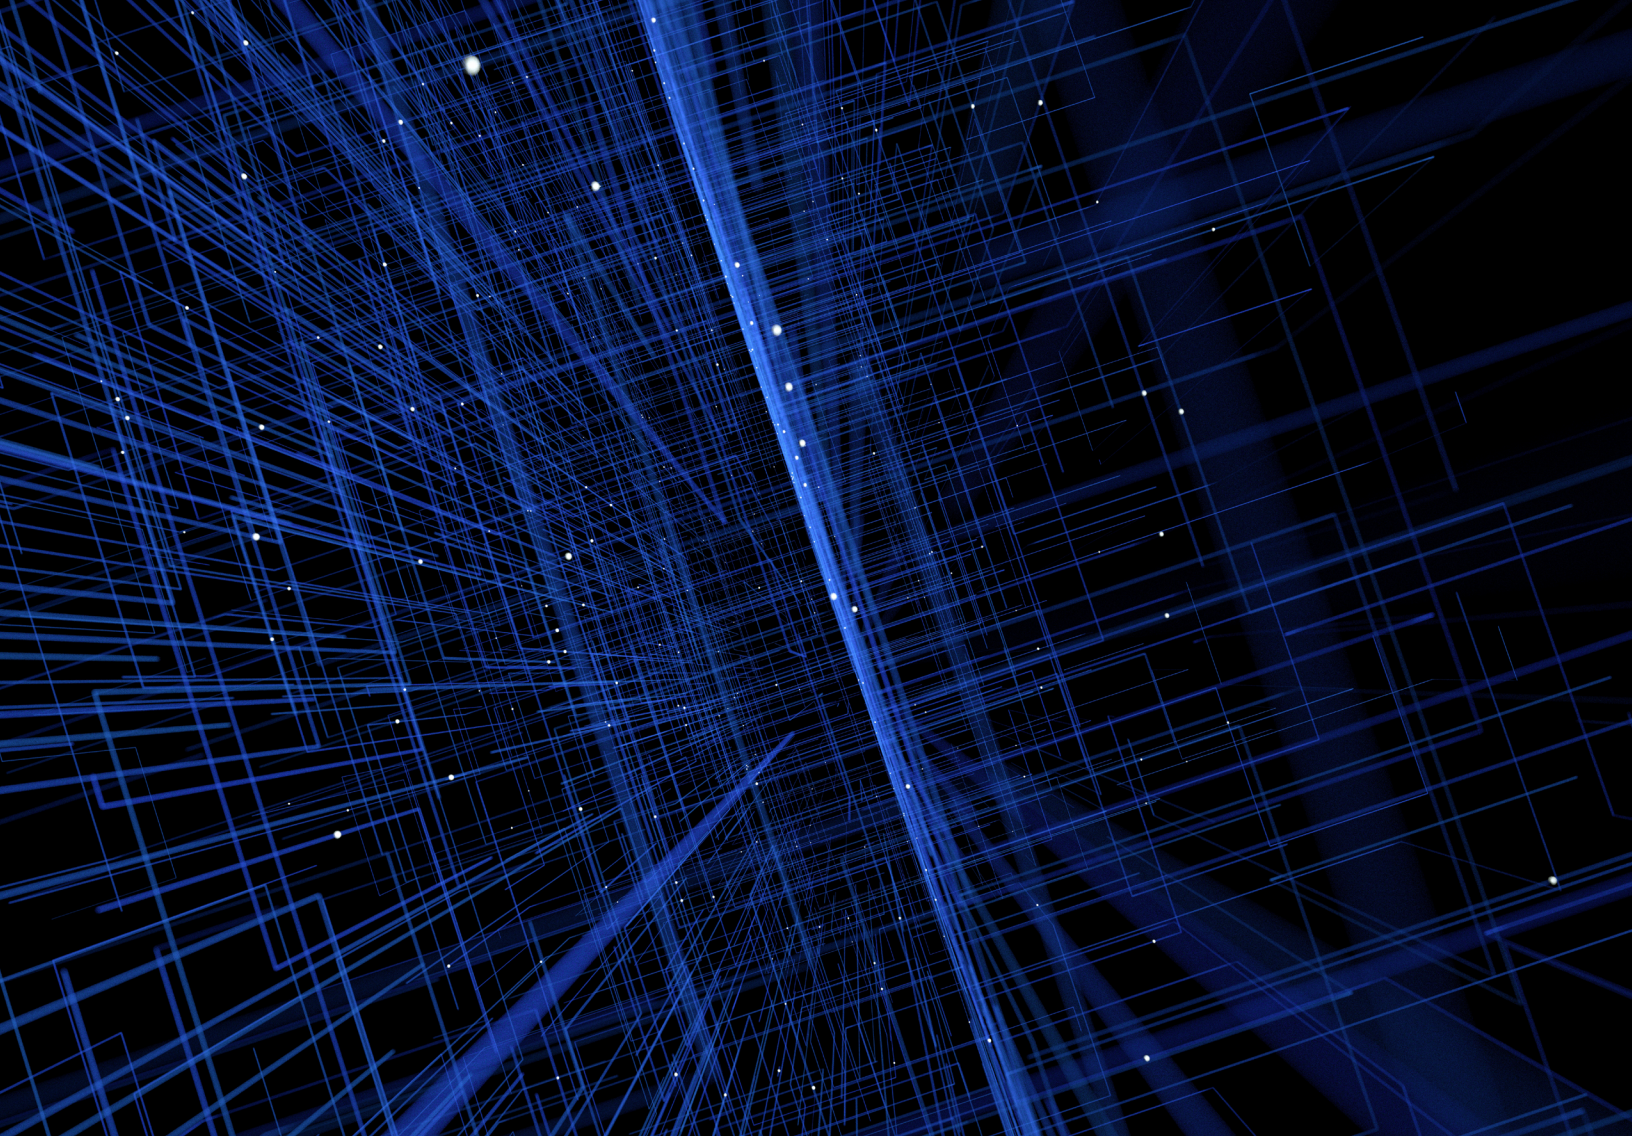

SOLUTION BRIEF

Real-Time Search with Redis Enterprise

Enhance customer experience with immediate results

redis

Unleash the power of data for your business with a real-time indexing, querying, and full-text search engine. Redis Enterprise includes a flexible search engine that delivers sub-millisecond results for enhanced customer experiences and better business intelligence.

Search engine for modern applications

No one wants to wait. Developers and architects who build modern applications turn to real-time search to provide the performance required to keep their users happy. Redis Enterprise is highly optimized to deliver fast results by indexing for queries such as full-text search, complex filter expressions, secondary key lookups, numeric or geo-ranges, aggregation functions, and ranking of search results. Applications built for cybersecurity, fraud detection, financial services, e-commerce, gaming, and many others rely on Redis Enterprise for instant search results at scale.

What are the benefits of Redis Enterprise real-time search?

TBusiness and customer expectations have never been higher. Your competition is always just a click away. Redis Enterprise provides a flexible, real-time search engine that delivers:

- **Enhanced customer experience:** Immediate and accurate search results keep customers engaged and loyal. Real-time search, auto-suggestions, and faceted search help users locate what they are looking for quickly.
- **More accurate business intelligence:** Redis provides instant search results for real-time analytics for process and service optimization. Redis is capable of high performance for immutable and mutable data performing well in heavy read and heavy update use cases.
- **Lower costs:** Do more with less infrastructure thanks to the in-memory database technology and built-in search engine.
- **Support for microservices frameworks:** Redis Enterprise low-latency search results support modern applications that rely on intra-microservice communications where any additional latency is detrimental to user experience.
- **Scalability:** Linearly scale ingest and distribute database indexes across shards and geographies to support user search of tens of millions of keys in real time.
- **Geo-distribution:** Redis Active-Active Geo-Distribution is a data resiliency architecture that distributes the database information over multiple data centers via independent and geographically distributed clusters and nodes. This ensures local low-latency on read/write operations, regardless of the number of geo-replicated regions and their distance from each other.

Redis search overcomes common application challenges

- **Network Latency:** Network latency is most often the major challenge for real-time search applications—especially if data has to travel far. Redis Enterprise can run geo-distributed on cloud and on premises, providing data locality and ensuring low-latency search results anywhere around the globe.
- **Ingest performance:** To search fast you need to ingest and index data fast. Ingesting massive amounts of data and scaling to millions of writes per second with sub-millisecond latency is impossible for traditional disk-based databases. Redis Enterprise runs in-memory for fast ingest performance and scalability.

