

Nine Essential Database Capabilities

To keep up with the evolving needs of modern enterprises, your database has to do more than ever—starting with these capabilities.



Today's databases have to do more



For the last 30 years, a deluge of transformative trends has continually remade and accelerated the database industry—and changed the very nature of the databases in your organization's technology stack. From relatively staid mainframe programs only handling carefully constructed data to modern microservices architectures working with a wide variety of unstructured data, database technology continues to morph and leverage new technologies, tackle ever harder problems, and power even more mission-critical use cases.

To fully understand what's happening in the database world right now (and how it affects your technology choices), it's important to work through the impact of these trends. Today's databases face unprecedented challenges. They must work with mountains of disparate data, return results in the blink of an eye, and cope with hugely complex technology environments, all while remaining easy to work with and highly available. While many databases exist that have just one or even a few of these critical features, having all nine is vital to successfully manage the challenges they are facing today.

Real-Time Performance

Modern databases have to respond in less than 1 millisecond to avoid becoming a performance bottleneck.




Real-Time Performance

In today's digital economy, application performance is business performance. Your customers demand a real-time response from applications, and if they don't get it, they'll jump to a competitor.

But what does “real-time response” really mean? From a user perspective, real-time is defined as the actual time something happens, meaning there has to be no perceivable delay. Research into human response indicates applications have roughly 100 milliseconds (ms)—one third of the time it takes to blink—before users feel like they're waiting for a response. Except a round trip from when a user makes a request to when they receive a response can easily take 100ms. This leaves somewhere between 0ms and 1ms for the database to respond.

Unfortunately, your legacy databases running on good old-fashioned spinning disks can take up to 100ms to respond. Your team can significantly decrease latency by adding a caching layer such as open source Redis, but adding a caching layer also increases the architecture complexity. Some modern databases are faster, but still need a cache because they still take 10–100ms to deliver results.

To meet customer needs, deliver real-time performance, and have a simplified, enterprise-enabled architecture, more and more companies are turning to in-memory databases such as Redis Enterprise, which can respond in less than 1ms with no need for a separate caching layer (though Redis Enterprise is still often used for caching on other databases).



Scalability

Linear scaling is critical to dealing with overall growth and seasonal surges without requiring massive infrastructure investments.



Scalability

Modern applications are increasingly distributed, and you must plan for cost-effective ways to scale your database to meet overall growth while also coping with seasonal surges and unexpected demand peaks. **Linear scaling of database performance** (performance increases are proportionally equal to the amount of infrastructure capacity added) is critical to making that possible without massive infrastructure investments. While many cloud and on-premises databases claim the ability to scale linearly, they often struggle to prove it with solid benchmarks.

Redis Enterprise enables infinite linear scaling by simply adding shards and nodes. In 2019, Redis demonstrated Redis Enterprise extending its **industry-leading linear scalability** with 200 Million ops/sec with less than 1ms latency on only 40 AWS instances.



High Availability

The inability to quickly recover from a database failure can negatively impact the business and its reputation.



High Availability

As your organization becomes more reliant on applications for mission-critical use cases, those applications, and the databases behind them, have to be available at all times. And unfortunately, no matter how reliable your systems and software, problems can occur. Downtime can reduce revenue, raise costs, and degrade the customer experience, but that's only part of the problem. The inability to quickly recover from a database failure can result in loss of data and millions of operations. And for mission-critical use cases (like online orders or financial transactions), that's unacceptable.

Redis Enterprise offers **uninterrupted high availability** with five-nines (99.999%) uptime. By using technologies like diskless replication, instant failure detection, and single-digit-seconds failover across racks, zones, and geographies, Redis Enterprise delivers high availability while being more cost effective than competing technologies.



Cloud Native

(including hybrid cloud, multcloud, and inter-cloud)

Have maximum flexibility
and choice in your
deployment options.



Cloud Native (including hybrid cloud, multcloud, and inter-cloud)

With legacy databases, you couldn't just buy it, install it, and run it. You bought it and then hired an army of skilled database administrators to install, run, and support it for you. The rise of cloud computing changed everything. Now that your databases are provisioned in the cloud, you no longer have to manage the internal workings. With a fully managed database service (a Database-as-a-Service, or DBaaS) delivered via the cloud, you get an endpoint to send and retrieve data from as needed—and someone else deals with everything beyond that.

But taking full advantage of the cloud is about more than just buying a DBaaS. To gain maximum cloud benefits, every element of your technology stack has to work with a wide variety of resources delivered via the cloud. Your databases have to play with microservices architectures as well as containers and container orchestration systems like Docker and Kubernetes.

And working well with a single cloud provider is only the beginning. Your databases must be able to work across multcloud, hybrid cloud, and inter-cloud environments.

The cloud enables streamlined global distribution of apps, but this comes with additional database challenges. As databases are distributed globally, the question of how remote data consistency is achieved. [Redis Enterprise Active-Active](#) geo-distribution replication enables global applications to provide real-time consistency at local latencies. [Active-active geo-distribution](#) allows you to place your Redis database replicas close to where your users are, no matter where they are. The eventual consistency is achieved via conflict-free replicated data types, a family of replicated data types with a common set of properties that enable operations performed on them to always converge to a final state consistent among all database replicas.

[Redis Enterprise Cloud](#) is offered as a fully managed serverless cloud database service on [AWS](#), [Google Cloud](#), and [Microsoft Azure](#). Redis Enterprise Cloud enables real-time data, at any scale, anywhere.



4

Tiered-Memory Support



Database designers are turning to tiered storage methods, putting the hottest data in DRAM while warm data resides on lower cost memory.

