

REPORT REPRINT

Macrometa emerges with planetary-scale database for apps and web services at the edge

**MATT ASLETT, JAMES CURTIS, CRAIG MATSUMOTO,
CSILLA ZSIGRI**

21 JAN 2019

The company plans to engage with CDNs, public cloud vendors, telecom operators and other service providers to make its technology available globally to support distributed edge data processing requirements.

THIS REPORT, LICENSED TO MACROMETA, 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.



©2019 451 Research, LLC | WWW.451RESEARCH.COM

Macrometa has emerged from stealth with a new distributed database platform and a plan to engage with CDNs and other datacenter and service providers to make its platform available across the globe to support distributed edge data processing requirements. The company's Global Edge Fabric is already available across 25 locations and is based on a combination of multi-model distributed database and general ledger capabilities, as well as data streaming, containers and functions-as-a-service, designed to deliver low-latency, hyper-local data processing at the edge.

THE 451 TAKE

Macrometa is by no means the first company to develop a globally distributed database, but it is differentiated by making the edge a focal point of its development, as well as by its approach of partnering with service providers to make the service available on a planetary scale. We have observed that it can take time for enterprises to realize the application requirements for globally distributed databases, which has slowed adoption to date. We anticipate that CDN providers would therefore have the potential to accelerate the adoption cycle for Macrometa, but much will depend on how many partners it can sign up to its Global Edge Fabric and how quickly it can get them offering services based on the platform.

CONTEXT

Macrometa was founded in late 2017 by CEO Chetan Venkatesh and chief architect and CTO Durga Gokina, who had previously worked together at hyperconverged infrastructure and software-defined storage specialist Atlantis Computing – where Venkatesh served as both CTO and CEO over the years, while Gokina was chief architect and director of engineering. They continued to collaborate following the acquisition of Atlantis' technology and customer assets by HivelO in mid-2017 – with Venkatesh serving as an advisor to Benhamou Global Ventures (BGV) and Lakewood & Co, and Gokina serving as director of engineering at Nutanix – then came back together in the fall of 2017 to form Macrometa, which has recently raised a fresh funding round (the value of which is currently undisclosed) led by BGV along with Partech Ventures, Shasta Ventures & Sway Ventures.

The new company has now emerged from stealth with its Global Edge Fabric, which combines in-house-developed multi-model distributed database and general ledger capabilities, along with data streaming, containers and functions-as-a-service (FaaS), designed to deliver a platform for low-latency hyper-local data processing at the edge. The company also has an ambitious strategy to work with CDNs, public cloud vendors, telecom operators and other datacenter and service providers to make its platform available across the globe – providing them with the ability to support dynamic data-driven applications in return. The company's Global Edge Fabric is currently available at 25 locations, with Macrometa aiming for more than 100 by the middle of next year. It currently has approximately a dozen customers on trial.

TECHNOLOGY

Macrometa is positioning its Global Edge Fabric offering as a geo-distributed data platform for global applications that is available in the form of a PaaS that can run atop public cloud infrastructure – AWS, Azure and Google Cloud Platform, for instance. However, the company is also looking to work with CDN providers. CDN systems typically are very good at replicating data across many nodes, but these systems typically do not maintain state. Macrometa's Global Edge Fabric is based on a new decentralized edge-native architecture built on a specialized streaming pub-sub engine that handles all of the data replication from the edge nodes. But a globally distributed database requires query and state management. So Macrometa has developed a multi-model NoSQL database that incorporates a document data model but can also handle other data models, such as key-value, time-series and graph. As such, Macrometa has effectively blended CDN capabilities for distributed replication with a NoSQL database for data persistence.

Architecturally, Macrometa employs a decentralized environment in which every edge node acts as a complete state repository, which means that every edge node can handle both reads and writes. It uses a distributed ledger-

type architecture, in order to record database state changes locally. Consistency is provided by what Macrometa calls a coordination-free convergence approach (as opposed to coordination or consensus), which broadcasts state changes at any particular edge node to the other edge nodes, which then accept and merge the data accordingly. More specifically, Macrometa does not leverage event time-stamps, but determines data ordering based on causality, retracing the cause-and-effect ordering of the events. Macrometa claims the resulting convergence is deterministic, adding a 3-5% delay over network latency. This is a somewhat different approach from that of emerging 'blockchain databases,' which are also decentralized but use consensus for consistency while providing immutability.