- **Freshness:** Modern applications require current and immediately usable data. Slow search engines force developers to pre-calculate, pre-aggregate, or manipulate the data before using it. This causes data to become stale, losing accuracy and credibility. Stale data invalidates operational or real-time analytics. Redis delivers fresh data with sub-millisecond search results for accurate intelligence
- **Concurrency:** Data concurrency allows multiple users to process reads and writes simultaneously within a single database structure. It's common to find analytic use cases that need 1,000+ queries per second (QPS) performance. At the same time, a common developer goal is to provide application users with sub 200ms response time to avoid performance complaints. This is where Redis Enterprise in-memory database supports real-time search use cases when other disk-based search engines fail. An in-memory database can consistently scale search and query with low-latency response times distributed across database shards, providing high concurrency and data protection.

Better developer experience with Redis indexing

With Redis secondary indexing, developers can index any field to create unique data views and provide more accurate search results. It also makes it easy to create and manage complex multi-field queries with no application code changes. Once defined, Redis indices are automatically updated.

- Secondary indices of Redis Enterprise data can be created and queried with no changes to application code.
- Eliminates need to manage and distribute sorted sets to provide fast queries.
- Global auto-indexing keeps indices fresh everywhere.
- Maintains indices automatically and allows you to query secondary keys and across data structures in a clustered database.
- Provides several indexing strategies for the value part of the key, including full-text, geo-location, numbers, and tags.

WITHOUT REDIS SEARCH

Data stored via Hash

HMSET person: 1 id 1 username John age 25
HMSET person: 2 id 2 username John age 42
HMSET person: 3 id 3 username John age 54

Multiple Indices via Sorted Sets

ZADD nameldx 0 John: 1
ZADD nameldx 0 Jane: 2
ZADD nameldx 0 Joey: 3
ZADD ageldx 25 1
ZADD ageldx 42 2
ZADD ageldx 54 3

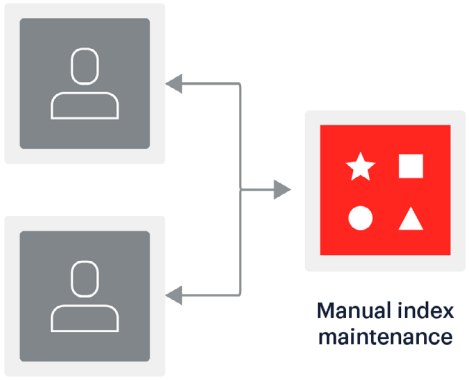

Manual index maintenance

Cumbersome multi-key search

ZRANGE ageldx 20 30 **BYSCORE**
ZRANGE nameldx [Jo + **BYLEX**...plus logic to perform set intersection...

WITH REDIS SEARCH

Native JSON storage

JSON.SET person: 1 $'{"username":"John","age": 25}'
JSON.SET person: 2 $'{"username":"Jane","age": 42}'
JSON.SET person: 3 $'{"username":"Joey","age": 54}'

Multiple secondary keys, Native Index

FT.CREATE idx **ON JSON PREFIX 1** person: **SCHEMA**
$.username **AS** username **TEXT SORTABLE**
$.age **AS** age **NUMERIC SORTABLE**

Native, optimized search

Automatic index maintenance

FT.SEARCH idx "@age[20 30] @username:Jo"

Scaling applications and analytics with real-time search

There are three common application use cases where low-latency search is a must:

- **Real-time analytics**
- **Master data table lookups**
- **Customer 360°**

Real-time analytics

Traditional databases are too slow to provide fresh data for real-time analytics and business intelligence. When data is constantly changing, their queries become slow and can't scale. For accurate real-time analytics, consistent, high-performing search results are required. Redis Enterprise can act as an in-memory data fabric, consolidating siloed sources of record databases to provide real-time metrics to answer questions such as "What is the top product purchased in the past 10 minutes?"

