

## REPORT REPRINT

# RedisConf18: Redis rolls out updates on graph, Kubernetes and persistent memory

**JAMES CURTIS**

**12 JULY 2018**

RedisConf18 served as the platform for Redis to announce a bevy of product updates, including enhanced modules on graph and search, Kubernetes support, and increased memory capabilities based on Intel's 3D Crosspoint technology.

---

THIS REPORT, LICENSED TO REDIS, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2018 451 Research, LLC | [WWW.451RESEARCH.COM](http://WWW.451RESEARCH.COM)

At its annual user conference, Redis has always made key announcements regarding its enterprise version of Redis, the in-memory NoSQL database. True to form, the company continued its approach with several product updates at the recent RedisConf18. This year, however, the company's product announcements showed an additive approach to previous updates, proving the maturing and growing adoption of Redis. Product announcements included enhanced modules on graph and search, Kubernetes support, and increased memory capabilities based on Intel's 3D Crosspoint technology. Additionally, Redis reported healthy growth numbers for 2017, which followed the company's announcement of a series D round of \$44m in August 2017.

---

## THE 451 TAKE

Redis is showing good growth since our previous coverage, primarily in customer count and the types of customers it is securing. Investors are also showing confidence in the company, as evidenced by the recent series D round. While Redis may be challenged by the perception that it is used primarily for performance caching, the company's recent RedisConf18 announcements might prove otherwise as they enhance Redis' ability to function as a primary data store. The RedisConf18 announcements show the database company building upon existing functionality and maturing on a number of fronts, particularly with its graph and search modules, which are now included, making Redis a multi-model database out of the box. Integration with Intel's persistent memory, available in the future, should further strengthen the company's appeal.

---

## CONTEXT

Redis is the largest and most recognized commercial supporter of the open source Redis NoSQL database. The company was founded in 2011 by Ofer Bengal and Yiftach Shoolman, who now serve as CEO and CTO, respectively. In 2015, the company brought on Salvatore Sanfilippo to lead open source development efforts; he is also credited with creating the first version of the Redis NoSQL database.

Even though the company is not yet 10 years old, it is considered one of the leaders in the NoSQL space with its in-memory key-value-based database offering. The company's most recent announcements (discussed later in this report) also reflect this – many of its updates are further enhancements of existing features that Redis is driving to maturity.

While Redis is privately held, the company did publicly announce earlier this year its tenth straight quarter of double-digit growth (based on fiscal year ending 2017), including ARR increasing by more than 100%. The company provides both a cloud service and an enterprise version of Redis, and reports a collective total of 8,500 customers. Total employees come in at 200.

In August 2017, Redis secured a series D round of \$44m led by Goldman Sachs Private Capital Investing and existing investors Bain Capital Ventures and Carmel Ventures, with participation from Dell Technologies Capital. With this round, the company's total funding comes to \$86m.

## PRODUCTS

At RedisConf18, the company made a number of product announcements, and as noted, most were iterative versions of previous capabilities. For instance, during last year's event, Redis rolled out its active-active replication capability for geo-distributed workloads, leveraging the company's CRDT (conflict-free replicated data types) functionality. This year, Redis further that and announced its ability to handle up to 32 replications, as well as the ability to handle lists and sorted sets, essentially rounding out the primary data types supported in Redis. It should also be noted that CRDTs now support the company's casual consistency functionality, which translates to the ability to provide data consistency in globally distributed scenarios, should network or node failures occur.

Redis is enhancing two previously released modules: Redis Graph and Redisearch. With Redis Graph, the company is adding a new graph engine based on GraphBLAS, which leverages sparse matrix multiplication that stores the graph data efficiently in memory as 1s and 0s. To the user, the enhancement is relatively abstracted out, but the benefit is significant performance results, based on the company's internal benchmarks.

Regarding the company's Redisearch Module, this module is now able to handle aggregate queries. To the user, this capability should reduce previous manual efforts to carry out aggregations. This module now includes built-in capabilities for grouping, sorting and filtering, as well as the ability to do min, max, average, quantile and so forth. Redisearch and Redis Graph, including ReJSON, ReBloom and Reds-ML, all ship as part of Redis Enterprise.

