

REPORT REPRINT

This Impact Report was published by 451 Research as part of our syndicated Market Insight subscription service and subsequently licensed for commercial use by Pivotal.

In-memory specialist Pivotal unveils a Cloud Cache

JASON STAMPER

20 JUNE 2017

The company has launched an in-memory caching platform for cloud-native applications, and it's horizontally scalable. Pivotal was thought in some circles to have considered an IPO in prior years, but instead has raised a healthy \$1.7bn in funding across several rounds.

THIS REPORT, LICENSED EXCLUSIVELY TO PIVOTAL, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, 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.



©2017 451 Research, LLC | WWW.451RESEARCH.COM

Pivotal, which was spun off of EMC and VMware four years ago, was thought in some circles to have considered an IPO in prior years, but instead it has raised a healthy \$1.7bn in funding across several rounds. Now the company has added a Cloud Cache to its Cloud Foundry platform that can be deployed in a variety of clouds, including Amazon Web Services, Microsoft Azure, VMware vSphere and Google Cloud.

THE 451 TAKE

While Pivotal's new Cloud Cache can be used in a range of clouds to improve the performance of applications (and hence the end-user experience), it seems likely that the company will have the most success, at least initially, selling it to existing Cloud Foundry customers. There are one or two items on the roadmap that will make it even more compelling – such as cross-datacenter WAN replication – but most of the common requirements are already available: ease of configuration, a 'look aside' caching pattern out of the box, and high availability.

CONTEXT

Founded in San Francisco, Pivotal Software was formed in April 2013 after spinning off of EMC and VMware. It was created through the combination of Pivotal Labs, which had been acquired by EMC in 2012, with additional technologies from EMC and VMware. Since then, EMC has been purchased by Dell, but Dell has maintained Pivotal as a stand-alone brand.

Pivotal is very well funded. It considered an IPO in 2016, but in the end raised a fourth round of funding – a sizeable \$653m series C round led by Ford Motor (a customer), with additional investors including Microsoft, GE and Pivotal's two founding firms, VMware and EMC.

In terms of leadership, Paul Maritz, previously CEO of VMware, initially headed the company, but in August 2015 Maritz became chairman, and Rob Mee, the founder of Pivotal Labs, became CEO. Mee had been on the founding team of Pivotal Software alongside Maritz, Scott Yara (now president and head of products) and Bill Cook (now president and COO).

The company is relatively coy about revenue. But it has told us in the past that the Pivotal Data Suite, which includes the in-memory data fabric technology GemFire (powered by Apache Geode), as well as the Pivotal Greenplum data warehouse and its Hadoop-native SQL engine, Pivotal HDB, is a \$100m+ business. It has open-sourced the cores of all three products.

The other big revenue stream for Pivotal – alongside the Pivotal Data Suite – is Pivotal Cloud Foundry (PCF), which has in fact been its primary growth driver since its commercial release in 2013. PCF is Pivotal's PaaS that helps companies or cloud providers build public or private cloud-native applications. The company did disclose PCF-related revenue (or at least bookings) details in 2016: \$270m.

TECHNOLOGY

Pivotal recently took the wraps off of Cloud Cache, an in-memory caching platform for cloud-native applications. The firm explains that cloud-native applications receive a flood of concurrent requests from ever-larger numbers of users and applications. Cloud Cache makes use of some of the technology that was already available with its GemFire in-memory data grid. However, with Cloud Cache, that technology is optimized as a service to run in the cloud. GemFire, on the other hand, is prepackaged and not available as a service. Cloud Cache is available in the PCF platform, which supports AWS, Google Cloud, OpenStack, VMware vSphere and Microsoft Azure.

Pivotal notes that it is possible to scale out Cloud Cache horizontally – adding nodes for increased capacity and throughput. But that's all in one particular cloud (such as AWS). There are plans, however, to be able to do replication across data-centers for additional flexibility and resiliency – although the company says that it already offers strong high availability.

It also reports that Cloud Cache maintains ACID compliance where necessary – ACID being an important feature of transactional systems (it stands for atomicity, consistency, isolation and durability). Cloud Cache can 'inherit' durability from the database that it sits alongside. It is available through the Pivotal Network, where it can be accessed and installed as a service on Ops Manager. The service is then exposed to developers in a Marketplace. If added capacity is required, Cloud Cache can be scaled dynamically without any impact on client applications, and without any downtime.

Pivotal notes that Cloud Cache comes with several common caching patterns preconfigured. The first release was built for the ‘look aside’ caching pattern. With a look-aside cache, developers write code in their original applications, which manages the role of the cache. A more specific option related to http Session State Caching is targeted for release in late June or early July. The session state option allows applications to store state information externally in a cache for fast retrieval. This adheres to the 12-factor principles for building cloud-native apps that require apps to be stateless – i.e., state should be stored externally. Currently, it is not possible to replicate the cache from one datacenter to another, but WAN replication is on the roadmap.

Cloud Cache also enables operators to configure a range of service plans to the needs of their application developers, Pivotal explains. Capacity consumption and the associated costs can be controlled by setting instance quotas. Quotas can be set at two levels: the maximum number of instances, or the maximum number of servers, depending on the anticipated workloads. The company also emphasizes that Cloud Cache is ideal for microservice architectures, particularly when used in conjunction with PCF.

COMPETITION

The nearest competition to Pivotal’s GemFire in-memory data grid comes from Oracle Coherence, IBM WebSphere eXtreme Scale and Software AG’s Terracotta. There’s also Red Hat with its JBoss Data Grid, Hazelcast and ScaleOut Software. On the caching front, we’d expect Pivotal to vie with Alachisoft, as well as GigaSpaces with its XAP in-memory data fabric. In the cloud, AWS has ElastiCache. There is another open source project called Varnish, which is supported by Varnish Software if users buy into the Varnish Plus version. There is also some overlap between the data-caching and NoSQL database sectors since NoSQL can scale out horizontally – for example, the likes of Couchbase and Redis Labs claim that they can be used as replacements for Memcached, which is the open source cache employed mostly for Web applications. Redis itself can also be used as a cache.

SWOT ANALYSIS

STRENGTHS

Pivotal offers a cache for the cloud that has the potential to speed application performance and hence improve the end-user experience, particularly for latency-sensitive applications (of which there are more and more), as well as good out-of-the-box configuration options.

WEAKNESSES

Some will consider adding a cache in the cloud an additional layer of complexity. Meanwhile, cross-datacenter replication is on the roadmap but not quite here yet (although high availability within a datacenter is already supported).

OPPORTUNITIES

The lowest-hanging fruit must surely be Pivotal’s existing Cloud Foundry customers, which range across several cloud platforms. We see more and more latency-sensitive applications that could benefit from Cloud Cache.

THREATS

There is some risk that in the medium to long term, caching will be expected to be part and parcel of any cloud deployment, which will in turn effectively commoditize the caching sector.