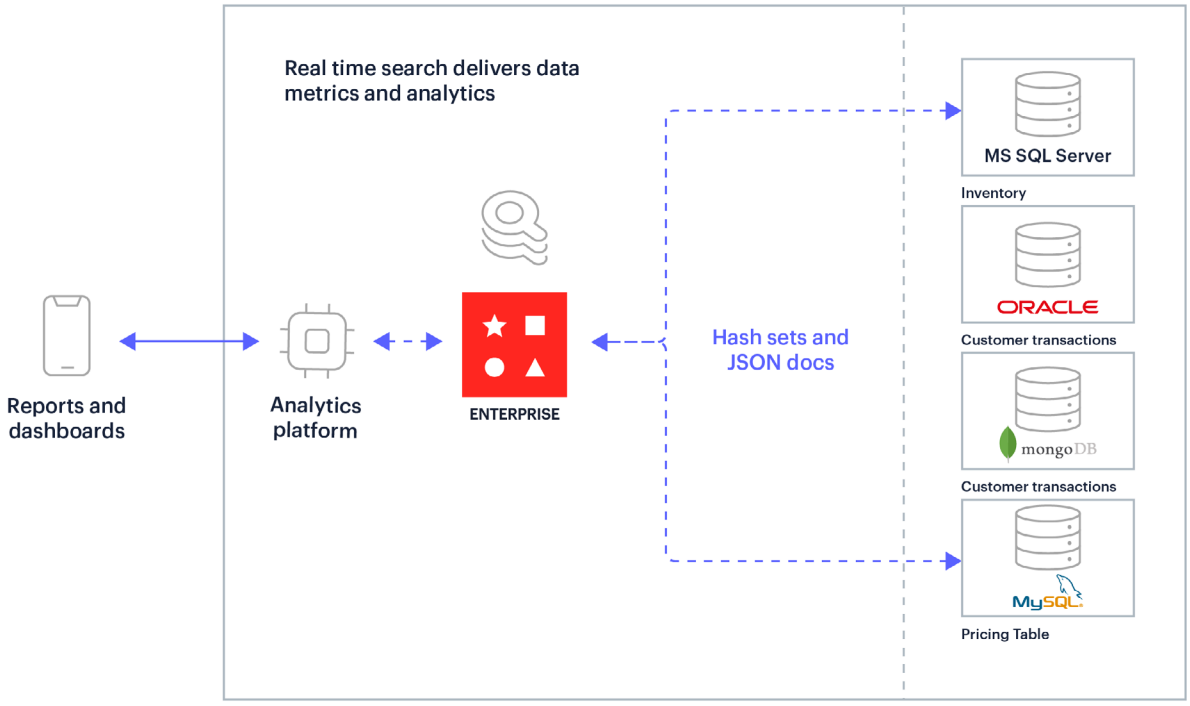

Master data table lookups

Many applications, such as e-commerce applications, require repetitive, high-speed lookups of master data tables. For example, a product master table might contain every product ID, product name, description, and price. As users shop, they are all simultaneously accessing the product master data table to calculate the value of their carts. As the tables grow, the search results tend to get too slow and customer experience suffers.

Redis Enterprise easily eliminates these common application bottlenecks by scaling real-time search. Supporting tens of millions of unique keys, Redis Enterprise enhances customer experience with intuitive and fast auto-suggestion and full-text search capabilities that delight online customers.

