

redfly.ai

Fully Automated Redis Caching for SQL Server

Most Enterprises have adopted Redis as the de facto standard for Distributed Caching. This is an ongoing trend as all databases, including NoSQL databases, face challenges with low latency read access and high scalability. The speed of data retrieval from disks has a threshold, and the database engine itself has become a bottleneck.

The fastest adoption rate for Redis is among RDBMS databases like SQL Server, Oracle & MySQL. The prominent use cases for Redis are User session store, high-speed transactions, content caching, Messaging, Real-time analytics, Notifications, Search/ secondary indexing, Fast data ingest, and Geospatial indexing. Caching significantly improves user engagement by improving application responsiveness. This is a critical factor for all successful applications – especially mobile.

The main factor inhibiting Redis adoption is manual cache management. Every team implements its own decentralized, custom caching solution, which, along with imperfect cache invalidation schemes, causes stale data and bugs. Redis adoption is hindered as each feature goes through implementation + QA and still has bugs because of imperfect cache invalidation schemes, which is the norm. This can be a significant loss of opportunity because of a buggy, unresponsive user experience, which is quite typical.

redfly provides schema-agnostic, end-to-end automation for cache management. All the major players like [Uber](#), Stack Overflow, Microsoft & Google have spent many years and millions of dollars with dedicated teams to implement custom Redis integration solutions that cannot be reused even within these companies.

- Our general-purpose, schema-agnostic design provides a single, universal solution for caching within any application.
- Redfly reduces cloud/ IT costs by 70% in real-world scenarios. This is consistent with the findings from [Uber CacheFront](#).
- Redfly APIs are 4,000x faster than SQL.
- Our SaaS services automate database synchronization to Redis.
- redfly requires no complex configuration, setup, or mapping.
- redfly APIs integrate caching, search indexing, and data access within the data access layer, enabling developers to focus on business logic rather than infrastructure plumbing.
- Our PolyLang compiler generates high-level C# code for the data access layer, which automates weeks/ months of complex, custom software development without the bugs caused by human error.
- Our platform is designed for resilience and extensively tested to work on dirt-cheap hardware, including micro-VMs and high-latency networks.
- Our AI-powered multi-threading framework accelerates multi-threaded workloads within all our products.
- Our current solution supports SQL Server, Redis, MongoDB, Azure Cognitive Search & Azure.