



**innovation**  
hub sessions

Organizado por:



**redislabs**  
HOME OF REDIS

15:25 a 16:05h

## Build Highly Resilient Applications with Redis Enterprise Clustering



**Manuel Hurtado**

Solution Architect en Redis Labs.



# Build Highly Resilient Applications with Redis Enterprise Clustering


MAY 2019 | MANUEL HURTADO

# Redis Introduction

# Most Loved Databases 2017, 2018 & 2019


Stack Overflow survey, among >100K developers

*% of devs who expressed interest in continuing to develop with a database*



# Most Popular Database Container

Number of containers (in Billions) launched at Docker Hub (as of Dec 2018)




# Redis Top Differentiators

1

## Performance


*NoSQL Benchmark*



2

## Simplicity


*Redis Data Structures*



3


## Extensibility

*Redis Modules*




# 1 Performance: The Most Powerful Database

Highest Throughput at Lowest Latency  
in High Volume of Writes Scenario




Benchmarks performed by Avalon Consulting Group

Least Servers Needed to  
Deliver 1 Million Writes/Sec



Benchmarks published in the Google blog

## 2 Simplicity: Data Structures - Redis' Building Blocks



"Retrieve the e-mail address of the user with the highest bid in an auction that started on July 24th at 11:00pm PST"

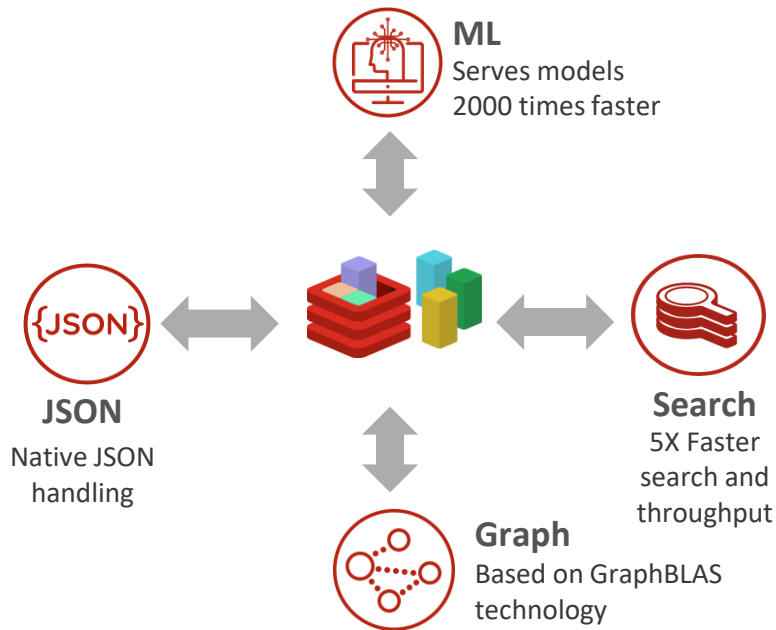
=

**ZREVRANGE 07242015\_2300 0 0**



### 3 Extensibility: True Multi-Model Functionality for All Use Cases

- Implemented by Redis Modules, independent of the Redis core
- Add-ons that use a Redis API to seamlessly support additional use cases and data structures
- Loosely coupled design, i.e. load only models needed for your use case
- Optimal data structure implementation for JSON, Graph, Search (and other) functionality, not just APIs
- Add new capabilities and data structures to Redis – in speeds similar to normal Redis commands
- Redis Enterprise Modules are tested and certified by Redis Labs



# Uniquely Suited to Modern Use Cases

A full range of capabilities that simplify and accelerate next generation applications



Real Time  
Analytics



User Session  
Store




Real Time  
Data Ingest



High Speed  
Transactions



Job & Queue  
Management



Time Series Data



Complex  
Statistical  
Analysis



Notifications



Distributed Lock



Content Caching



Geospatial Data



Streaming Data




Machine  
Learning




Search

# Highest Performance at Any Use Case


## NoSQL – x150




## Search – x5




## Graph – x400



## Streams – x500



## ML Serving – x2,000






# Redis Enterprise Architecture


# What is a Redis Enterprise cluster

## Database Types

- Simple database (single master shard)
- Highly available (HA) database
- Clustered database (no replication)
- Clustered HA database




