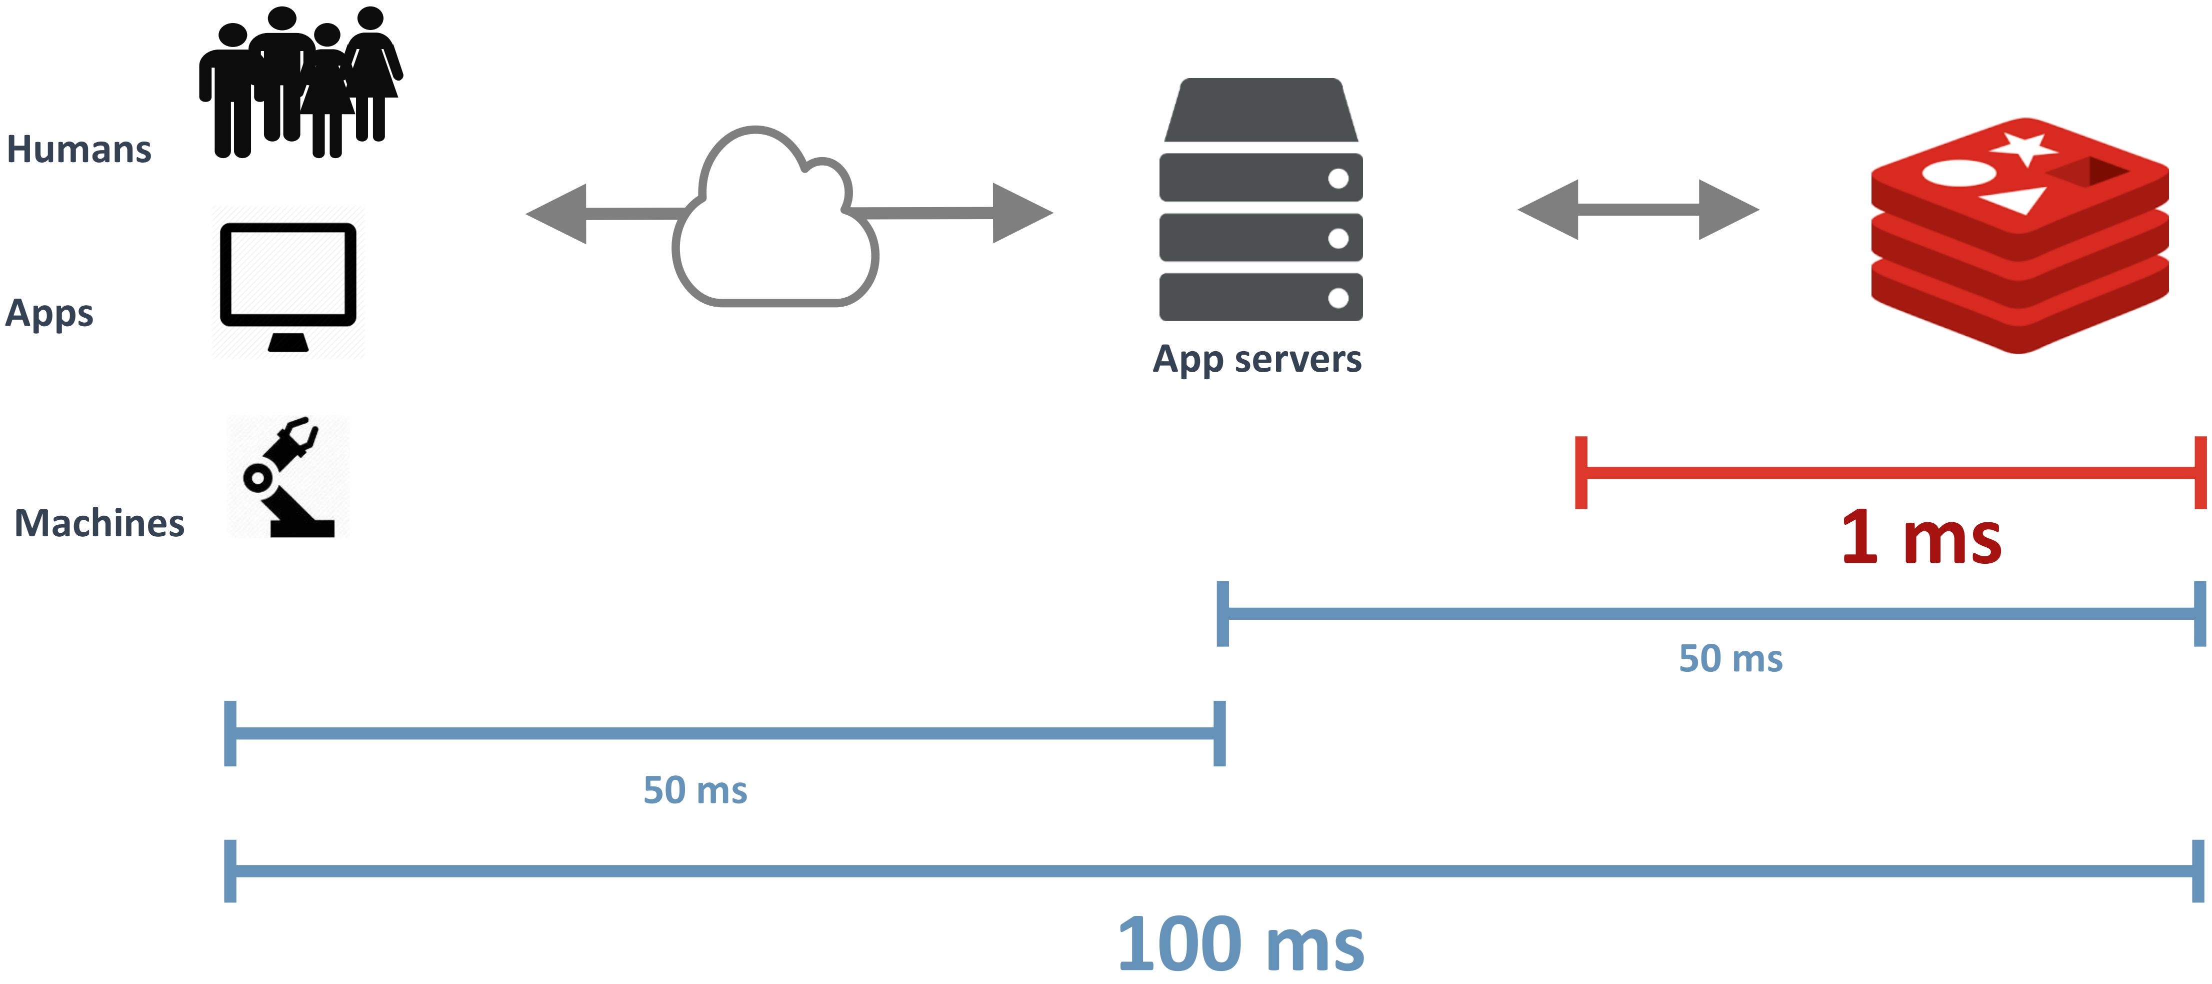

Humans

Apps

Machines

App servers

1 ms

50 ms

50 ms

100 ms

AllianceData.

Achieving More through Microservices with Redis Enterprise

Agenda

- Introduction
- Problem
- Solution
- Results

- Leading tailored loyalty and marketing solutions
- Manage more than **365 branded credit programs**
- Customer-tailored marketing communication solutions
- Handle end-to-end acquisition to create a **seamless experience** for consumers

AllianceData

Our Family of 365 Brand Partners

IKEA

GIANT EAGLE®

HSN

BARNEYS NEW YORK

BED BATH & BEYOND

Limitedbrands

ANN TAYLOR

EXPRESS

NEW YORK & COMPANY

meijer

PINK
VICTORIA'S SECRET

Crate&Barrel

AllianceData℠
CARD SERVICES

MARATHON

wayfair

GameStop
POWER TO THE PLAYERS

TigerDirect.com

Justice

Virgin america

VIKING

DSW
DESIGNER SHOE WAREHOUSE®

TOYOTA

ZALES
THE DIAMOND STORE®

VICTORIA'S SECRET

AllianceData

Current State of Business

- Content management system **SDL Tridion**
- Repository for millions of individual content pieces for each client

Challenges

- **Staying competitive in the market**
  - Many fintech companies are challenging the industry with innovation
  - Staying nimble and leading the industry

Need

- A **high-performance**, **scalable**, and **stable** data platform/engine that would **streamline** our publishing system

Redis Enterprise within Microservices

- Content Management system is 10 years old, original architecture was **monolithic**
- Recently transitioned to a **microservice** architecture

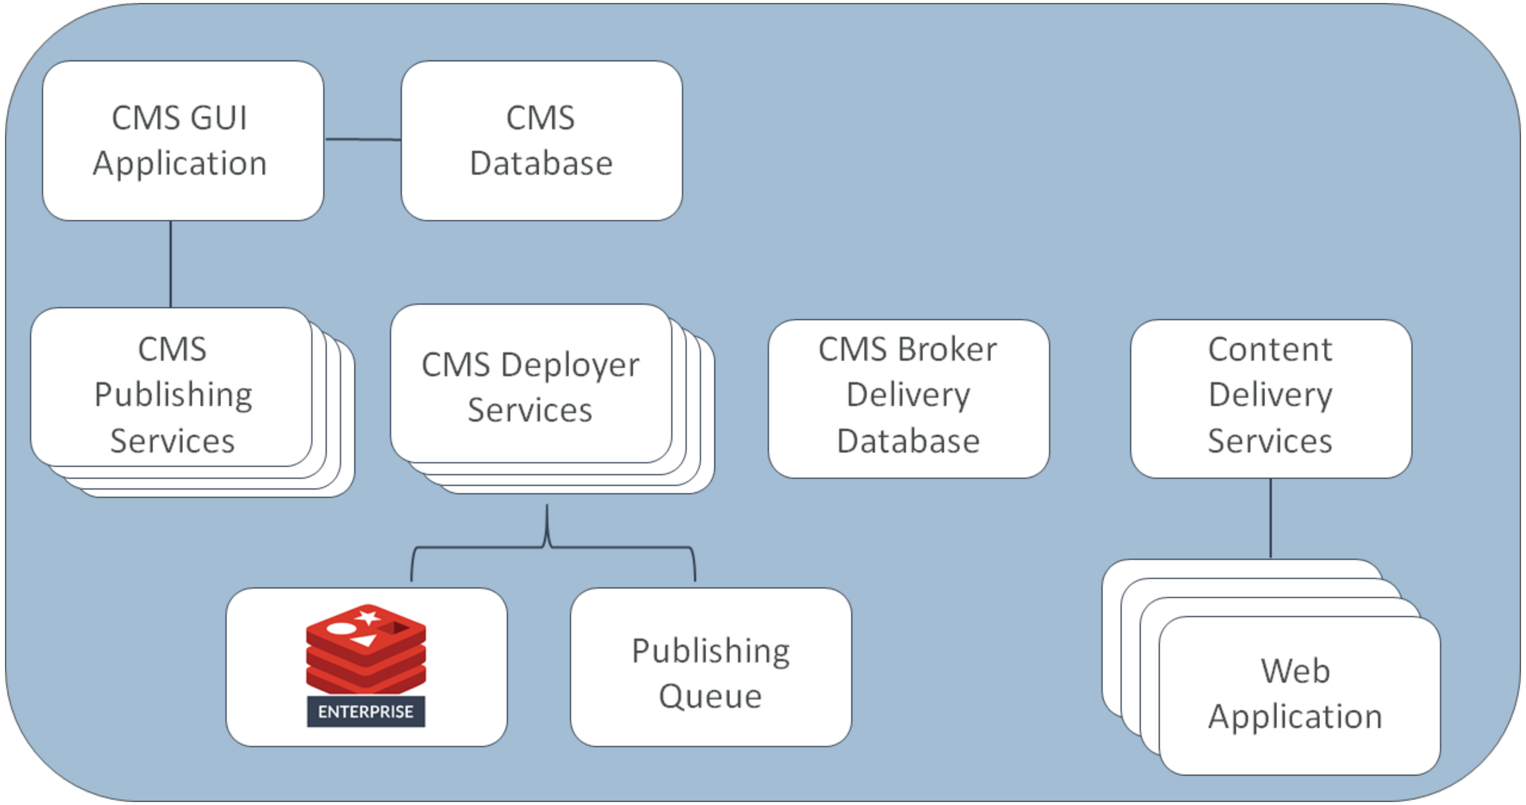

AllianceData

New Architecture with Redis

- We're pushing **200K** publishing transactions per hour; Redis is using **~1,000 Ops/S** with incoming traffic **~41 KB/S** and outgoing traffic **~1.2 MB/S**
- Currently running in two data centers
  - 6 publishing servers communicating to 6 deploying servers
  - Publishing server intercommunication is synced off six 3-node cluster of Redis Enterprise
  - 198 web application servers serving up the self-service application

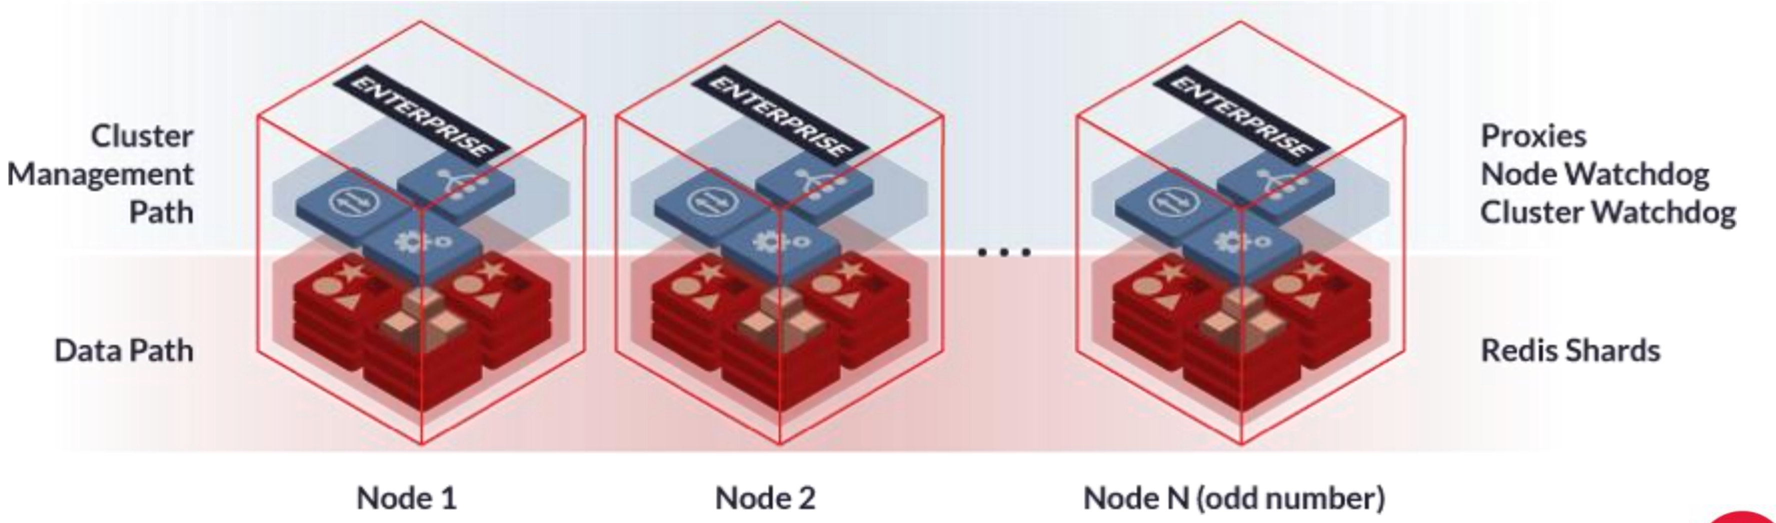

AllianceData

Why Redis Enterprise?