Tiered-Memory Support

Given the widespread and growing need for real-time performance at scale, many companies are looking to store their data in faster in-memory database architectures. But this isn't always a financially viable option given the cost of DRAM and the size of the data set.

One solution is to tier data to alternative memory technologies that deliver memory-like performance at significantly less than the cost of DRAM. There are a variety of options, including Intel NVMe, but the key requirement is that the automated, intelligent tiering ensures that the sub-millisecond latencies are retained.

By keeping the hottest data in memory while warm data resides in persistent memory options, like flash or Intel Optane NVMe, you can get cost savings of over 70%. This significantly expands the use cases for an in-memory database such

as Redis Enterprise. With intelligent tiering, you can tackle a variety of use cases, such as managing a significant growth in my game users, or being able to expand personalization in applications. Applications that have large amounts of data and require hundreds of thousands of operations per second with sub-millisecond latency.

Redis has long been a leader in [tiered-database storage](#), supporting Redis on Flash since 2016. In 2019, Redis Enterprise announced support for Intel Optane DC Persistent Memory, which delivers a new persistent-memory tier between DRAM and SSD. The combination of these two technologies make it cost effective to serve and analyze very large datasets in real-time using in-memory databases.



Simplicity and Extensibility



Because no single data model works well with all kinds of data, you can end up deploying multiple databases in the same organization—or even the same application.

Simplicity and Extensibility

Developers and operations teams increasingly seek operational simplicity in their databases. After all, we all want to use solutions that are easy and optimal for our particular use case. But because no single data model can work well with all kinds of data, you can end up deploying multiple databases in the same organization or application. And nobody wants that.

As a true [multi-model database](#), Redis Enterprise offers multiple data structures in the same database, cutting complexity and overhead by reducing the need to provision and maintain a different database for each data model required. Redis Enterprise incorporates 10 different data structures (including key-value, document, graph, and time-series data models) in a single, highly versatile multi-model database.

Much of that versatility comes from Redis Enterprise features like [Search and Query](#), [JSON](#), [Time Series](#), [Probabilistic](#), and many others. Redis Enterprise features provide ultimate flexibility and extensibility to effortlessly support probabilistic data structures, comprehensive search functionality, stream processing, deep learning, artificial intelligence, and more.

Developers know that Redis is extraordinarily easy for them to use, which is a key reason they love it. On Stack Overflow, Redis has been rated the [#1 most-loved database](#) for five years running. And that means there's a huge number of experienced Redis developers with the skills required to create powerful new applications.



NoSQL for the Future



SQL databases will never go away, but NoSQL is the future—that's where your most innovative and valuable use cases are being addressed.

NoSQL for the Future

Traditionally, SQL databases model highly structured data in tables. With the rise of the internet and the explosion of new kinds of data, that's no longer good enough. Most of the data in your databases today probably aren't structured—and would be very difficult to organize in a table. Also, the lack of flexibility and high latency of SQL databases means they're unable to meet the requirements of online applications and meet user expectations.

To move beyond that, modern [NoSQL databases](#) put unstructured data in a document, then use key-value data storage and strong search capabilities to find the right document and the right item in the document.

Does that mean SQL is dead? Of course not! There will always be cases where SQL databases are the best way to order the data, not to mention that it can take decades for legacy solutions to age out of the market. Increasingly, though, NoSQL is the future—especially cloud-native NoSQL. According to [Redis' 2021 Digital Transformation Index report](#), usage of every type of NoSQL database is on the rise, while relational (SQL) usage is stagnant.

Redis was recognized in the inaugural [2020 Gartner Magic Quadrant](#) for Cloud Database Management Systems. Redis was positioned furthest to the right on the Completeness of Vision axis in the Challenger quadrant.



Developer Tools



There's increasing demand for a graphical user interface to manage the database and let developers explore and interact with their data visually.

Developer Tools

In addition to being simple and versatile, your databases must also be easy for developers to use. As maximizing the value of data becomes even more important, it's critical that your database makes it easy to quickly massage and analyze that data.

While command line interfaces (CLIs) remain the most common way to work with a database, there's increasing demand for a graphical user interface to make it easier to manage

the system and support analysts who want to explore and interact with their data visually.

RedisInsight is a free, consolidated tool combining a CLI and GUI to make it easier to write commands for all Redis data structures and perform data visualization. RedisInsight also **automates common bulk actions** and helps analyze and reduce memory usage.



Open Source DNA



Even the most conservative enterprises are now using commercial solutions rooted in open source—and that includes databases.

Open Source DNA

It seems silly in 2021, but not long ago enterprise solutions based on open source software weren't seen as up to the task of running mission-critical enterprise applications. Those days are gone. Even the most conservative enterprises now rely on open source based solutions for a variety of use cases—as long as it's a good product with good support.

The most innovative and powerful technologies are coming from open source roots. Having open source DNA is essential: it's the only way to

ensure a technology is as extensible as possible and gets updated as quickly as possible. And roots in open source mean developers address what the community really needs, instead of what a vendor can do or wants to do.

Redis Enterprise is a DBaaS built to deliver the best Redis experience. It maintains the simplicity and high performance of open source Redis, while adding enterprise-grade capabilities.



About Redis

Data is the lifeline of every business, and Redis helps organizations reimagine how fast they can process, analyze, make predictions, and take action on the data they generate. Redis provides a competitive edge to any business by delivering **open source** and **enterprise-grade** data platforms to power applications that drive real-time experiences at any scale. Developers rely on Redis to build performance, scalability, reliability, and security into their applications.

Born in the cloud-native era, Redis uniquely enables users to unify data across multi-cloud, hybrid and global applications to maximize business potential. Learn how Redis can give you this edge at redis.com.

