

RediSearch

Build Modern Applications with Interactive Search Experiences

RediSearch is a powerful indexing, querying, and search engine that enables you to query your Redis data to answer a wide variety of complex questions. Use it as a secondary index for data hosted in Redis or search across data in other data stores, as a fast text search or auto-complete engine, as well as an engine for light-speed aggregations and faceted queries.

Customers are demanding real-time answers

Making data searchable in Redis is burdensome

Redis lacks out of the box secondary indexing. Creating and maintaining indexes manually is complex and error-prone.

Existing full-text search engines are unable to support real-time use cases

Existing search engines are disk-based and depend on JVM heap for their functionality. Due to inefficient runtimes, slow writes, and unpredictable performance, especially during garbage collection, they are unable to support real-time use cases.

RediSearch unlocks your data at lightning speed

Purpose-built for high performance

- Uses in-memory data structures and modern distributed indexing algorithms
- Supports low latency indexing and ingestion
- Can scale out and partition indexes over several shards and nodes

Index and query data in Redis

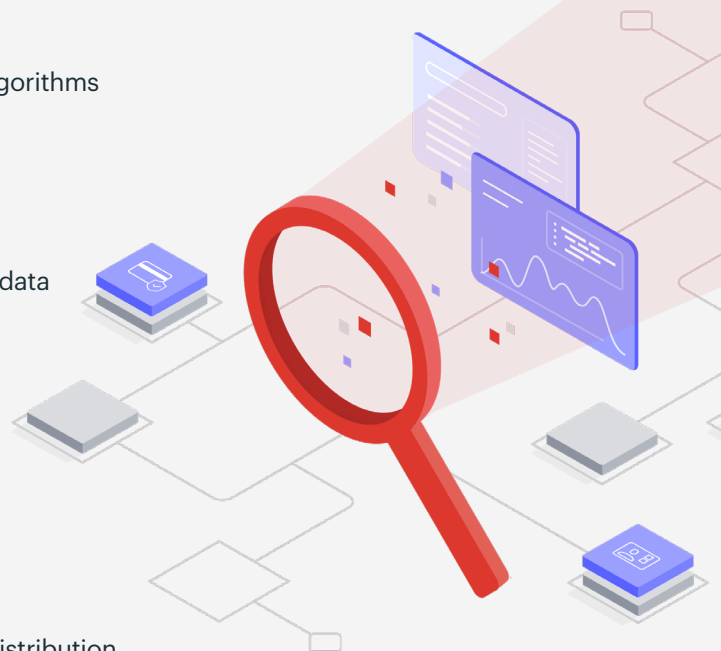
- Rapidly create, populate, and expire indexes for frequently updated data
- Rich query language supports complex aggregations, and can filter by properties, numeric ranges, and geo distance

Perform full-text and fuzzy search

- Search through relevant documents in the shortest possible time
- Enrich search experiences with context-aware suggestions and 'fuzzy' searching.

Enjoy enterprise-grade availability and scalability

- Enjoy continued operations in any scenario with Active-Active geo-distribution
- Available on premises and in the cloud, including multicloud and hybrid cloud deployments
- Fully managed support for AWS, Google Cloud, and Microsoft Azure



Use Cases

Redisearch is ideal for building modern applications, delivering real-time performance for use cases such as real-time inventory, content recommendations, customer insights, dynamic catalogs, secondary indices for large databases, and more.

Enterprise Search

Deliver fast search for site, app, and workplace search.

External Database Indexing

Build relationships between records and perform aggregate searches while keeping raw data in external databases.

Real-time inventory

Search for products in stock with real-time performance and aggregation over millions of stores and SKU combinations.

Customer service and insights

Look up customer history during support interactions or match specific profiles and behaviors to provide relevant help.

Key Features



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



Rich querying

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



Active-Active geo-distribution

- Redisearch can work as a multi-site, Active-Active search engine
- Automatically indexes and copies data across replicas
- Delivers reliability and continuity in any scenario



Aggregate

- Intuitive aggregation operations can be combined in any order across millions of documents
- Operations can groupby/reduce, map, sort and filter



Fast incremental inserting

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



Auto-complete

- Context-aware suggestions
- Fuzzy prefix matching

“GoMechanic needed a fast data layer to search across 10 million spare parts and 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



Get started today on Redis Enterprise Cloud

redislabs.com/modules/redisearch-quick-start