Top notch 24x7 expert support

Fewer resources, lower overhead

Simple, Seamless Scaling and Clustering

Stable and Predictable High Performance

Benefits Observed with Redis Enterprise

Numeric Results

- **4x** throughput compared to what we had before
- Went from an average of 70% uptime to **97%–99% uptime**

Qualitative Benefits

- **Speed to market**—two-day set-up means business operations can be put to the back of the mind and we don’t have to worry
- **Streamline client onboarding and easier ability to scale**—we can consider clients we would have passed over otherwise
- **Ease of training**—a developer can easily transition from never having heard of Redis to being fully hands-on

Next Steps

- Standardizing Redis Enterprise as the primary over coherence

- Building new application with Redis Enterprise Initial designs for chatbot solution include Redis Enterprise use case

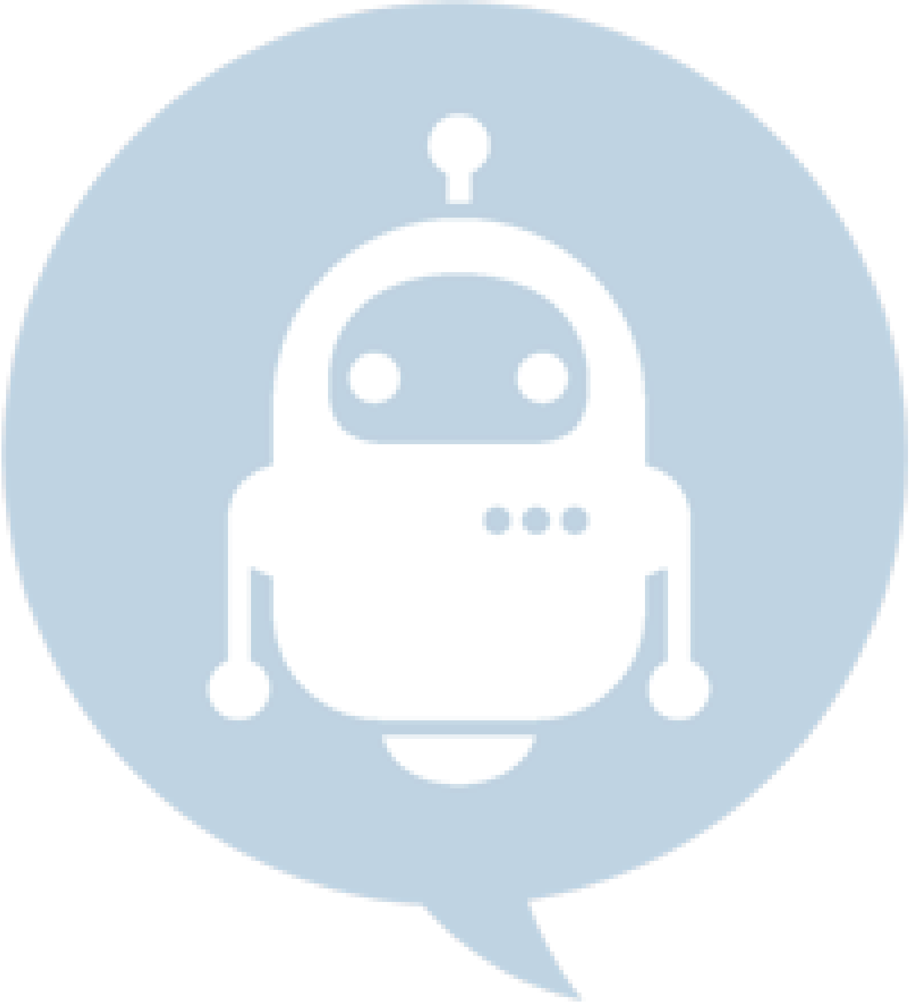

Scaling >100x with Redis Labs

Agenda

- **Introductions**
- **Gap Inc. International footprint**
- **Our Journey**
  - What we set out to achieve / business requirements?
  - What were the challenges?
  - How did we approach the solution?
  - And the results…
  - Why Redis Labs?
  - What did we learn along the way?
  - It’s not over yet…

Who are we?

Junaid Fakhruddin

Director IT - Product Architecture @ Gap Inc.

https://www.linkedin.com/in/junaid-fakhruddin-aa39332

Bhilhanan A Jeyaram

Software Engineer II @ Gap Inc.

linkedin.com/in/bhilhananajeyaram

Gap Inc. International footprint

- *Over 3,100 stores and 9 distribution centers globally*
- *Online e-commerce features various fulfillment programs like*
  - *Ship from Store*
  - *Buy Online, Pick Up in Store*
  - *Order in Store*
- *Current data volume representing reads/writes, IOPS*
  - *Total # of Node/SKU combinations ~ 50M–80M at any time.*
  - *Number of Write Requests to Node/SKU combinations - 5M updates/hour; 47M updates/day*
  - *Number of Read Requests to Node/SKU combinations - 2M reads/hour; 15M reads/day*
  - *Latency for reads ~ 40 ms*

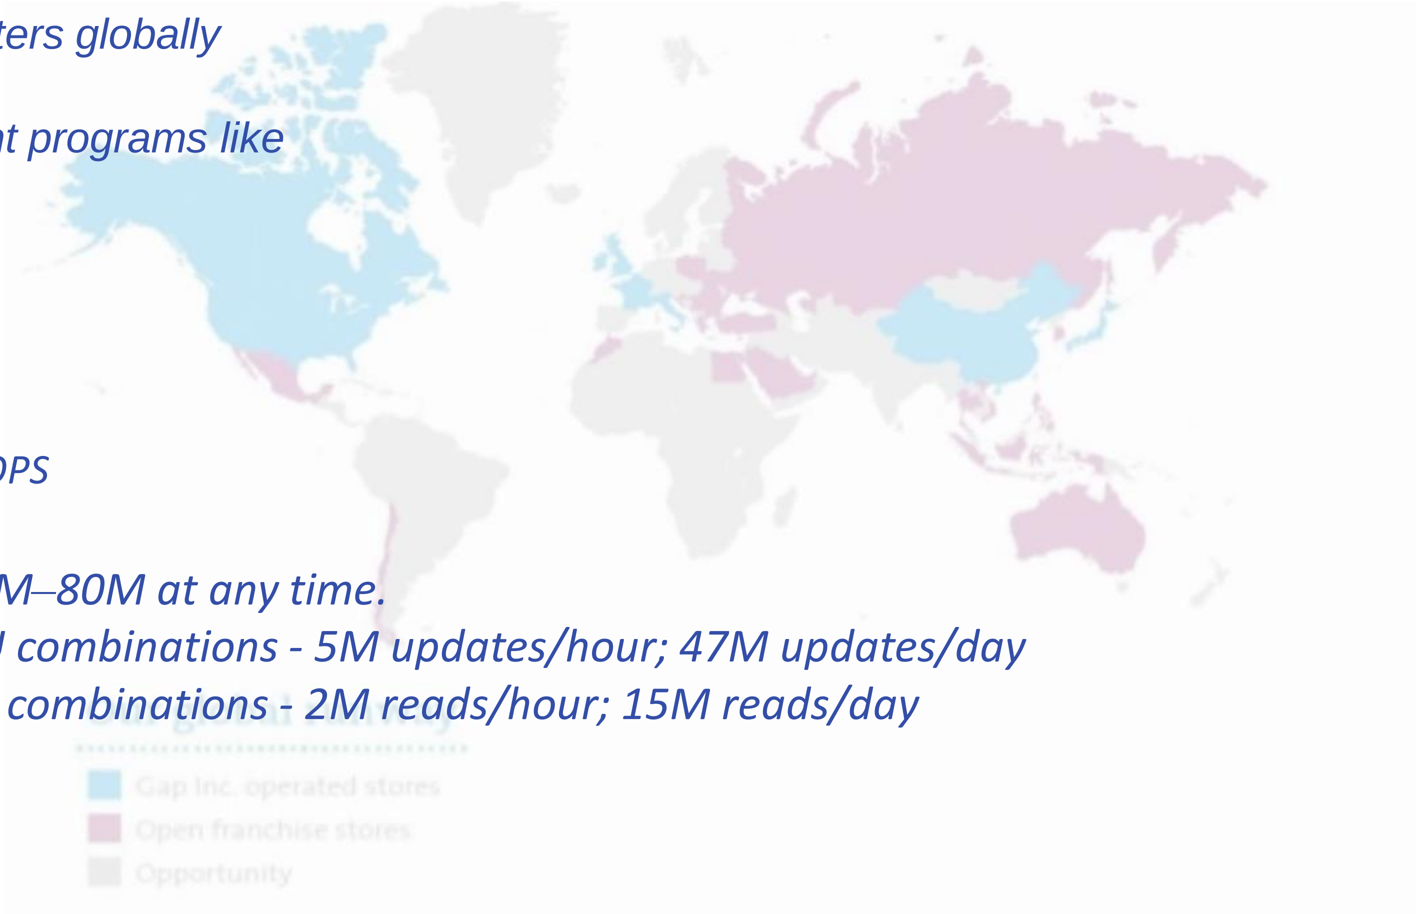

What did we set out to achieve as a business requirement?

“To optimize shipping operations and cost while providing customers with fulfillment transparency and choice in real time.”

What does it mean for the stakeholders?

Knowing how to connect the consumer with the product as early as the search page, given all the factors involved including, but not limited to:

1. Inventory availability visibility in total and at each location
2. Inventory protection and capacity based on historical and current operating run-rates
3. Location to customer fulfillment capabilities & options at a global scale
4. Distance to the customer from all possible ship-from locations
5. Pickup and delivery dates that are accurate, reliable and based on historical and current performance

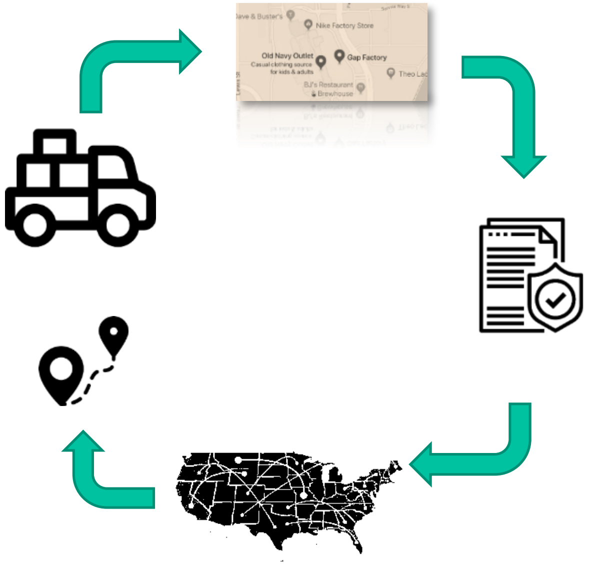

What were the challenges?

Performance

Large codebase

Non-standard integrations

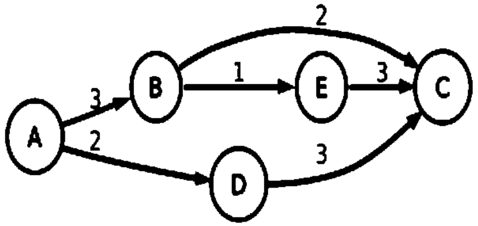

Critical path system

How did we approach the solution?

Our solution consisted of 2 major elements using

1. **Event sourcing pattern -** *to move the data into a fast layer providing APIs for real time inventory access.*

2. **Read optimized data layer -** *for fast searching and aggregations over millions of inventory events.*

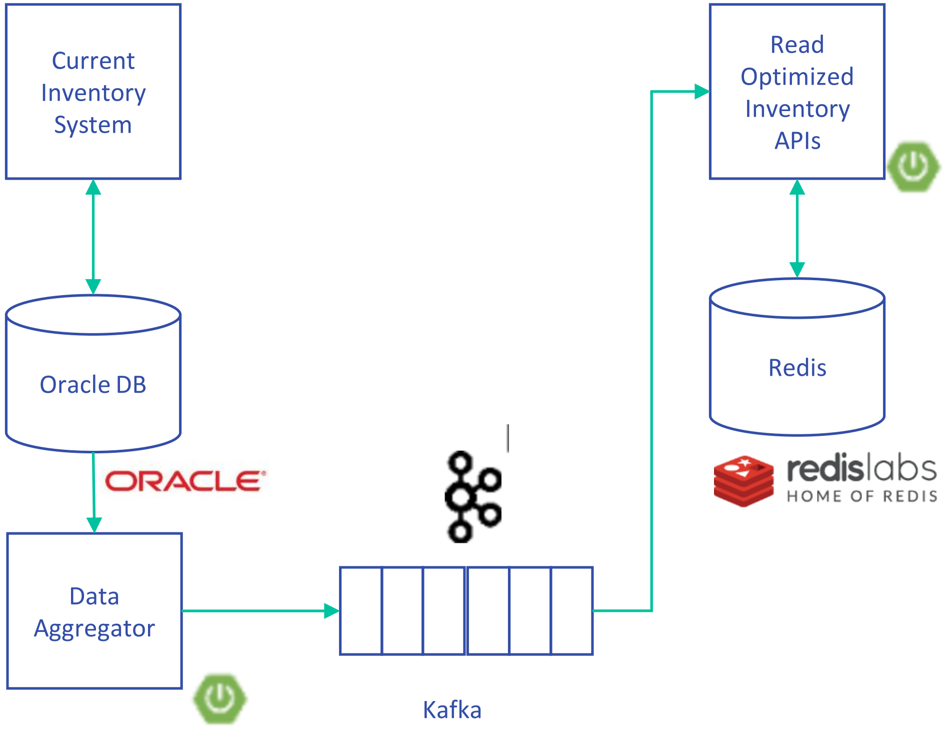

And the results…

~3 ms response time for all relevant queries

=

>100x improvement

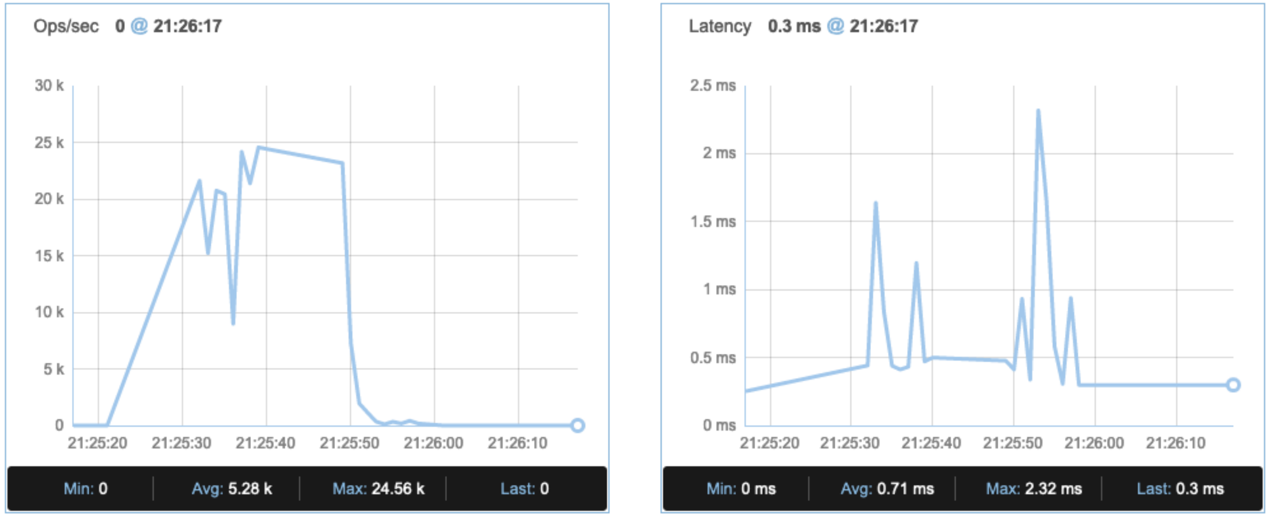

Figure 1 – Redis Labs monitoring representing ~3ms read latency under sustained load