# Shared-Nothing, Linearly Scalable, Multi-Tenant, Symmetric Architecture



Separation between the data path and the control path components. Benefits:

- Performance
- Availability
- Security
- Manageability

# Redis Enterprise Cluster Components




# Linear Scaling with Redis Enterprise




# Scaling Shards and Nodes

## Scaling Up




# Scaling Shards and Nodes

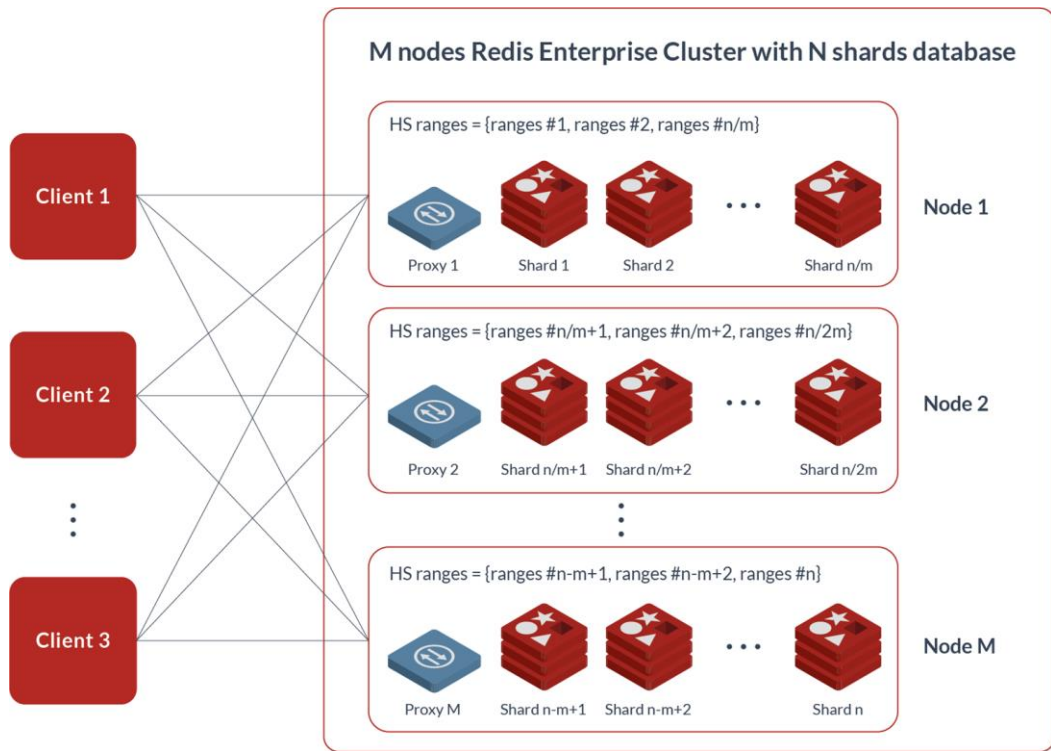
## Scale out, Rebalancing, Resharding



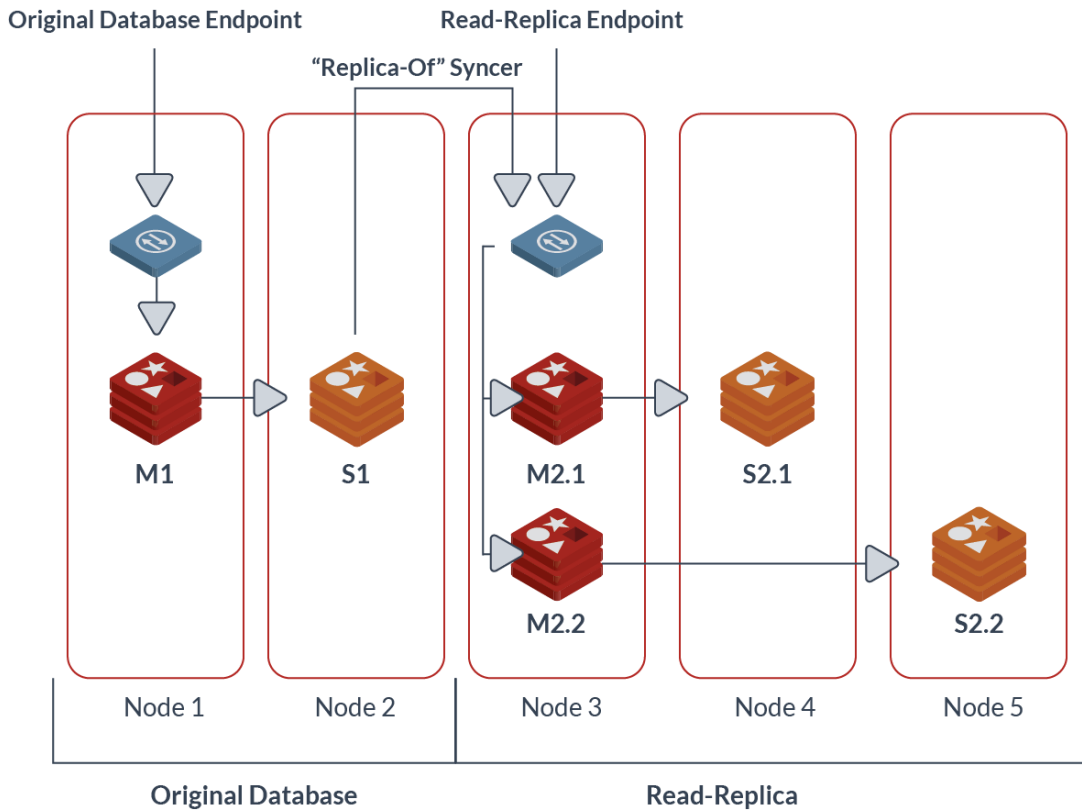