Because Macrometa provides a PaaS, which serves as a platform framework, the company enables enterprises to drive applications over this platform with API integration, which can then be run on public cloud infrastructure. Currently, Macrometa reports approximately 12 paying customers, running a variety of use cases. For instance, some customers are deploying a type of 'edge data lake,' where the service allows for processing at the edge to drive localized recommendation engines or to carry out targeted marketing. Another use case for Macrometa is compliance, such as for GDPR, since the system enables geo-fencing of the data, but the data can also be masked and moved to carry out other activities, including data analytics. Additionally, a large global telecom operator is building a telecom edge cloud using Macrometa as a PaaS for IoT, smart cities, blockchain and gaming use cases in anticipation of its imminent 5G rollout.

STRATEGY

The Global Edge Fabric is broadly applicable, so Macrometa is hoping to keep focused on immediately applicable markets – CDNs are one early target. Major CDNs are grooming edge-compute services: Akamai has EdgeWorkers and Cloudlets; Cloudflare offers Workers that run on JavaScript; and Fastly lets users run programs in Varnish Configuration Language (VCL) and plans to launch a FaaS sandbox this year. Macrometa could add full-fledged database capabilities and state awareness to these efforts, allowing database queries via APIs, for instance, from these edge compute services.

The company also sees an enterprise play around GDPR, because its platform executes fine-grained control over where data is stored. One medium-term target is likely to be the gaming market, which keeps user experience streamlined by avoiding functions such as database lookups. Macrometa could be useful for providing add-on features such as real-time leaderboards, the ability to match up opponents of the fly, and real-time data analytics to help drive in-game monetization in multiplayer freemium games.

COMPETITION

Given Macrometa's geo-distributed architecture, the majority of its competition will come from the NoSQL database providers, some of which are also specifically targeting global distributed database opportunities. But competition is also likely to come for peer CDN vendors as well as blockchain databases.

Because Macrometa positions its offering as a cloud service, Microsoft Azure Cosmos DB is a primary competitor. Cosmos DB supports multiple data models – document, key-value, wide-column and graph – but the core model is based on what Microsoft calls the ARS, or atom-record-sequence, data model. Other offerings from cloud vendors include AWS's DynamoDB (key-value, document), Google's Datastore (document), and the Google Cloud Spanner distributed relational database. There is also the Oracle NoSQL Database offered on its cloud platform, as well as the recently released Oracle Autonomous Transaction Processing (ATP) service. IBM has its Cloudant NoSQL database offered on IBM Cloud, while FaunaDB and Cockroach Labs are in the mix here as well.

Meanwhile, other NoSQL database vendors that support multiple data models include DataStax, Redis Labs, MongoDB, Couchbase, and MapR with MapR-DB. Also noteworthy is recent startup YugaByte, which specifically positions itself for globally distributed applications. Additionally, blockchain databases such as BigchainDB and FlureeDB combine features of a conventional distributed database and blockchain technology to allow multiple organizations to share a common set of data in a decentralized way.

CDNs such as Akamai, Cloudflare and Fastly have ambitious edge computing plans, but these efforts are early and resemble FaaS, in that they seem most appropriate for small, stateless applications. As such, Macrometa would seem a more likely partner than a competitor.

SWOT ANALYSIS

STRENGTHS

Macrometa offers a differentiated take on globally distributed database technology that is delivered as a more complete PaaS with additional general ledger capabilities, as well as data streaming, containers and functions-as-a-service

WEAKNESSES

The company's plan is ambitious, to say the least, and as an unknown quality at this stage, we expect many potential adopters to be skeptical about its ability to deliver on its promises until example use cases and reference customers emerge

OPPORTUNITIES

The plan to work with CDNs and other platform providers has the potential to accelerate time to market and enable enterprises to take advantage of the functionality without having to commit to large-scale internal deployment projects.

THREATS

The big cloud providers are likely to be the core competitors, and they boast mindshare that will be hard to shift. Again, the company's ability to leverage its emerging partner network will be key.