10K+ writes/second (**~0.5 ms** latency)

=

40x improvement

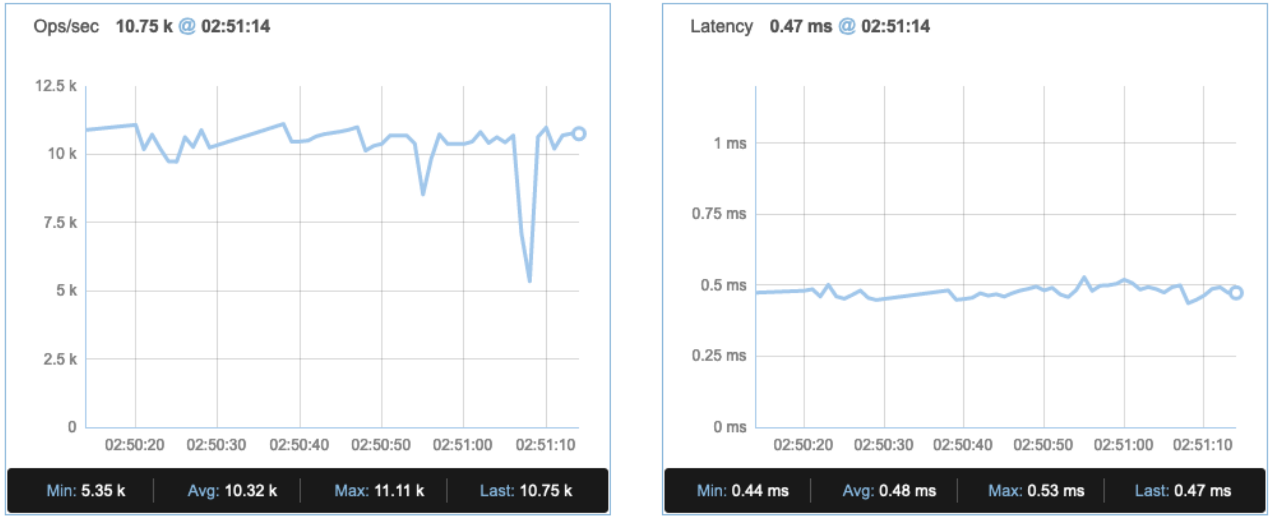

Figure 2 – Redis Labs monitoring representing 15K writes/per second under sustained load

Why Redis Labs?

After looking at various data stores that provide a competitive solve, we landed on using Redis Labs as enterprise offering for Redis per our requirements for its 8 major characteristics over OSS Redis

	Redis	RedisLabs
HA Architecture	✓	✓
Pluggable eco-system	✓	✓
RediSearch	✓	✓
Performance	✓	✓
Business Support		✓
Manage Multiple Databases easily		✓
Automatic Failover		✓
Monitoring & alerting		✓

What did we learn along the way?

Communication is key!

Making a change that is durable

Training, training, and more training

Share feedback with your vendor

It's not over yet...

We are a 24/7 operation with the business requirement of the most up to date inventory picture for our end consumers who keep our ship afloat.

What’s next

1. Disaster recovery failover configuration and testing

What’s desirable

1. Active-Active configuration of Redis supported with RediSearch capability

redislabs
HOME OF REDIS

Redis for Modern Data Models

Popular Data Models

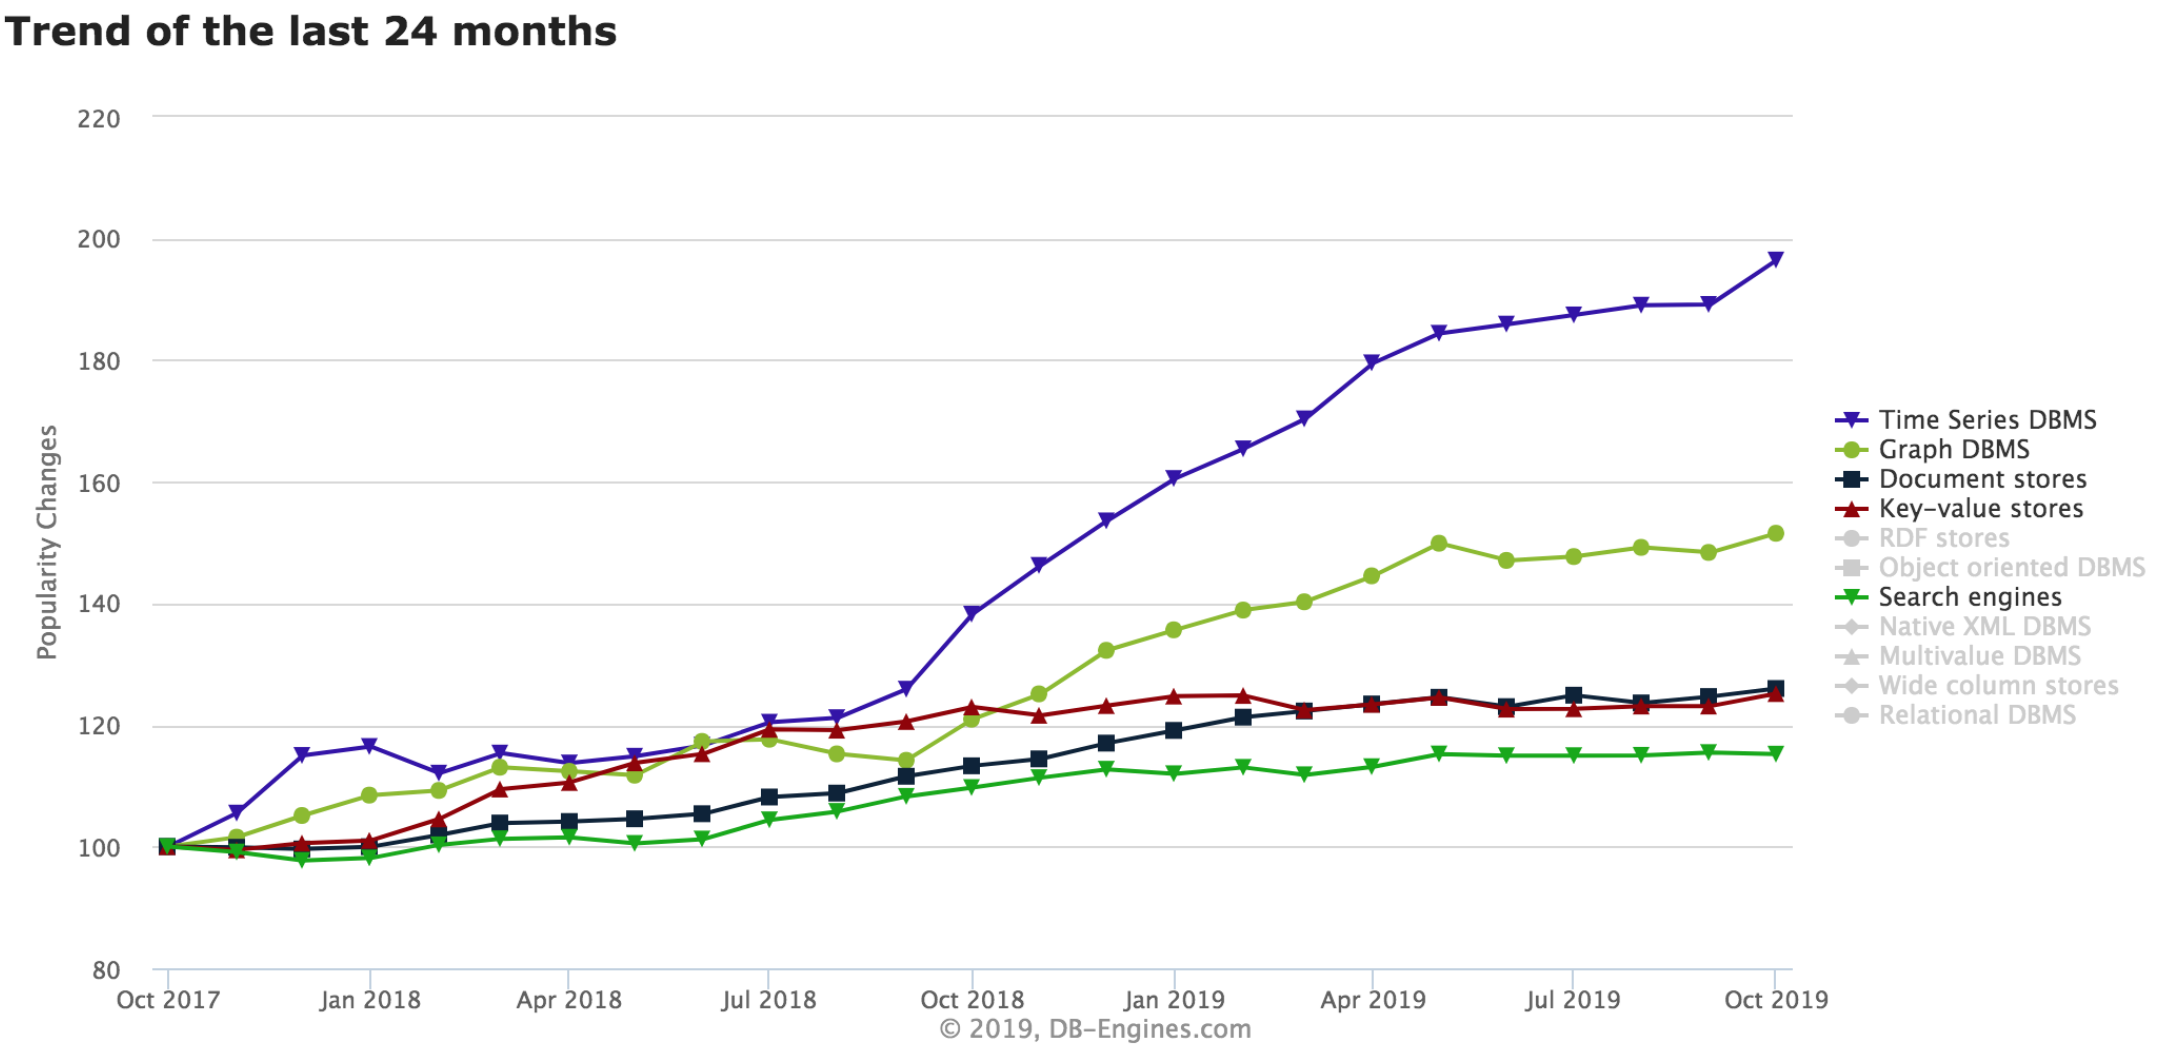

Modern Data Models in Redis

Streams

RediSearch

RedisGraph

RedisTimeSeries

RedisBloom

RedisJSON

RedisAI

Modern Data Models in Redis

Streams
(core)

Main capabilities

- Consumers groups
- Capped streams
- Removing item

Use cases

- Message broker
- Event sourcing, CDC

Modern Data Models in Redis

RediSearch
(module)

Main capabilities

- Full-text search
- Indexing, auto complete
- Aggregation, ranking, stemming
- Extendable and modular

Use cases

- Fast search
- Multi-tenant indexing

redisearch.io

Modern Data Models in Redis

RedisBloom
(module)