# Scaling Proxies



# Scalability with Open Source (OSS) Cluster API



# Scaling 'Read' using 'Replica-Of'




# Highly Available Redis

# Redis Enterprise: HA system with only two replicas


## The Cost of 90GB Dataset HA Deployment on AWS (reserved instances)

3 replicas Redis



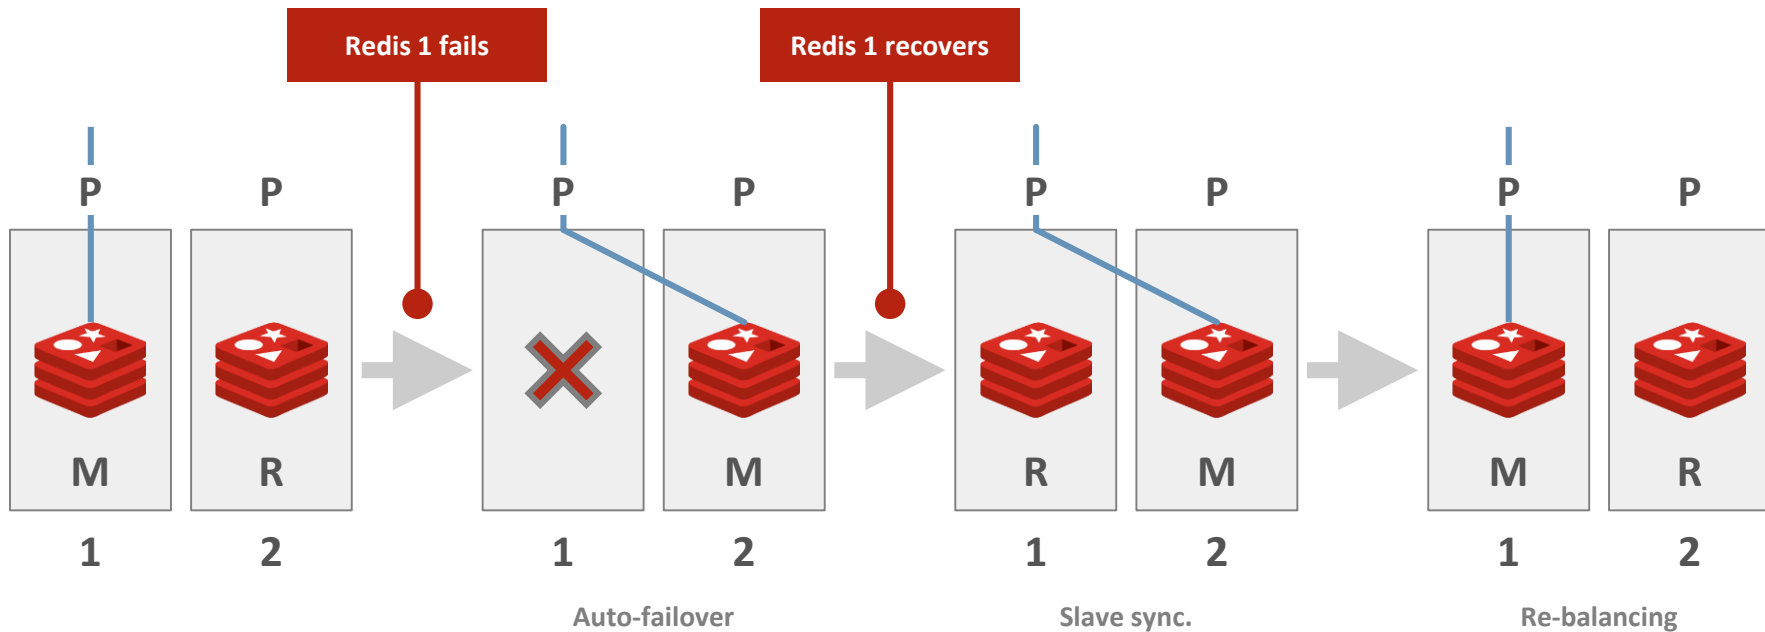
\$16,443/yr

Redis Enterprise




\$11,469/yr

# High Availability: Redis shard Failure

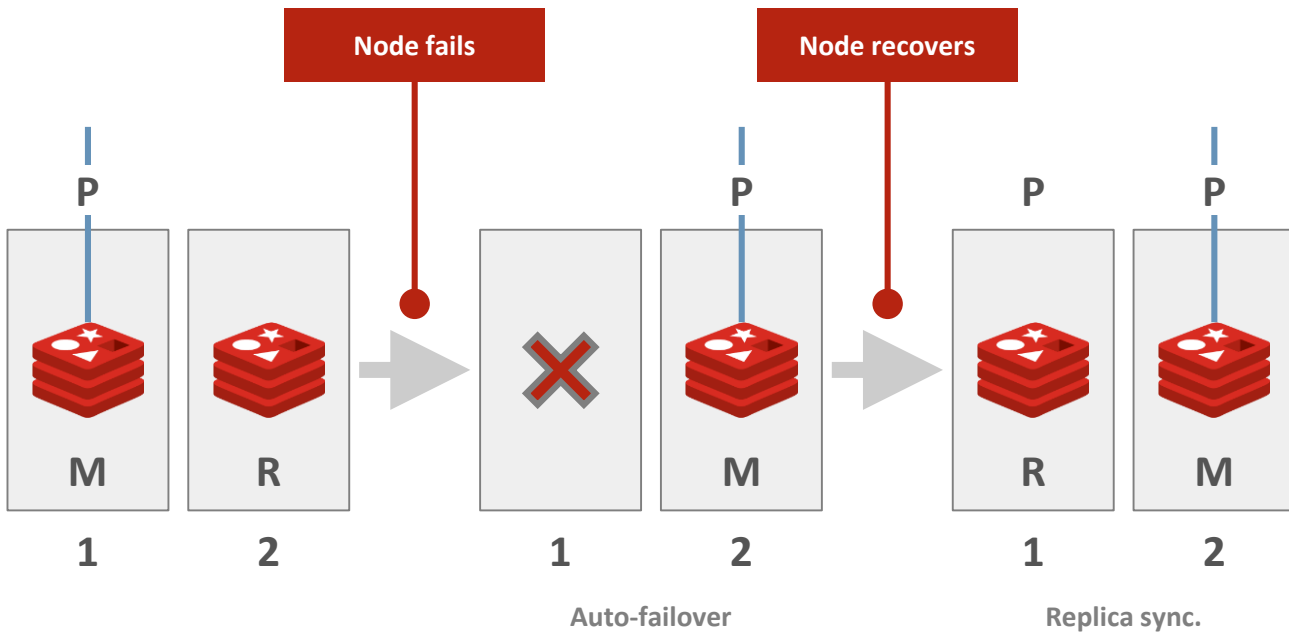




# High Availability: Proxy Failure



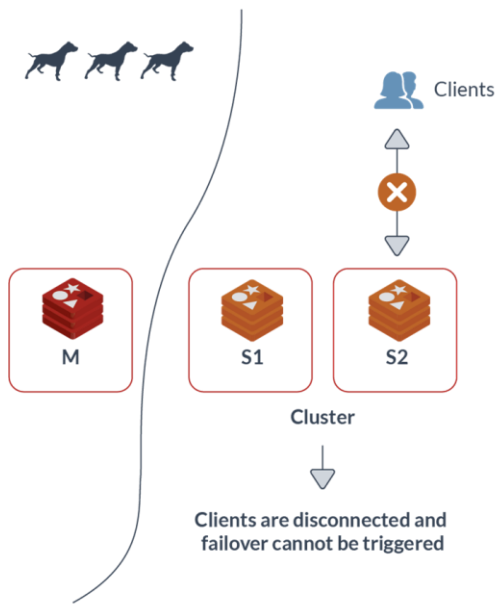
# High Availability: Node Failure



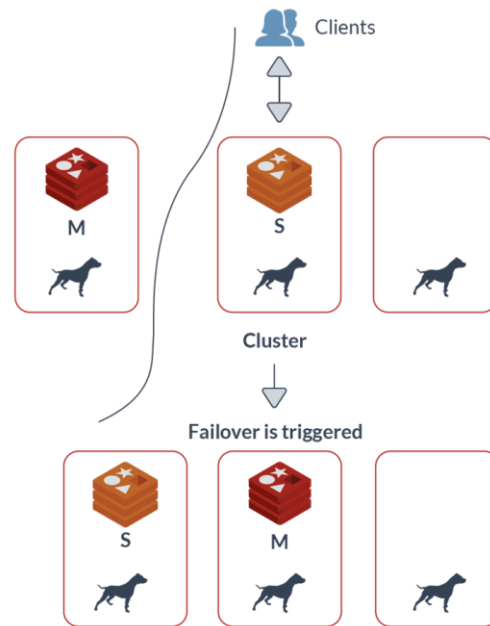