Redis also announced Kubernetes support, and more specifically the ability to provide persistent data services inside the Kubernetes environment. Kubernetes is known for its ability to manage containers at scale, and uses what are called 'primitives' as architectural building blocks that can be extended to provide broader container management capabilities. Leveraging these primitives, Redis is able to provide persistent storage and high-availability functionality such that a Redis database can continue uninterrupted in the event of a node failure or a network outage.

Finally, Redis announced support for and compatibility with Intel's new persistent memory based on 3D Xpoint technology – technology co-developed by Intel and Micron Technology. The persistent memory is a new type of media that sits between disk memory (DRAM) and flash memory, and can provide significantly greater amounts of memory storage over RAM while providing significantly lower latency than flash storage.

While there is no definitive date on when this new persistent memory from Intel will be available, Redis is accepting some trial customers now. Overall, the announcement broadens Redis existing strategy of leveraging other types of memory, and is expected to provide the benefits of SSD (high volumes of data) while providing performance close to DRAM. As a result, Redis has updated its Redis on Flash (ROF) technology, which provides the ability to extend RAM with flash memory, to take advantage of Intel's new persistent memory. Currently, ROF uses a series of I/O threads to access the flash memory, which then leverages the RocksDB storage engine underneath, since it works well with flash memory. This approach requires a good deal of overhead, so Redis will provide an integrated storage engine that will be able to access the new Intel persistent memory directly (without the need for I/O threads).

## COMPETITION

As an in-memory database also categorized among the NoSQL databases, Redis has a number of competitors within the NoSQL market, as well as in the broader database market.

On the NoSQL front, Redis competes with Couchbase, which also gets leveraged as a caching layer, although the company has been actively positioning as an engagement database and pressing its ability to drive mobile applications. Aerospike is another key-value-based competitor that promotes a differentiated SSD-based architecture for systems of engagement. DataStax, which leverages Cassandra, also provides in-memory and graph capabilities, among other features. MongoDB is another that provides in-memory capabilities and some graph functionality, although MongoDB primarily promotes the document data model.

Other vendors include MarkLogic, FairCom and Oracle NoSQL Database. Given that Redis can be used as a caching layer, it is not entirely uncommon to see it being paired with other NoSQL databases in some circumstances, as well, in which case it could also be compared with in-memory data cache and database functionality from the likes of GridGain, GigaSpaces, Pivotal, Hazelcast and ScaleOut Software.

Cloud-only NoSQL databases are another competitive field for Redis, given the company's Redis Enterprise Cloud and Redis Cloud services. Specific competitors include AWS DynamoDB, AWS ElastiCache, Azure Cosmos DB and Google Cloud Spanner, for instance.

Other non-NoSQL competitors that constitute in-memory data platform players include SAP HANA, MemSQL and VoltDB. And on the caching side, it might see Alachisoft, GigaSpaces and Microsoft Azure Redis Cache.

Also worth noting, given that Redis is open source, is that there are other firms providing commercial Redis support and services, such as Pivotal, along with a number of companies providing Redis as a hosted cloud service, including RedisGreen, OpenRedis, Redis To Go and ScaleGrid. IBM has Compose for Redis, and ObjectRocket has its Managed Redis service.

## SWOT ANALYSIS

### STRENGTHS

Redis maintains a strong reputation for performance given its in-memory architecture and loyal developer community following.

### WEAKNESSES

Redis is often perceived as used primarily for caching or for specific performance-related scenarios and applications, even though data can be persisted within Redis and it is often used as a primary data store.

### OPPORTUNITIES

The company's database will continue to be top of mind for many organizations that require performance as a key trait. However, the expansion of replication (active-active), Kubernetes and graph gives Redis a number of avenues in which to expand within its existing customer base.

### THREATS

The NoSQL database space continues to be competitive. While many of Redis' rivals have (or are in the process of launching) other data models to deal with different data types and broaden use cases, perhaps the more pressing threat is the cloud vendors that are targeting globally distributed scenarios.