Main capabilities

- Bloom Filter
- Cuckoo Filter
- Count-min sketch
- Top-K

Use cases

- Authentication
- Ad serving, leaderboards

redisbloom.io

Modern Data Models Are **Faster** in **Redis**

x500

Streams

x4 – x10

RediSearch

x6 – x600

RedisGraph

X4 – x100*

RedisTimeSeries

x17

RedisBloom

x4*

RedisJSON

Soon

RedisAI

Redis Enterprise

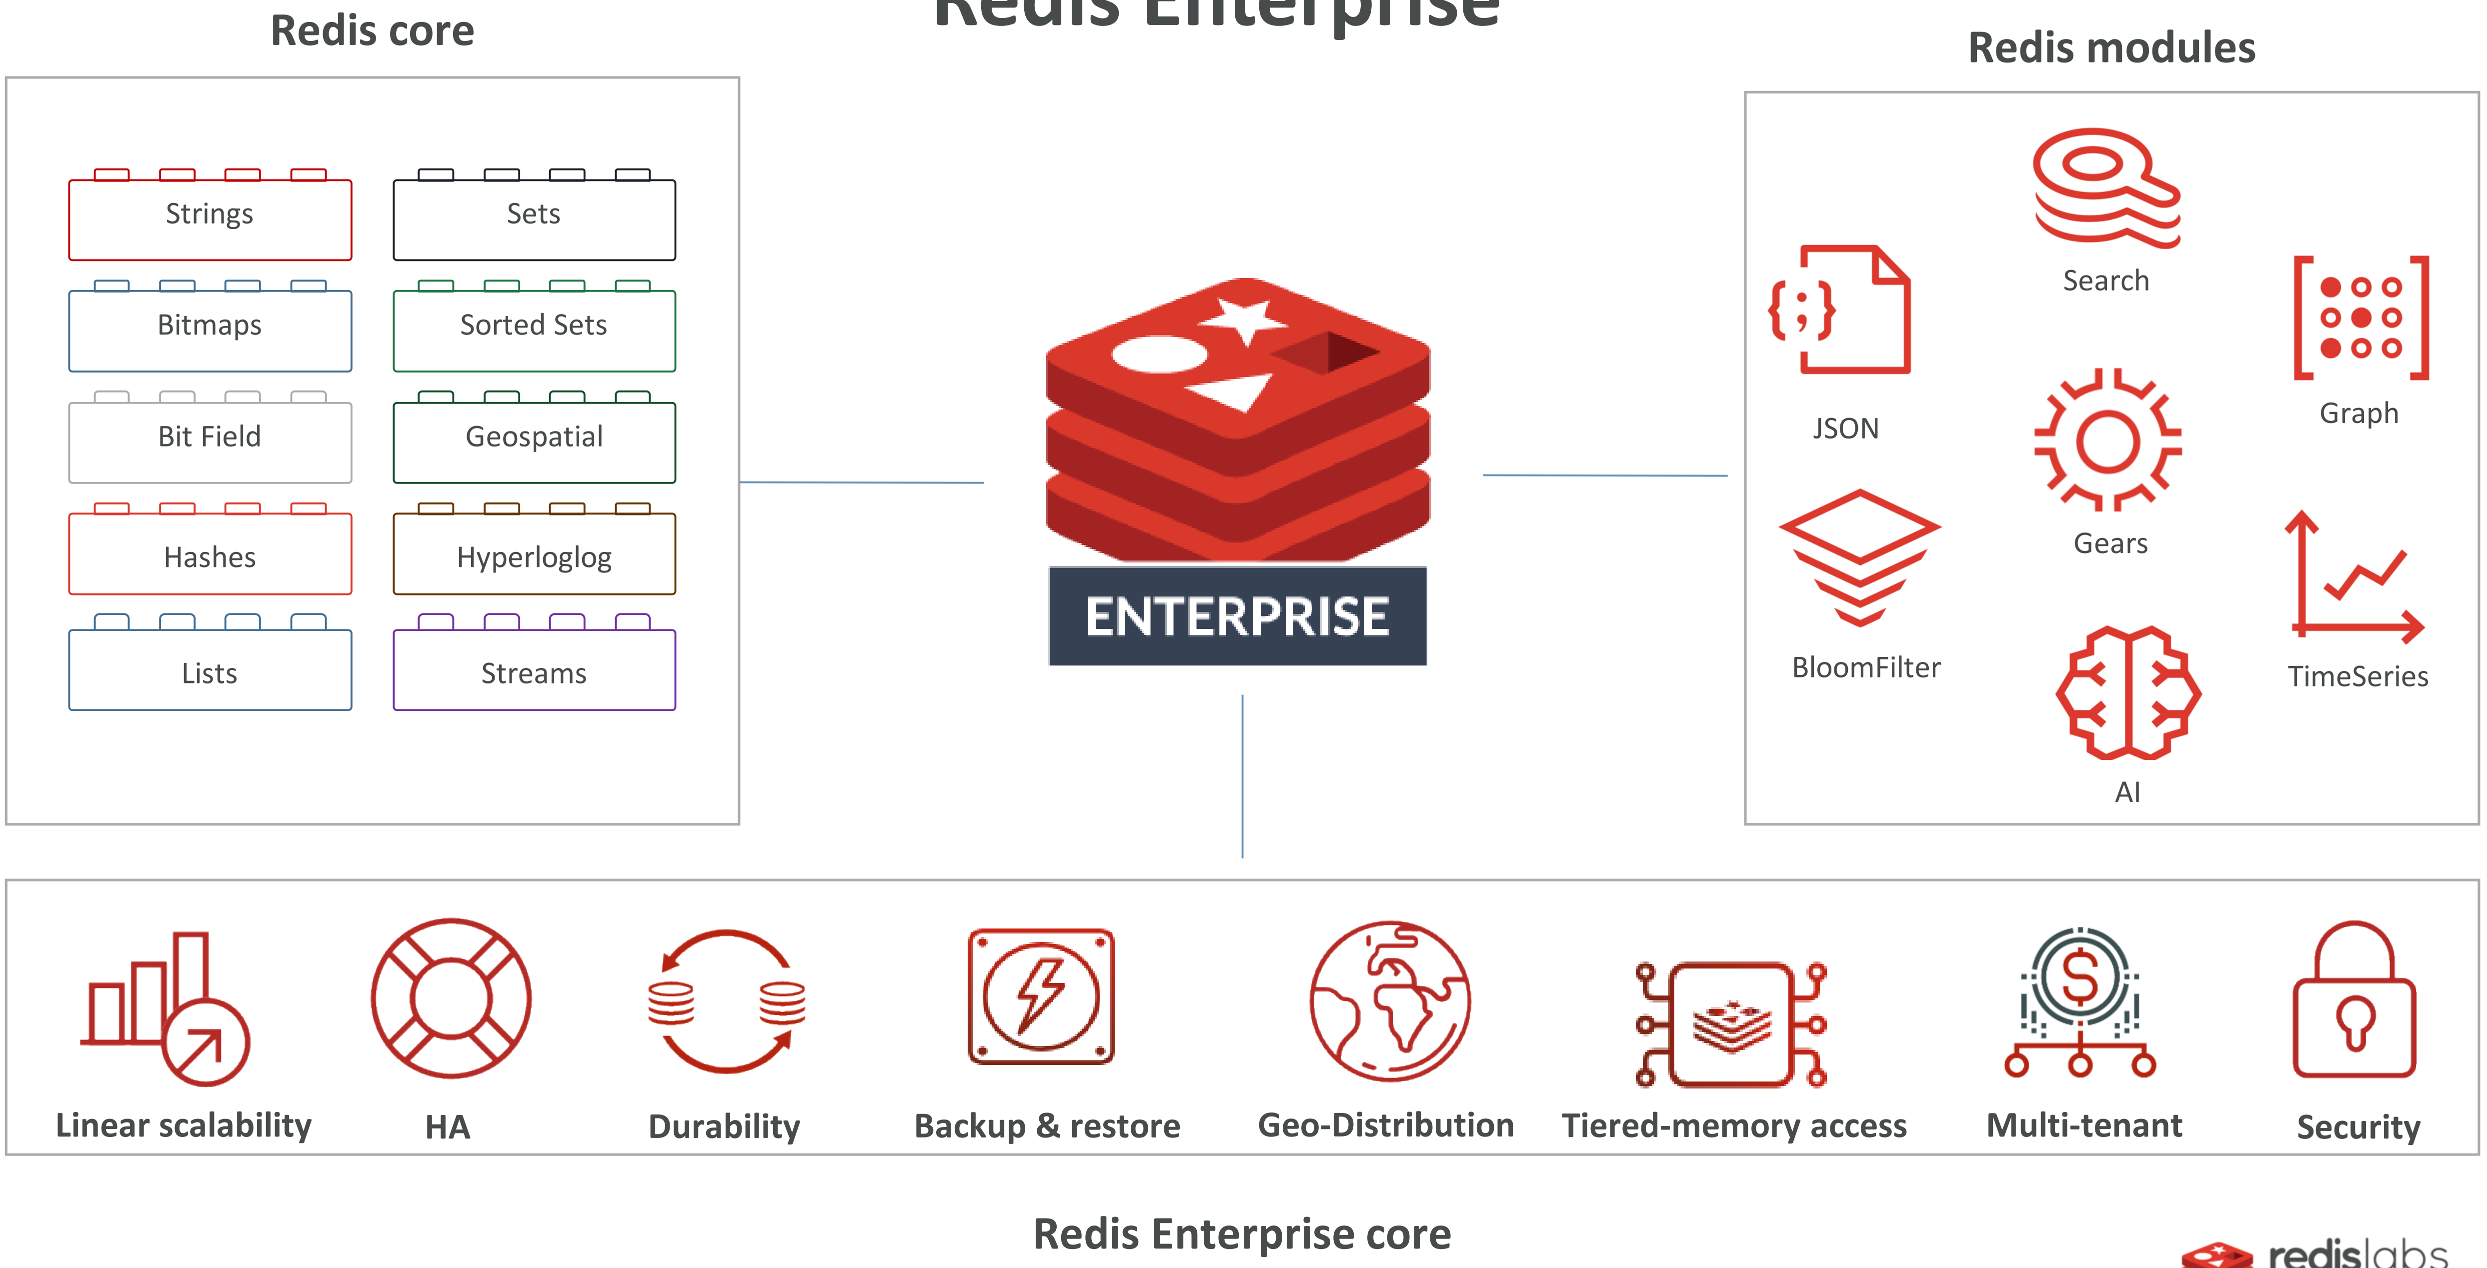

Redis core

- Strings
- Sets
- Bitmaps
- Sorted Sets
- Bit Field
- Geospatial
- Hashes
- Hyperloglog
- Lists
- Streams

ENTERPRISE

Redis modules

- Search
- JSON
- Graph
- Gears
- BloomFilter
- TimeSeries
- AI

Redis Enterprise core

- Linear scalability
- HA
- Durability
- Backup & restore
- Geo-Distribution
- Tiered-memory access
- Multi-tenant
- Security

