Enterprise Class Redis using a Database as a Service Delivery Model

© photo courtesy of Aaron Sylvan/LemonadeHeroes.com

Ofer Bengal
Co-founder & CEO
Redis Labs, Inc.

CEOCFO: *Mr. Bengal, what is the concept behind Redis Labs?*
Mr. Bengal: Redis Labs is the commercial company for Redis, which is an open source in-memory, NoSQL database. Most commercial companies behind open source products offer commercial support for the open source. We took a different approach. We provide enhancements for Redis, which make it more robust and much more suitable for enterprises. You can say that we provide enterprise-class Redis.

CEOCFO: *Have people been looking for a better solution?*
Mr. Bengal: Yes. The world is changing in terms of the amount of data and the expected performance from data. With the rise of modern Internet companies which deal with huge amounts of data - the traditional relational databases which served the world very nicely for the last few decades are being challenged as not being good enough for today's amounts of data and expected performance. Consequently, a new breed of databases which is called NoSQL ("not only SQL") has emerged. Among these you can find MongoDB, Cassandra, Redis and some others. The basic idea is to deal with very large amounts of data at a much better performance. Redis is outstanding and a little different because it is entirely served from RAM (Random Access Memory). This and the fact that Redis was built from the ground up for high performance, makes it extremely fast. By "performance" I am talking about the throughputs it can deliver and the average response time ("latency") under load, which are much, much better than any other database. This in part, is what made Redis extremely popular among developers, startup companies and today enterprises.

CEOCFO: *What was the challenge technologically? How did Redis Labs come about the concept?*
Mr. Bengal: Whenever you talk about modern databases there are a couple of things that need to be dealt with. First, there is scalability. You need to be able to scale out your database beyond the size of a single server and sometimes expand to very large clusters of servers. The other thing is high availability, or in other words; how to make sure your database is always on, even in cases of node failures or data center outages. And you need to take care of all of this while maintaining top performance. This is very difficult to achieve. In our case, where the entire database is served from RAM, if there is a node failure, you can lose your entire data. We had to build sophisticated technologies in order to cope with these challenges. In term of scalability, we can grow the size of the data set from a few megabytes to gigabytes to terabytes and even to petabytes, without any downtime and without any limitations on data types and commands of the database. At the same time we built a complete suite of high availability provisions to make sure you never, ever lose your data. You can replicate your data set in memory in the same data center, across data centers or across regions and even across different clouds. This is complemented by our instant auto-failover mechanism, which makes sure that if a node fails you always have a replica somewhere and the failover is instant and automatic. The database operator does not need to do anything. The replica is immediately connected to the application, and there is no downtime for users.

CEOCFO: *How are you approaching the enterprise market? What is the plan?*
Mr. Bengal: We started with a fairly new delivery model which is called Database-as-a-Service (DBaaS). About three years ago, we started to rent cloud servers on different public clouds, put our technology on them and offer customers to put their databases on these servers, or in other words, to outsource their entire database operation. This was considered a very advanced model, because customers basically put their data in our hands and must trust us in order to do so. I am happy to say that this became very successful. Today we have over 5,200 paying customers and approximately 25,000 free customers (we offer a free plan) on this service, which is called Redis Cloud. Later on we came up with a different product which is based on the same technology and is called Redis Labs Enterprise Cluster. This is enterprise software that anyone can download from our website and self-install on their own servers, whether they are located in a private

1

data center or a private cloud environment. The basic difference between this and the cloud service is that in this second model, the customer is operating the database. The third and most interesting delivery model has to do with customers which installed our enterprise software and asked us to remotely manage their databases within their environments. They give us access to these environments, whether on the cloud or in their private data centers and we manage their databases from remote. This again, is a very new and advanced way for operating databases.

CEOCFO: *I know you had a recent funding. How far will that take you?*
Mr. Bengal: We recently raised fifteen million dollars. I cannot disclose numbers, but our sales are climbing very nicely. We do not live on investors' money alone. The money we raised can take us a long way. It all depends on how aggressive we would like to be in terms of sales and marketing. It also depends on the global economic situation. If there is a market correction, which means investments become more difficult and with lower valuations, we would probably delay the next round of funding. Alternatively, if the market remains strong we may be more aggressive in trying to capture more market share, the current funds will take us a shorter way and we will need to raise additional funds.

CEOCFO: *How does the recent addition of Salvatore Sanfilippo, the Redis creator, enhance Redis Labs?*
Mr. Bengal: Salvatore Sanfilippo created open source Redis in 2009 and leads the open source community ever since. Redis started to become popular among developers around 2011 and really picked up in 2013. Salvatore contributes most of the code to open source Redis. Redis Labs was the second largest contributor. Salvatore has recently decided that the best way to make Redis a world leading database is to join forces with us. Therefore, now, after joining the company, we became responsible for managing the open source community and for contributing most of the open source code. That is a major milestone for the company.

> **"When it comes to performance there is only one database in the world that can respond in less than one millisecond under any load and this is Redis." - Ofer Bengal**

CEOCFO: *What surprised you as Redis Labs has grown and evolved?*
Mr. Bengal: When we started to deal with Redis in 2011, it was a big bet, because Redis started to show market traction only in 2013 and mainly with developers. We loved Redis from first sight and thought it was a great product. We said then "Hey, let us build a company around it." But there was no evidence at that time that Redis will pick up and become popular. We are amazed again and again to see how popular Redis became. A recent such surprise was a report by G2 Crowd, which review software products based on users surveys. They ran a survey on user satisfaction level and market presence of databases and Redis came out number one among the NoSQL databases and number four among all databases in the world! Another data point is a survey by devops.com and clusterHQ.com about database popularity in containers. Redis came out number one among the NoSQLs and number two among all databases, slightly behind MySQL, the very popular relational database. It always surprises me to see these evidences and it seems this momentum does not stop.

CEOCFO: *What is involved in switching to Redis? What is an enterprise looking at to make the change?*
Mr. Bengal: The optimal time to adopt Redis is when you start a new project, a new application, or a redesign of an existing application's architecture. If you wish to adapt Redis to an application which already runs on a relational database (like Oracle, MySQL or MS SQL Server) it tends to be somewhat more complex. A nice way around it is to use Redis as a cashing layer for accelerating performance of these databases. This is done in many cases and is almost transparent. You can, with very little effort, gain a lot in terms of performance.

CEOCFO: *Why is Redis Labs noteworthy?*
Mr. Bengal: Enterprises today look for performance. Performance is a major driver in business success. After Google made us all used to instant response no one can compromise for anything less. When it comes to performance there is only one database in the world that can respond in less than one millisecond under any load and this is Redis. Therefore, I would say that any company or any application which is performance aware should look at Redis and Redis Labs.

Interview conducted by: Lynn Fosse, Senior Editor, CEOCFO Magazine

For more information visit: www.redislabs.com