# Watchdogs

- **Node watchdog:** Monitors all processes running on a given node.
- **Cluster watchdog** – Responsible for the health of the cluster nodes and uses a gossip protocol to manage the membership of the nodes in the cluster.

Watchdogs are not running on the cluster's nodes




Watchdogs are part of the cluster



# Multi-AZ (Availability-Zone) / Rack Deployment

## Conditions required

- Three or more cluster nodes, with total number being uneven
- Three or more AZs/racks, with total number being uneven
- The number of nodes in a given AZ/rack should always be a minority
- Network latency between the AZ/rack should be <10ms




# Redis Enterprise Geo Distribution

# Active-Passive Geo Distribution

## Replica – of: unidirectional replication between DBs

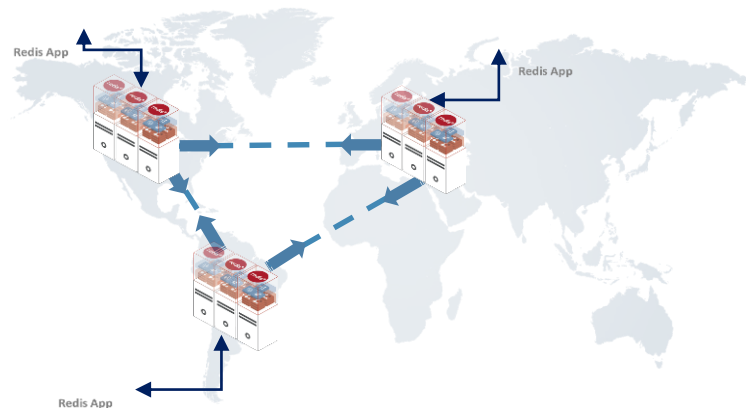
- Each source several destinations (one-to-many)
- Each destination replicated from several sources (many-to-one)
- Each destination can be a source (chained replication)
- Each database different database settings, and different cluster settings
- Automatic compression
- Optional encryption