redislabs
HOME OF REDIS

Redis Enterprise

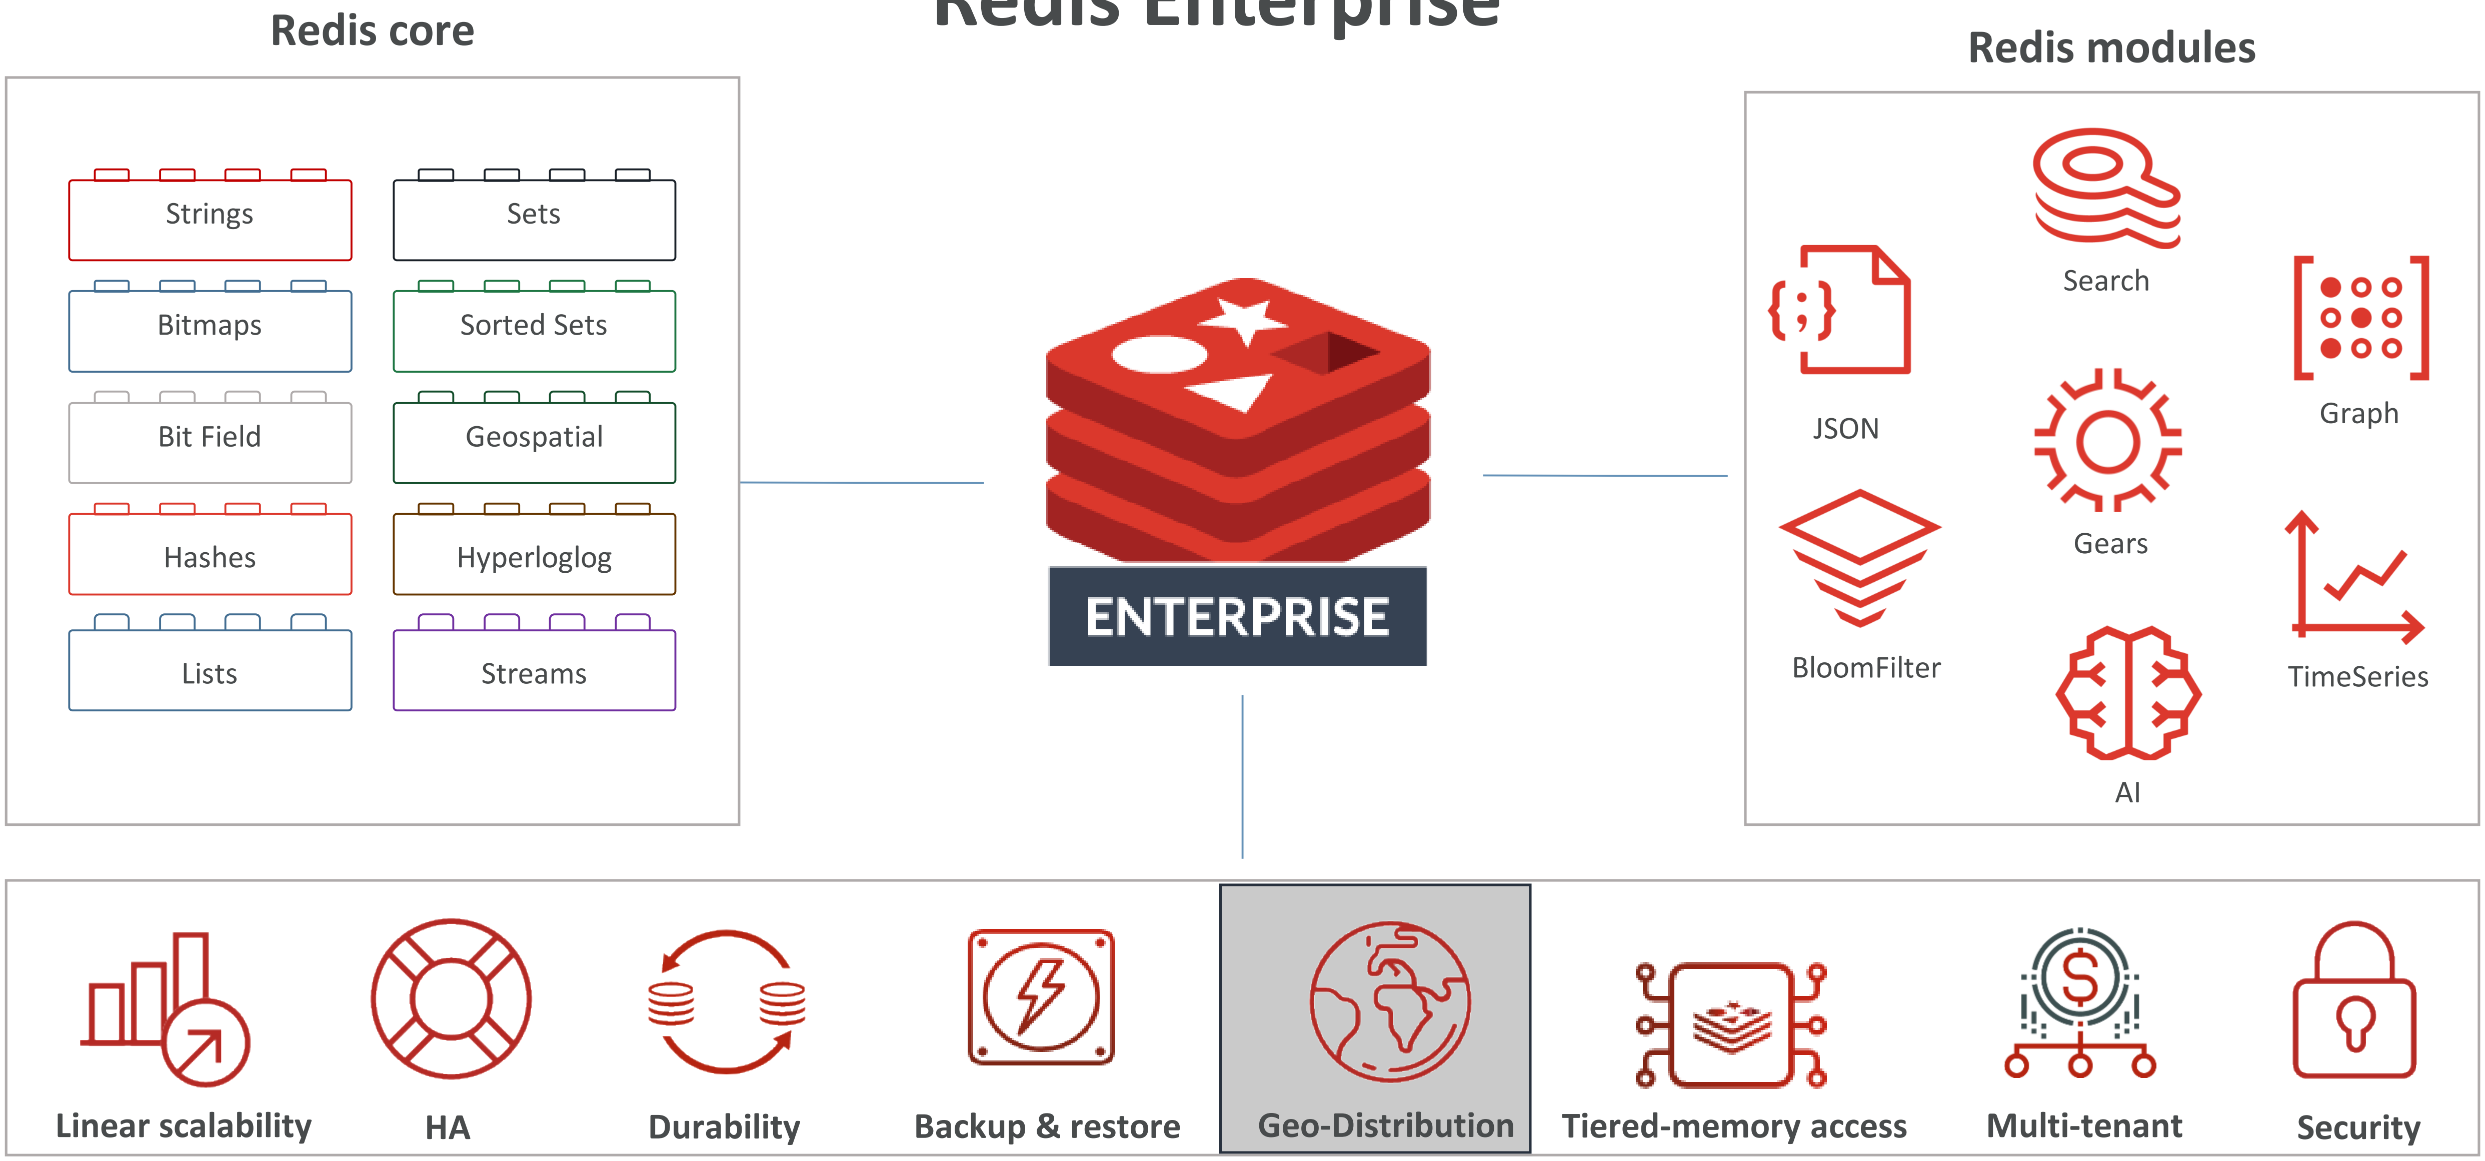

Redis core

Strings

Sets

Bitmaps

Sorted Sets

Bit Field

Geospatial

Hashes

Hyperloglog

Lists

Streams

ENTERPRISE

Redis modules

Search

JSON

Graph

Gears

BloomFilter

TimeSeries

AI

Linear scalability

HA

Durability

Backup & restore

Geo-Distribution

Tiered-memory access

Multi-tenant

Security

Redis Enterprise core

Why Do You Need Active-Active Redis ?

Disaster recovery

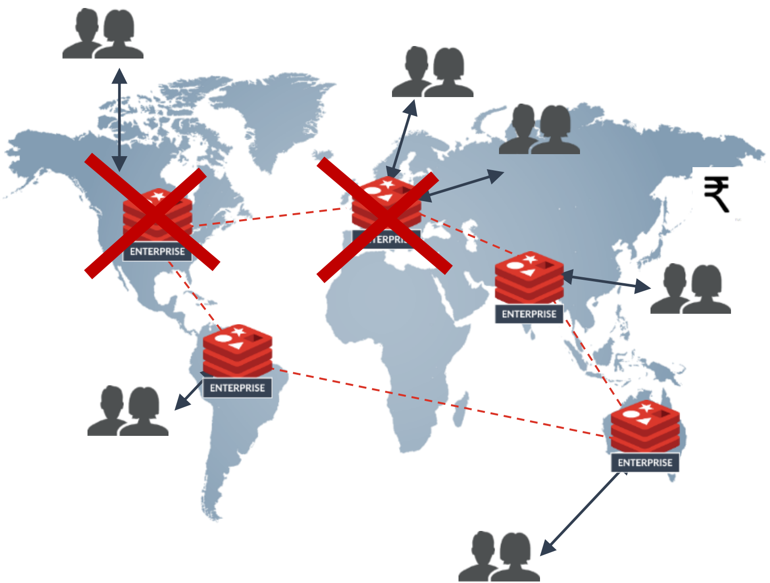

Why Do You Need Active-Active Redis ?

Disaster recovery

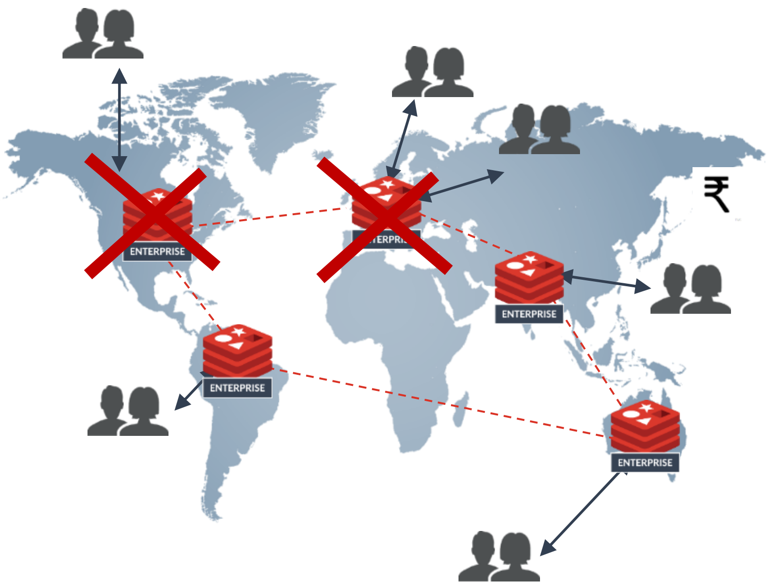

Session management

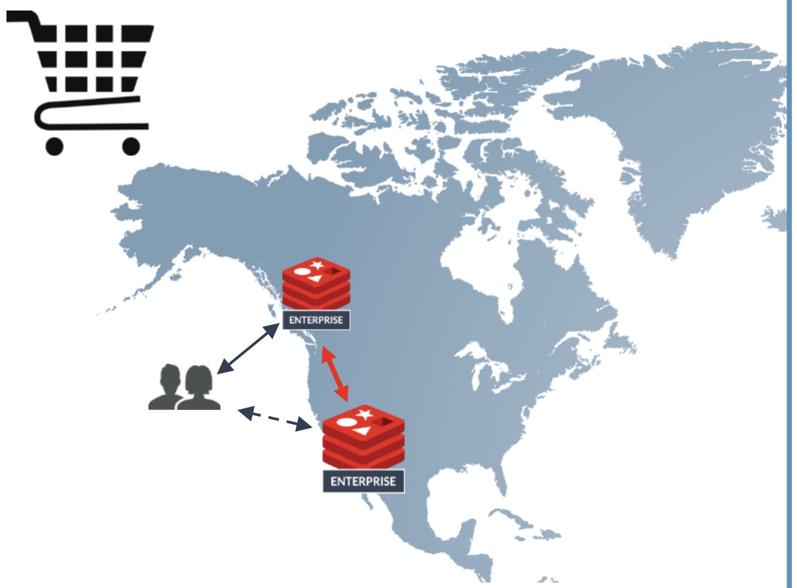

Why Do You Need Active-Active Redis ?

