Real Time Search

Deliver search and analytics in real time

Redis Enterprise provides a powerful indexing, querying, and full-text search engine for real-time data, available on-premises and as a managed service in the cloud. Use it as a secondary index for data hosted in Redis; consolidate data from other data stores, as a fast text search or auto-complete engine, as well as an engine for lightspeed aggregations and faceted queries.

Customers need answers fast

Expectations from customers and business users are extremely high. A scalable, high performance search engine delivers the customer experience required to keep them coming back. Redis Enterprise can help you deliver it.

Existing full-text search engines don't support real-time search

Most search engines are disk-based and depend on slower Java virtual machines for their functionality. That means they are unable to support real-time search, as they're held back by inefficient runtimes, slow writes, and unpredictable performance, especially during garbage collection.

Unique customer experiences require flexible search performance

Users can search on any field with fast and flexible secondary indexing. They can quickly find what they are looking with faceted, multi-category, search filtering on secondary keys like size, make, price, year, etc.

Real-time search use cases

Redis Enterprise real-time search capabilities are ideal for building modern applications. They help you deliver real-time performance for demanding use cases where responsiveness is paramount, and measured in sub-millisecond timelines even with huge data sets. Among those use cases are ecommerce, real-time analytics, master data table lookups and personalized customer 360° applications.

Ecommerce

Deliver accurate inventory records to enhance shopper experience. Deliver enhanced customer search experience with context aware auto-suggestions and fuzzy matching results.

Real-time analytics

Deliver accurate data for analytics with real-time performance and aggregation over millions of unique keys.

Master data table lookups

Eliminate application bottlenecks and scale master table lookups with sub-millisecond response on even the largest tables.

Customer 360° and chat history

Deliver accurate data for analytics with real-time performance and aggregation Deliver better quality customer service with super-fast access to customer records, transactions, and chat histories

> "GoMechanic needed a fast data layer to search across 10 million spare parts and to power our customer service chats. RediSearch on Google Cloud was not only four times faster than our existing database but also extremely robust, easy to scale, and inexpensive."
>
> **Prasenjit Singh, Vice President of Engineering, GoMechanic**

Key Features

Fast Search and filter

- Full-text indexing of multiple fields in documents
- Geo filtering using Redis' own geo-commands
- Retrieve full document content or just IDs
- Document ranking based on fields weights or ad hoc

Powerful querying

- Comprehensive, easy-to-use query language
- Support for prefix, fuzzy, synonyms, and phonetic searches
- Stem-based query expansion in multiple languages (using Snowball)

Active-active geo-distribution

- Delivers low latency, globally federated search through data locality, multi-site, active-active configurations
- Automatically indexes and copies data across
- Delivers reliability, high performance, and continuity in any scenario

Fast incremental inserting

- Partial and conditional document indexing in atomic operations
- Document deletion and updating with index garbage collection

Aggregate

- Use real-time aggregation to create new dynamic summary customer views, dashboards, and reports
- AFast and accurate queries with group by/ reduce, map, sort, and filter capabilities
- Delivers reliability, high performance, and continuity in any scenario

Auto-complete

- Context-aware suggestions
- Fuzzy prefix matching