# Active - Active Geo Distribution (CRDT-Based)

## CDRT (Conflict-free Replicated Data Types)

- **High Performance:** Read and write with low local sub-millisecond latency
- **Guaranteed data consistency:** CRDT based: The datatypes are conflict-free by design. All databases eventually converge automatically to the same state with strong eventual consistency. Supports causal consistency executing read and write operations in an order that reflects causality
- **Simplifies the app design:** Develop as if it's a single app in a single geo, we take care of all the rest



# Solutions that benefit from Redis Enterprise with Active-Active support

## Fraud Mitigation

- Geo Distributed Event Tracking: *Sets Gathering Geo Distributed Events*

## Social Engagement Apps

- Encoding Social Engagement: *Distributed Counters for “Likes”, “Shares”, “Retweets”*

## Collaboration Apps

- Constructing Smart Timelines: *Merged Lists Ordering Posts*
- Instant Messaging & Conversation Tracking: *Merged Lists Ordering Conversations*

## Geo Distributed Trading/Bidding

- Auctions, Bids/Asks: *Lists/Sorted Sets tracking Bids and Asks*

## Dashboards & Scoreboards

- Tracking Geo Distributed Scoreboards: *Sorted Sets tracking ordered scores*

## Real-time Metering Apps

- Tracking Usage/Consumption: *Sets/Lists Tracking Consumption Events*

And more.....



# Demo

# Demo

- DB Creation
- Client load (memtier\_benchmark)
- UI Console
- rladmin
- DB Scaling: sharding
- Failover scenarios

Not Secure | <https://node1.manuel.demo.redislabs.com:8443/#/nodes>

redis enterprise

by redis labs

Welcome manuel@redislabs.com

trial

version

cluster

nodes

databases

settings

log

Documentation

Support

Sign Out

nodes

Node ID / IP Address	Shards	Memory	Persistent storage	CPU	Network	Status
node: 1 / Multiple IPs ①	1	1.98 GB / 29.45 GB	1.62 GB / 38.71 GB	0.70%	14.1 KB / 80.83 KB	✓
node: 2 / Multiple IPs ①	2	1.7 GB / 29.45 GB	1.57 GB / 38.71 GB	0.60%	36.61 KB / 10.78 KB	✓
node: 3 / Multiple IPs ①	1	1.59 GB / 29.45 GB	2.29 GB / 38.71 GB	1.10%	43.58 KB / 9.01 KB	✓

© 2018 Redis Labs, Inc. All rights reserved.

```
f1ndmip> status
CLUSTER NODES:
NODE:ID    ROLE    ADDRESS          EXTERNAL ADDRESS    HOSTNAME             SHARDS    CORES    RAM              AVAILABLE RAM        VERSION    STATUS
#node:1    master  10.132.0.40      35.187.189.52       manuel-demo-1        1/100    8        27.42GB/29.45GB   28.95GB/24.15GB      5.4.0-19  OK
node:2     slave  10.132.0.42      144.148.24.159      manuel-demo-2        2/100    8        27.76GB/29.45GB   28.12GB/24.15GB      5.4.0-19  OK
node:3     slave  10.132.0.43      104.199.46.236      manuel-demo-3        1/100    8        27.85GB/29.45GB   21.39GB/24.15GB      5.4.0-19  OK
DATABASES:
DB:ID    NAME    TYPE    STATUS SHARDS PLACEMENT REPLICATION PERSISTENCE ENDPOINT
db:12    dbtest redis active 2      sparse enabled disabled redis-16041.internal.manuel.demo.redislabs.com:16041/redis-16041.manuel.demo.redislabs.com:16041
ENDPOINTS:
DB:ID    NAME    ID              NODE              ROLE              SSL
db:12    dbtest endpoint:12:1    node:1            single            No
SHARDS:
DB:ID    NAME    ID              NODE              ROLE              SLOTS              USED_MEMORY        STATUS
db:12    dbtest redis:38        node:1            slave             0-8191             90.29MB            OK
db:12    dbtest redis:39        node:2            slave             8192-16383          90.25MB            OK
db:12    dbtest redis:40        node:2            master            0-8191             90.24MB            OK
db:12    dbtest redis:41        node:3            master            8192-16383          90.24MB            OK
```

# Thank you!

[redislabs.com](https://redislabs.com)