Disaster recovery

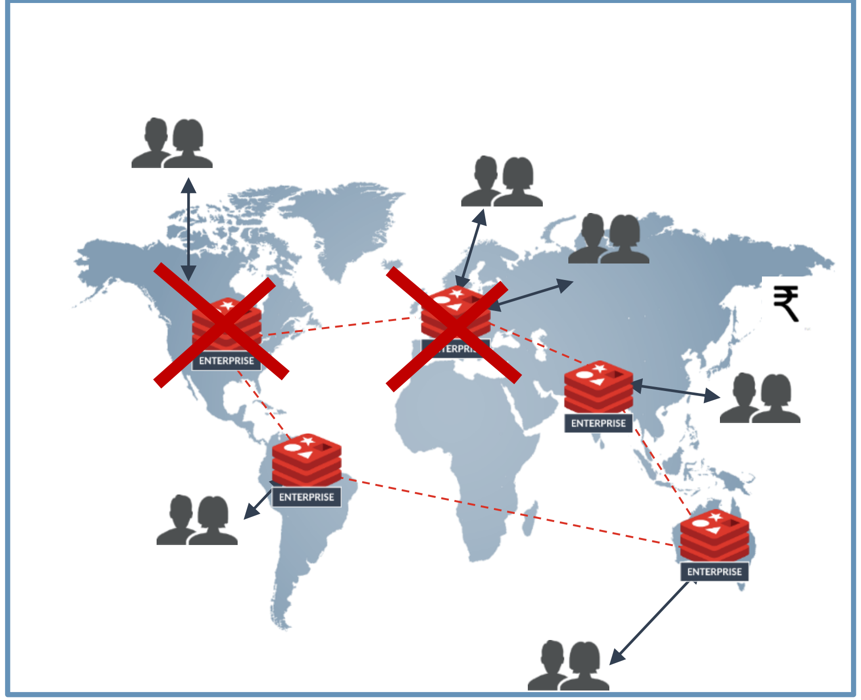

Session management

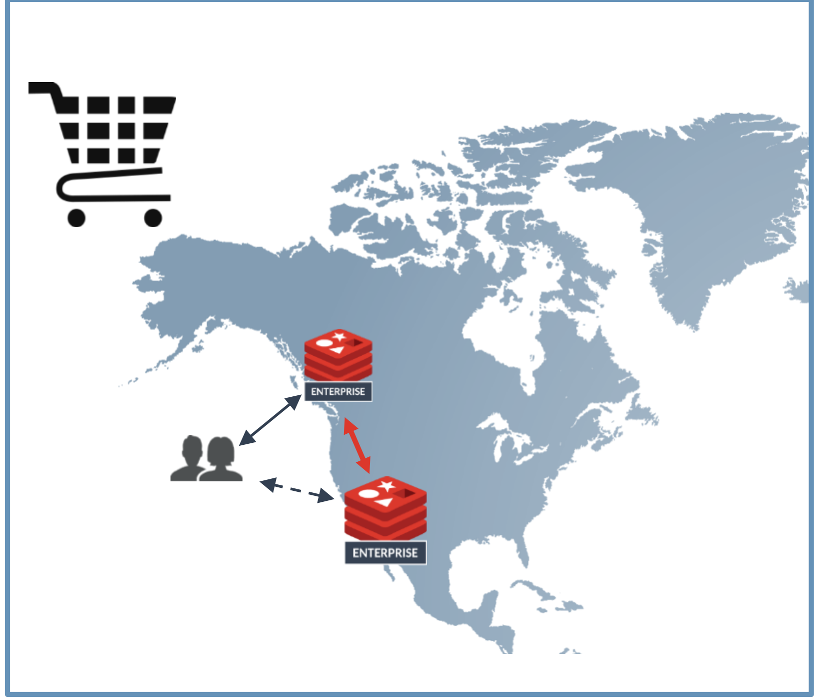

Globally distributed applications

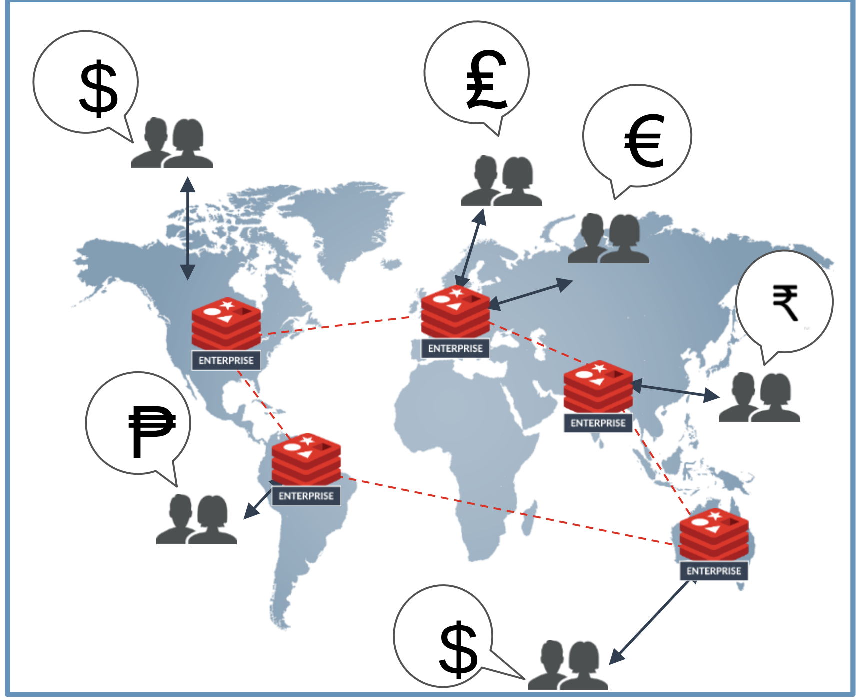

Active-Active Comparison

Cassandra/Amazon DynamoDB

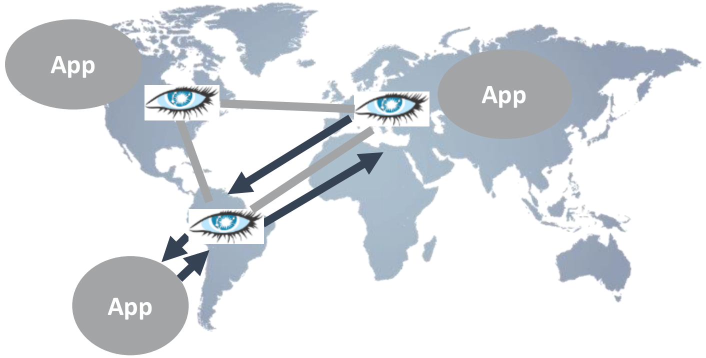

Eventual Consistency

100 ms

Spanner

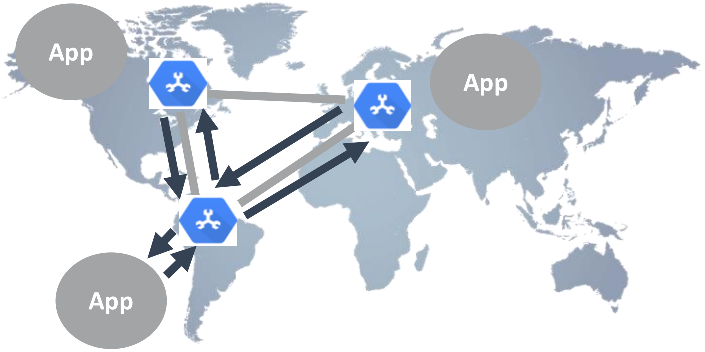

Strong Consistency

200 ms

Redis CRDTs

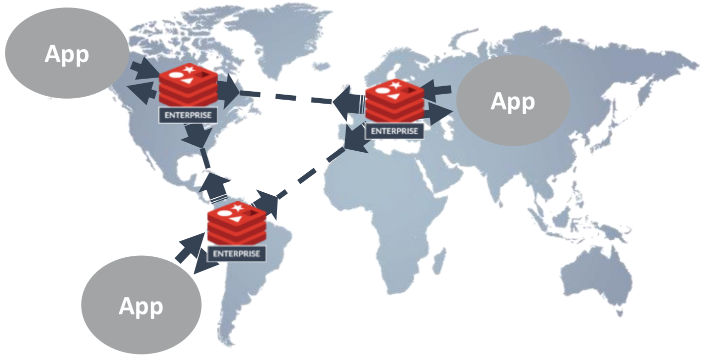

Strong Eventual Consistency
+
Causal Consistency

<1 ms

Redis Enterprise Cloud SLA

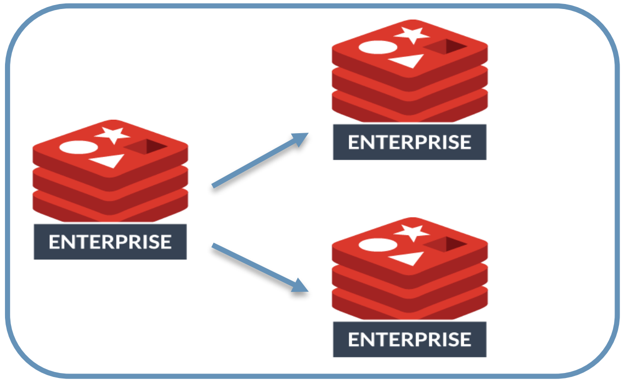

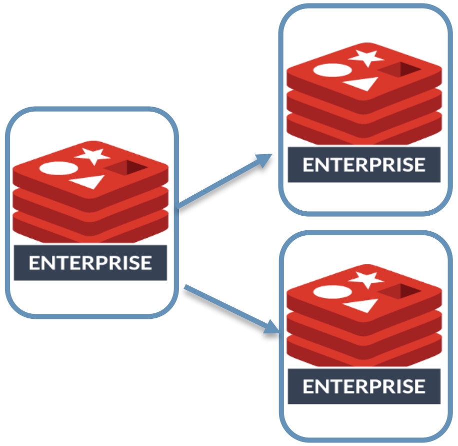

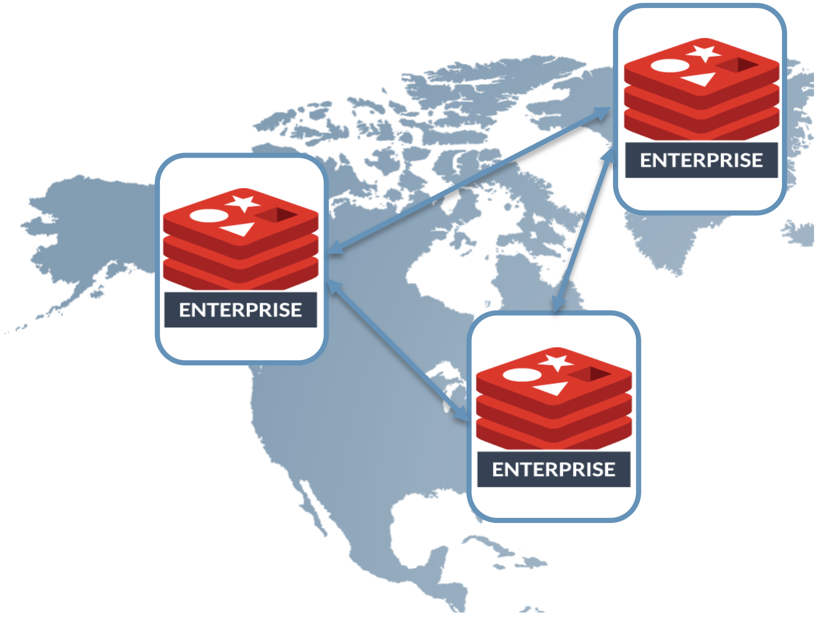

99.9% (<10m4.8s)
Active-passive replication
Single-AZ

99.99% (<4m23s)
Active-passive replication
Multi-AZ

99.999% (<26.3s)
Active-active replication