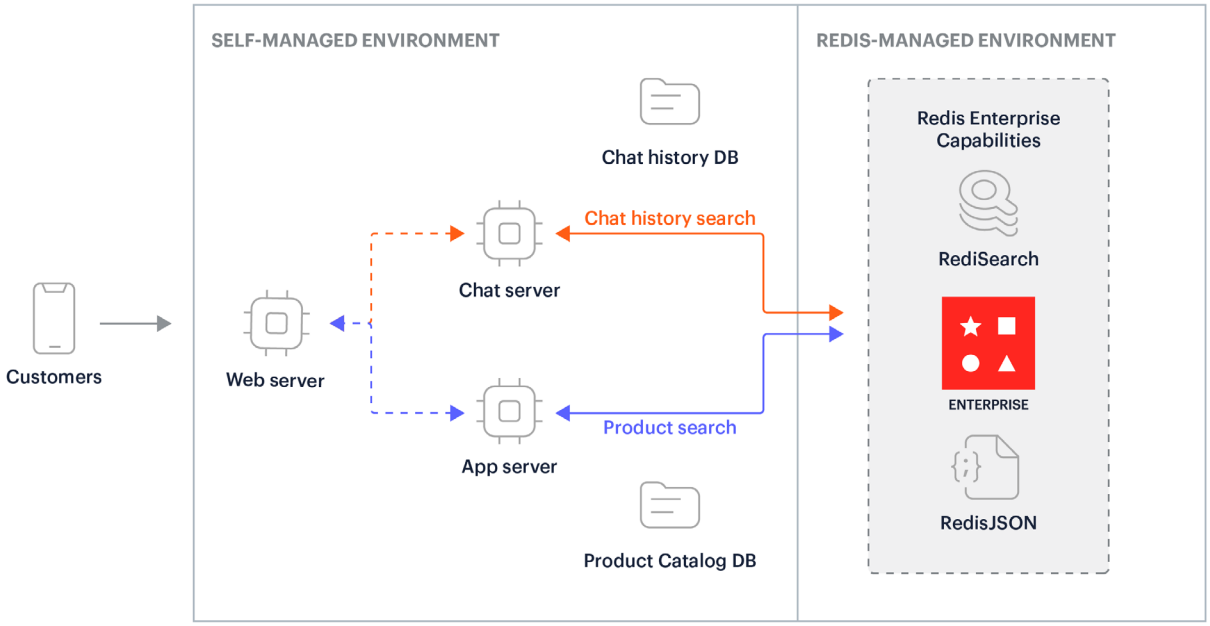

Customer 360°

Delivering real-time customer service is a challenging task. We have all been on hold waiting for customer service at one point. Customer service teams and Customer 360° applications require real-time information about customers. Redis Enterprise provides sub-millisecond customer lookups to scale online applications to tens of millions of users. Create new customer insights and summary views in real time with powerful commands like aggregation, filter, groupby, and sortby.

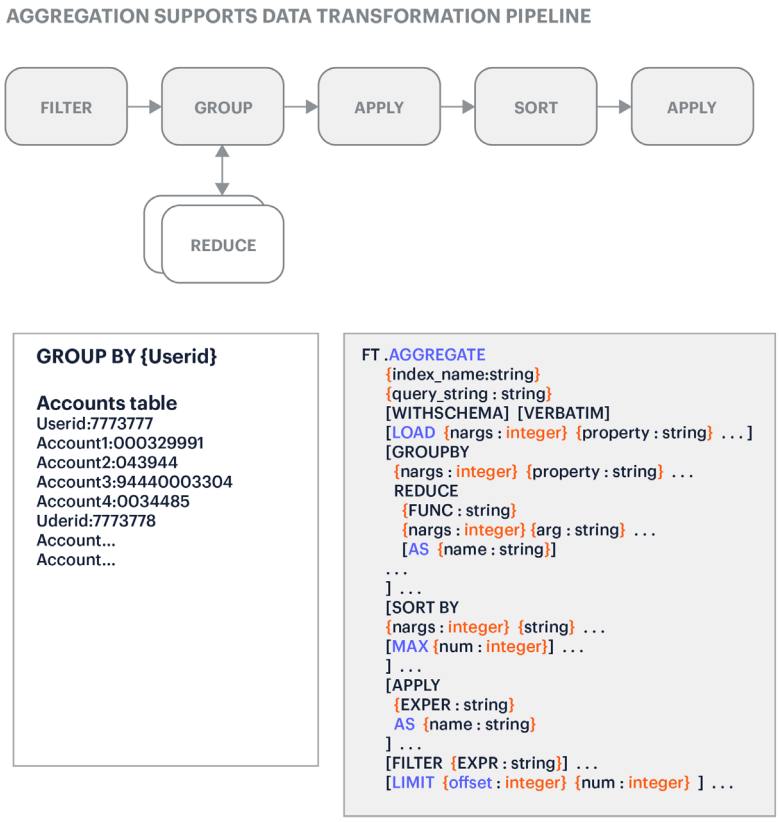

Results: Better customer experience from dynamic aggregated customer views

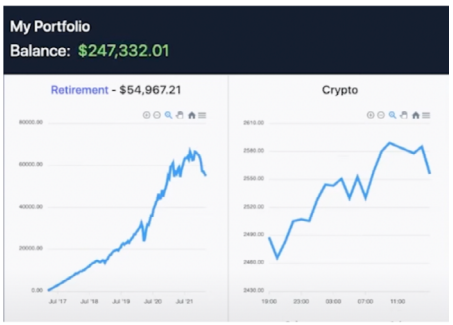

Want to learn more about real-time search?

Get started today creating enhanced customer experiences with a global unified in-memory data layer and search engine, Redis Enterprise.

- Book a meeting with Redis search expert
- Try Redis Enterprise Cloud with real-time search for free
- Experienced with Redis? Start here: Getting Started with Redis Search