


Fastest Multi-Model Database For Modern Applications and Instant Experiences

Large-scale computation-intensive applications use data in new ways to influence business-critical decision making. Digital footprints are growing, with massive amounts of data coming from different sources, in structurally variable formats. There's an increased need for systems that capture, store, visualize and analyze complex data with heterogeneous structures. While best-of-breed solutions may solve a particular business problem, fragmented environments create development complexity, increasing total cost over time.

Redis Enterprise delivers a unique serverless, multi-model solution that harnesses any data at high velocity, delivering complete interoperability at record-setting speed. This helps organizations break down data silos, simplify operations and build modern applications that deliver instant experiences.

Redis Enterprise's key building blocks

Redis Enterprise supports multiple data models and structures, so you can iterate applications quickly without worrying about schemas or indexes.



Gain ultimate flexibility and extensibility with Redis Modules

Search and Query

- Implement document and field scoring, numeric filtering, exact phrase matching, query expansion and advanced search operations.
- Use as a secondary indexing, suggestion and auto-complete engine.
- Outperform traditional search engines by 5X with both simple and complex full-text search.

JSON

- Organize data in a hierarchical, tree-like format.
- Manipulate JSON documents with serverside queries.
- Store complex nested schemas and access them with intuitive commands.

Probabilistic

- Use space-efficient probabilistic data structures to check set membership with certainty.
- Reduce code complexity and easily implement scalable bloom filters and cuckoo filters with out-of-the-box commands.
- Perform bulk operations speedily to more efficiently process data.

Time Series

- Capture, analyze and store metrics, events and time-based data from hundreds of data sources.
- Intelligently reduce the number of events to keep high precision data with out-of-the-box downsampling, aggregation, compression and configurable retention policies.
- Classify time-series data with labels and derive analytics and insights instantly.

Support heterogeneous data through versatile data structures




Redis' **data structures** support various key-value data types so you can develop applications faster with whichever data structure fits your needs. With Redis, developers gain the ability to run direct data retrieval commands with no queries or query planner abstractions in the middle. Its native structures include Strings, Hashes, Lists, Sets, Sorted Sets, Bitmaps, Hyperloglogs, Geospatial Indexes and Streams. Redis **Streams** collects large volumes of data arriving in high velocity (millions of keys per second) and stores this event data with full persistence-making data ingest, stream processing and stream consumption faster and easier.

Effortlessly connect data across all models

Manage multiple data models at the same time and easily connect all of these building blocks while operating on a single copy of your data set in a serverless environment. Using Redis Hash as a shared data structure, you can seamlessly extend or combine queries across all data in varying formats. With **Triggers and Functions**, automatically transform and process data, agnostic of its topology, without application-level heavy lifting. React to data changes, initiate functions based on preconfigured event triggers, and facilitate inter-model communication.

How Redis Enterprise benefits your modern applications

Redis Enterprise empowers application developers to choose the right model for each specific use case-e.g. key-value, document, graph or time-series data. The database effortlessly extends to support probabilistic data structures, comprehensive search functionality, stream processing, deep learning and AI.

 Easily deliver an instant experience	 Accelerate time-to-market	 Cut costs through reduced complexity
<p>Rapidly collect, process and analyze enormous volumes of variably formatted data at scale:</p> <ul style="list-style-type: none">• Linear scalability delivers consistently high and predictable sub-millisecond performance on large data sets.• Unified processing allows you to extend and combine queries across all data models without compromising data integrity.	<p>Automatically transform and blend data to execute business logic in response to data changes:</p> <ul style="list-style-type: none">• Flexible support for multiple models lets you store data in its natural format with native indexing, parsing and processing.• Real-time data interoperability provides different lenses on your data to power continuous analytics, search and artificial intelligence.	<p>Simplify application development and operations:</p> <ul style="list-style-type: none">• Globally distributed multi-region deployments enable cost-efficient multi-master replication with tunable consistency models and local latency on read and write operations.• With one multi-model database on any cloud, you can effortlessly and horizontally scale across all public clouds and on-premises, including multi-cloud and hybrid deployment options.

Get started today!

Contact a Redis expert: expert@redis.com

Try Redis for free: www.redis.com