Multi-region

Redis Enterprise

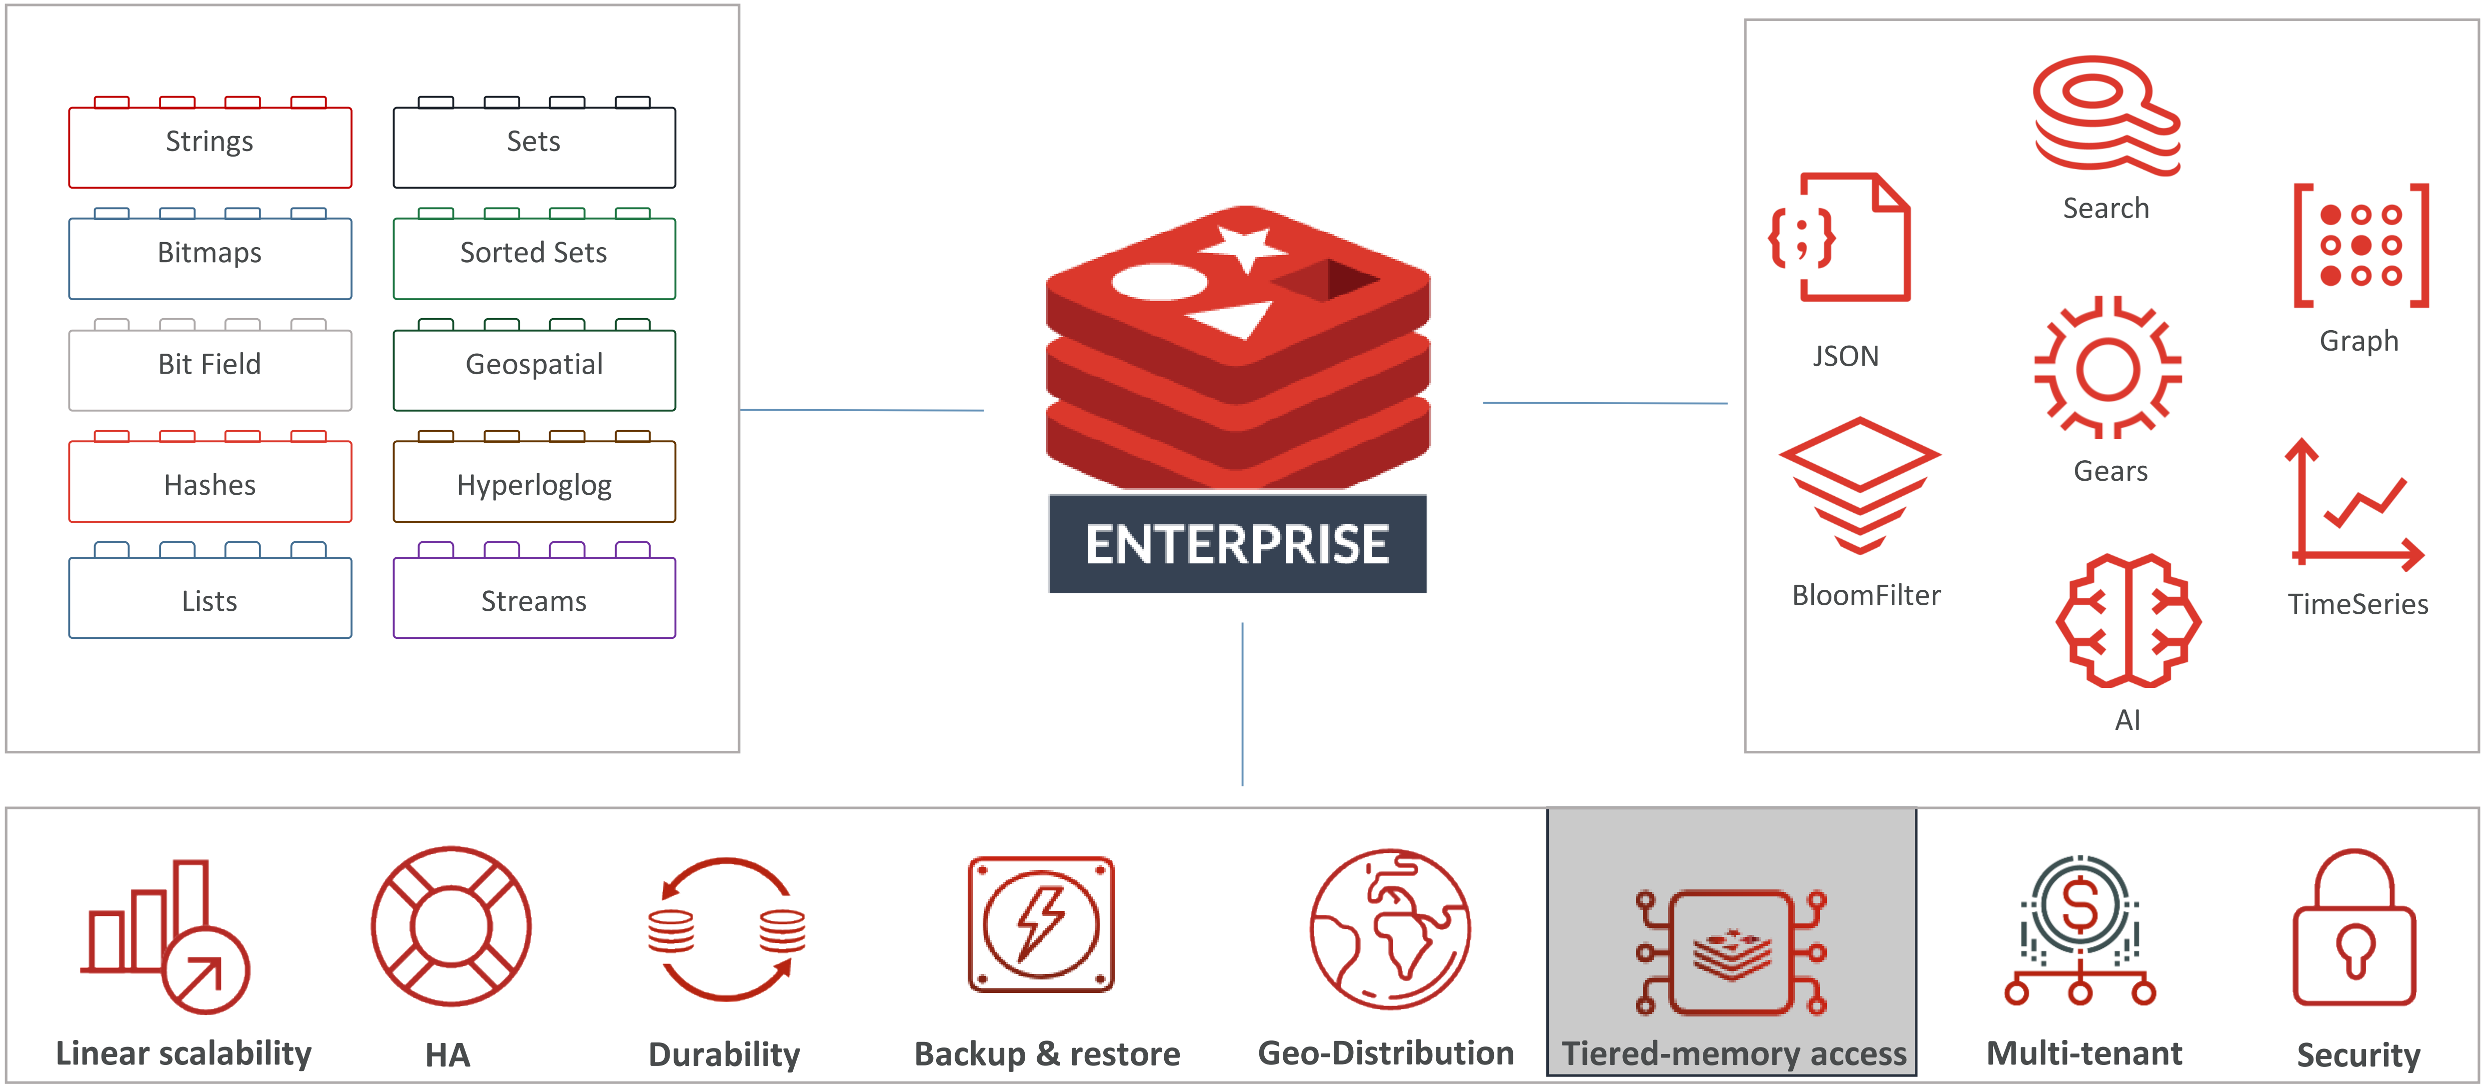

Redis core

Strings

Sets

Bitmaps

Sorted Sets

Bit Field

Geospatial

Hashes

Hyperloglog

Lists

Streams

ENTERPRISE

Redis modules

Search

JSON

Graph

Gears

BloomFilter

TimeSeries

AI

Linear scalability

HA

Durability

Backup & restore

Geo-Distribution

Tiered-memory access

Multi-tenant

Security

Redis Enterprise core

redislabs
HOME OF REDIS

GBs → TBs → PB

in

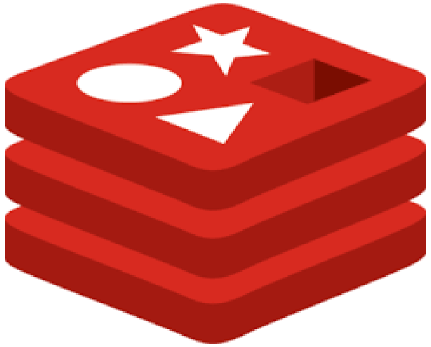

Running Redis on Multiple Memory Technologies

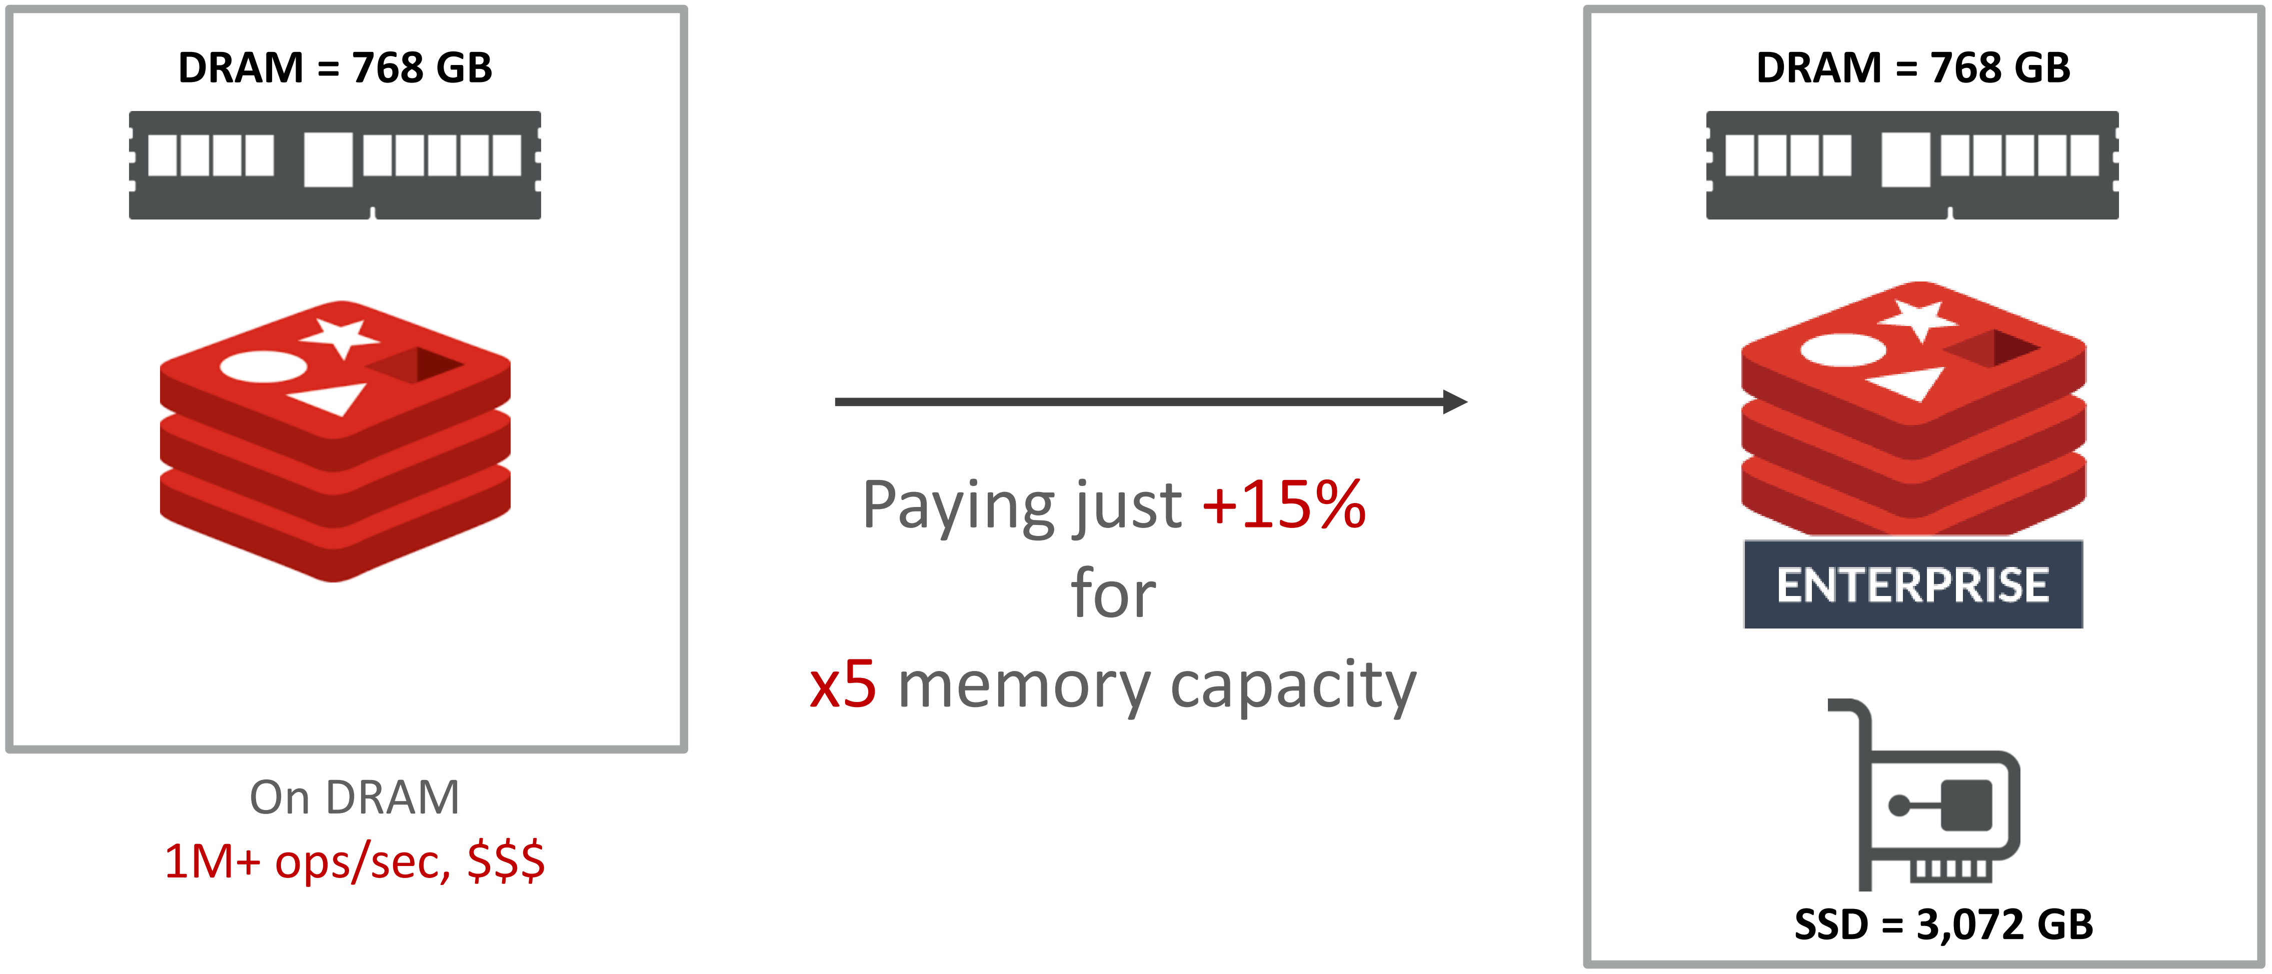

DRAM = 768 GB

On DRAM

1M+ ops/sec, $$$

Paying just +15%
for
x5 memory capacity

DRAM = 768 GB

ENTERPRISE

SSD = 3,072 GB

On Flash/SSD

100K+ ops/sec, $

1 TB Replicated Dataset (2 TB in total) @ 100K ops/sec

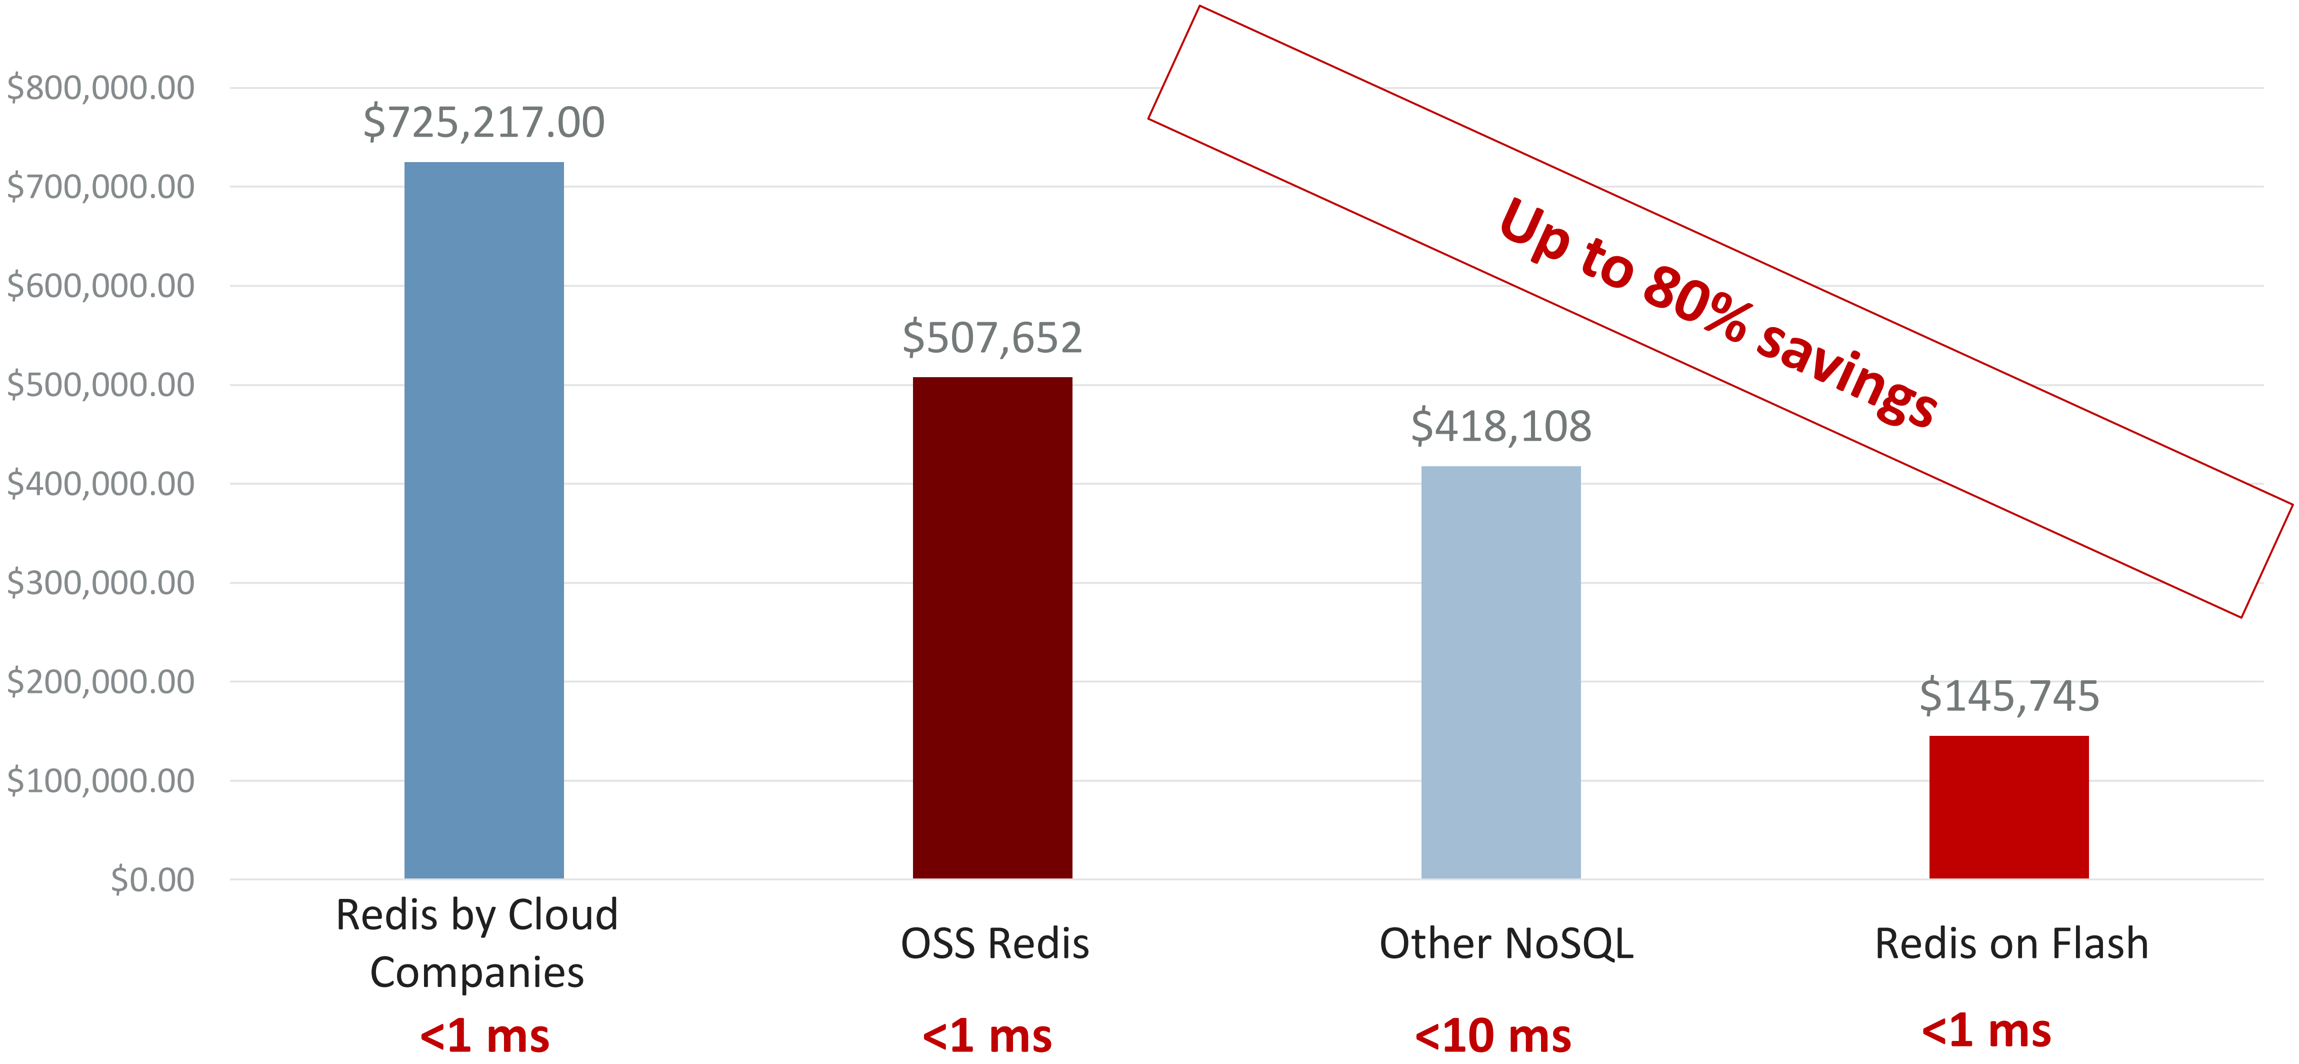

Up to 80% savings

presented by
redislabs
HOME OF REDIS

Redis Enterprise on Multiple Memory Technologies

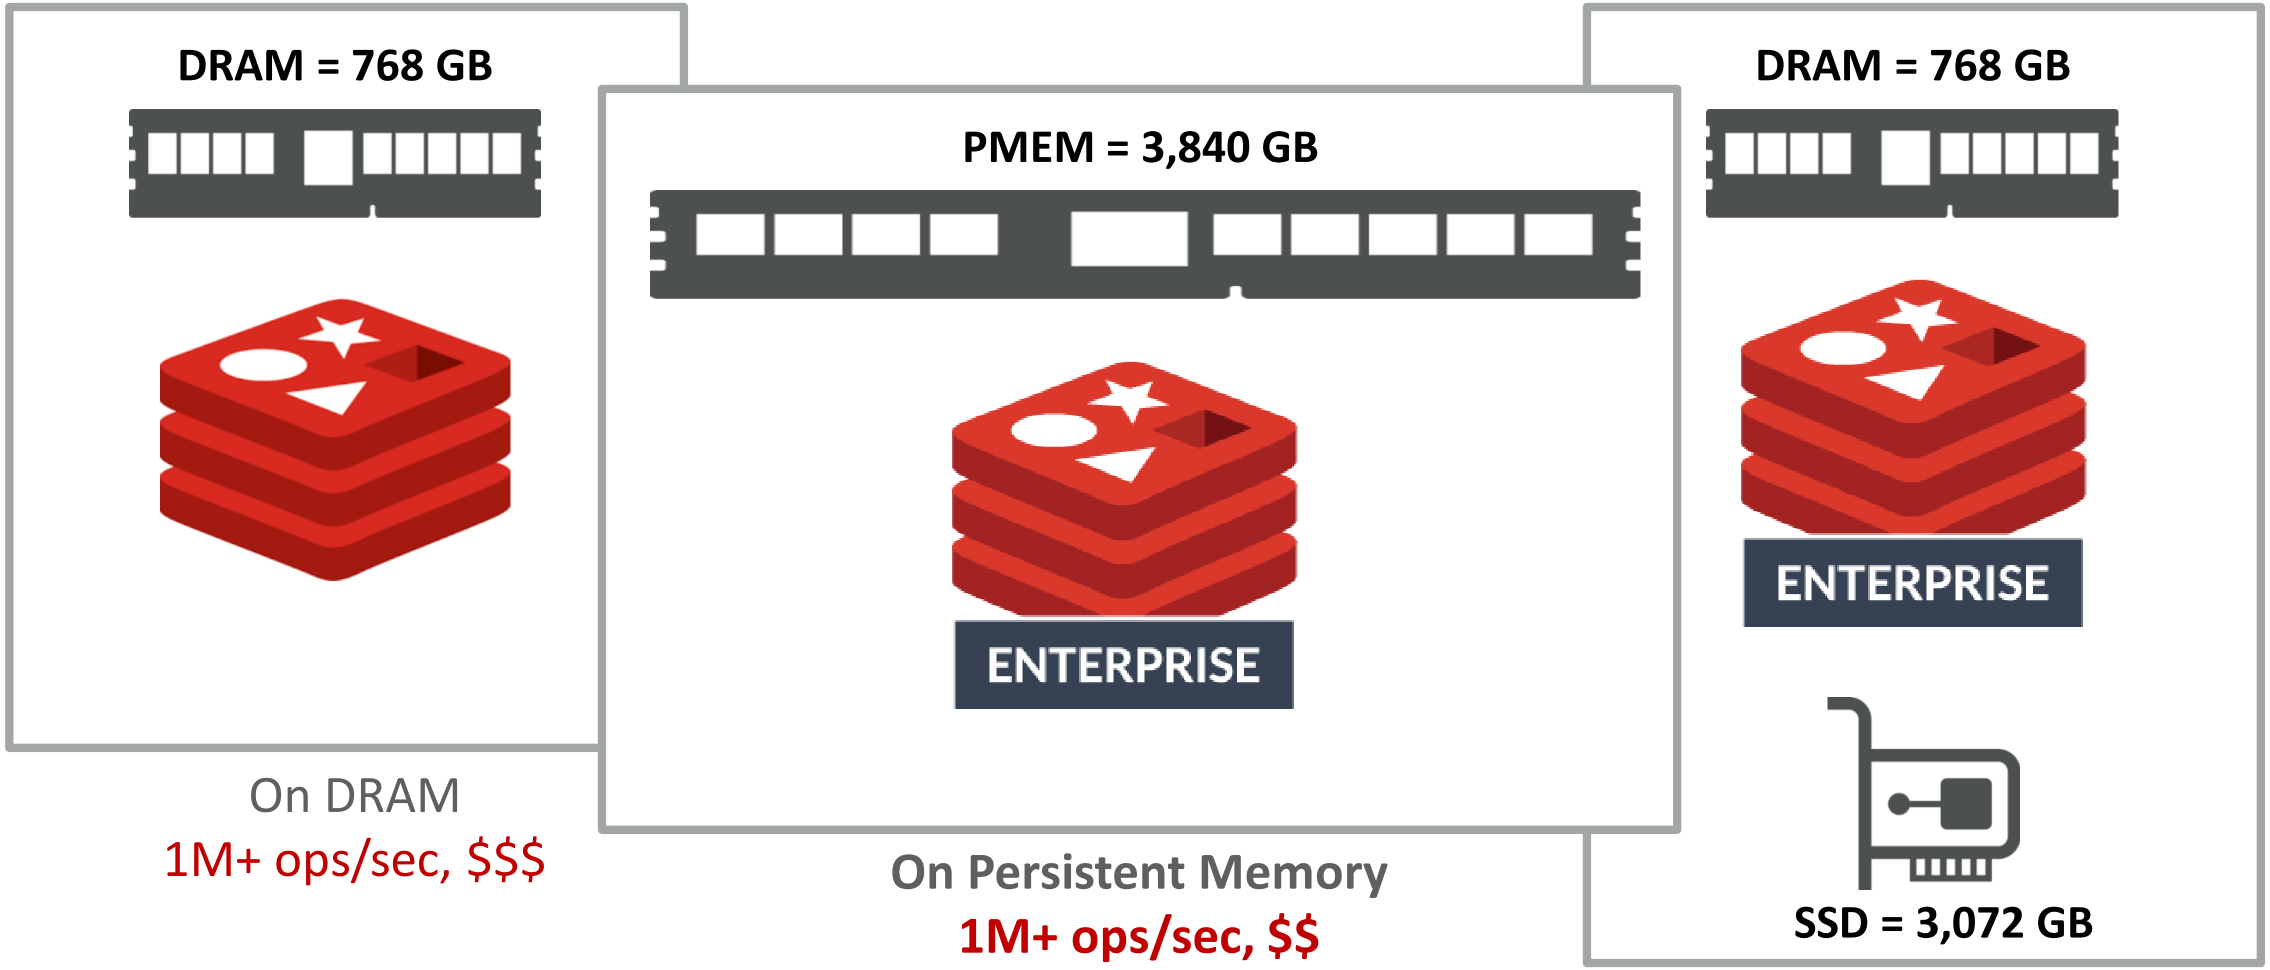

DRAM = 768 GB

On DRAM

1M+ ops/sec, $$$

PMEM = 3,840 GB

On Persistent Memory

1M+ ops/sec, $$

DRAM = 768 GB

SSD = 3,072 GB

On Flash/SSD

100K+ ops/sec, $

Why Redis for all your applications?

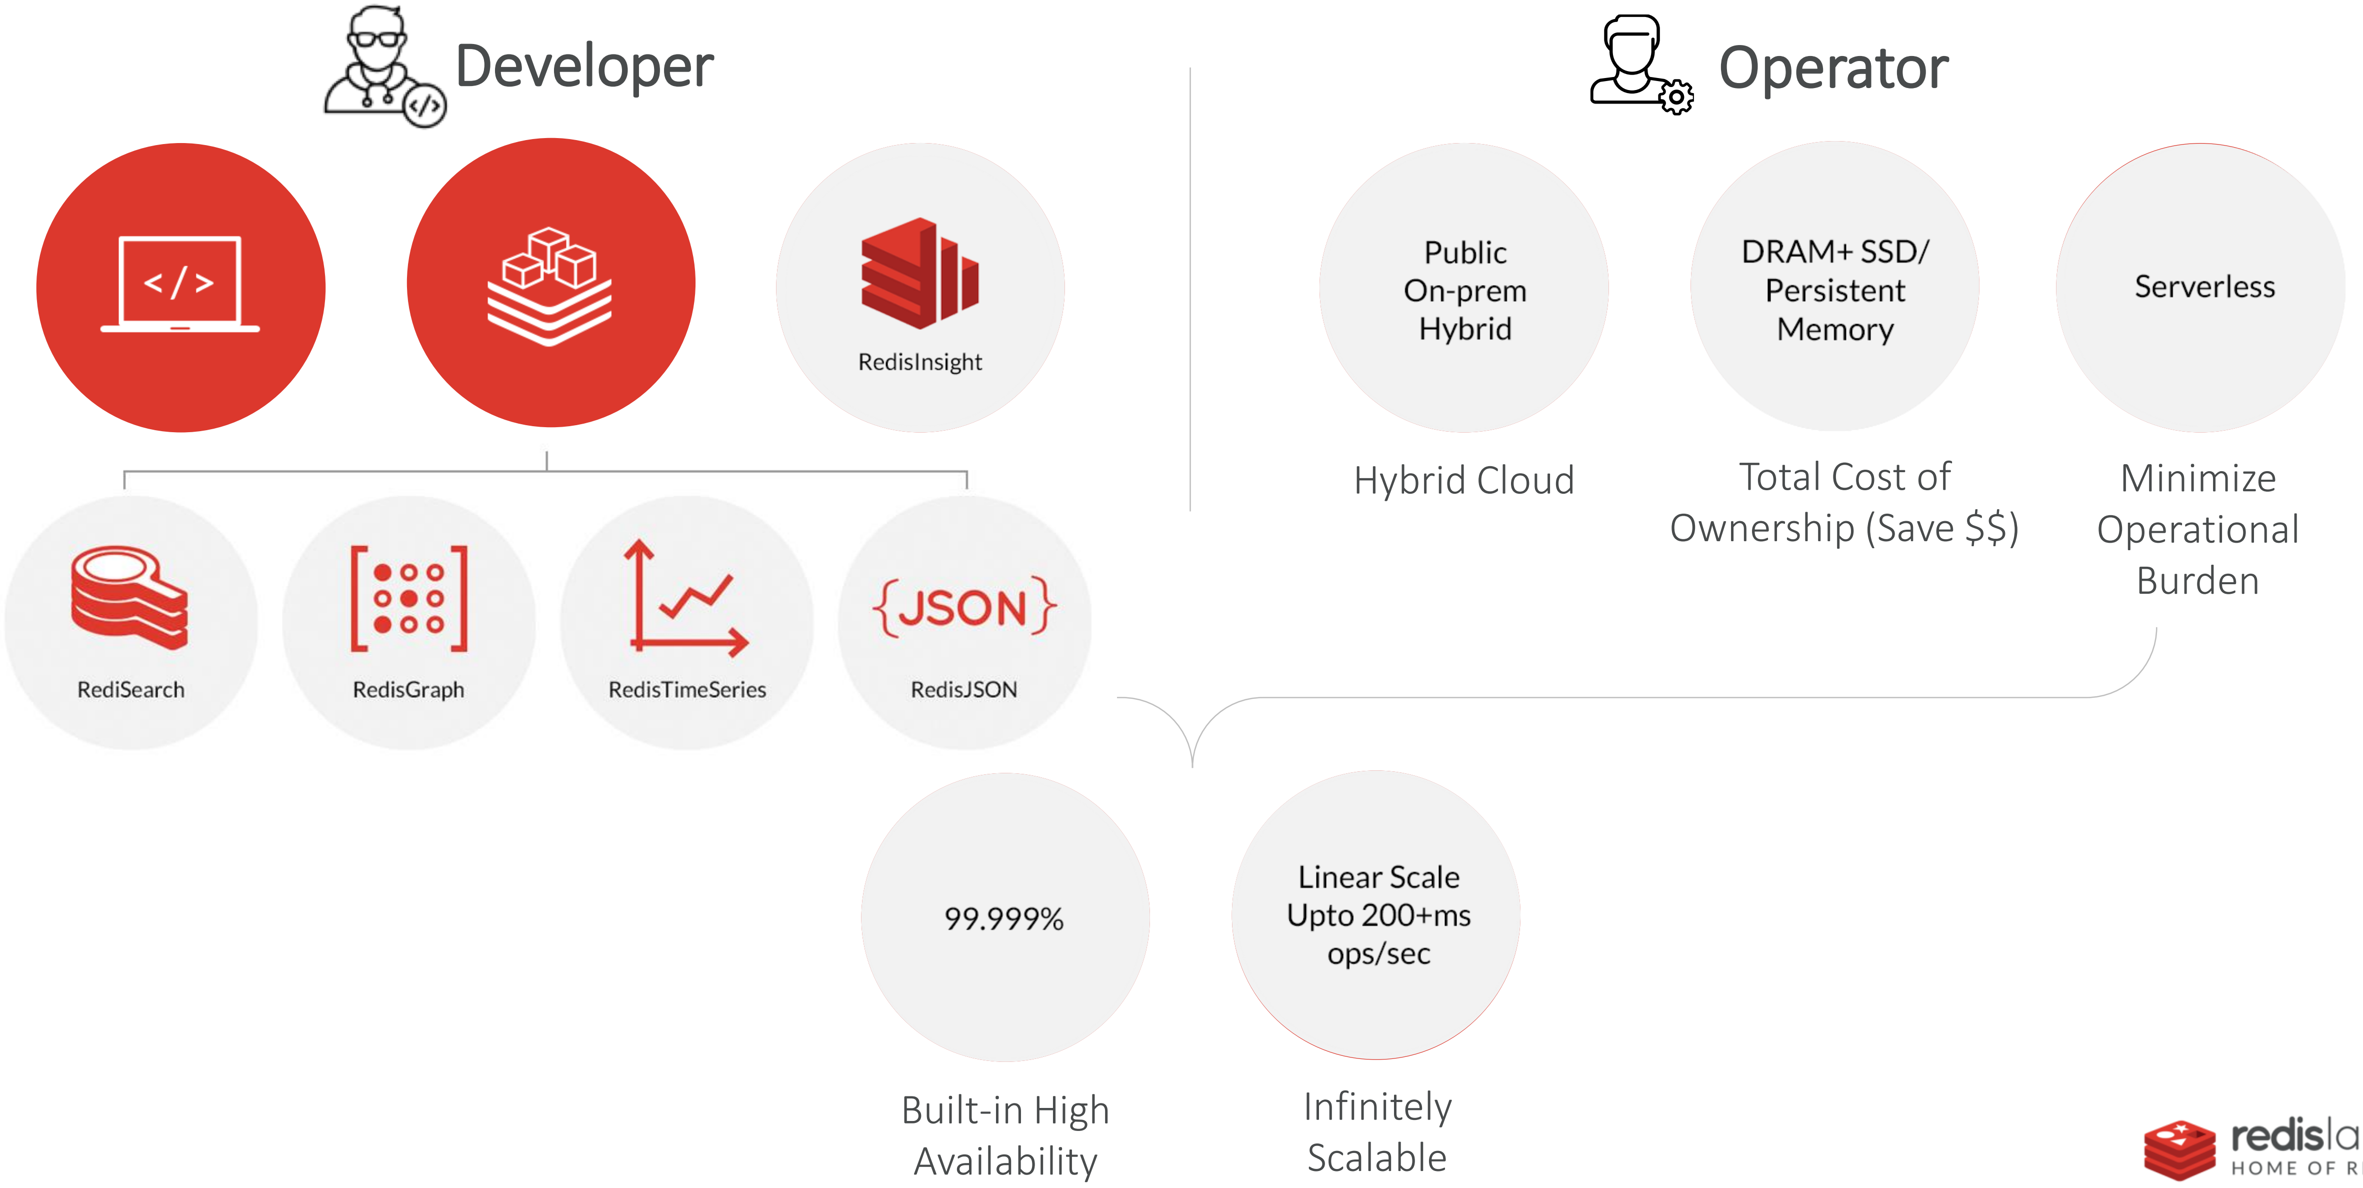

Developer

RedisInsight

RediSearch

RedisGraph

RedisTimeSeries

RedisJSON

Operator

Public
On-prem
Hybrid

Hybrid Cloud

DRAM+ SSD/
Persistent
Memory

Total Cost of
Ownership (Save $$)

Serverless

Minimize
Operational
Burden

99.999%

Built-in High
Availability

Linear Scale
Upto 200+ms
ops/sec

Infinitely
Scalable

redislabs
HOME OF REDIS

Q&A

Please complete the session survey in the mobile app.