IMPACT REPORT

iguazio highlights continuous analytics use cases for converged data services platform

APRIL 11 2017
BY MATT ASLETT

Iguazio offers a combination of file and object storage, stream processing, and NoSQL databases with fine-grained security and distributed nonvolatile memory, as well as compatibility with Apache Hadoop, Apache Spark, Apache Kafka and AWS Kinesis – among other things. The company is nearing general availability of its Enterprise Data Cloud, with a growing focus on continuous analytics projects

The 451 Take

We continue to see opportunities for iguazio driven by its ability to reduce the complexity of deploying multiple data processing technologies. The company's focus on continuous analytics speaks to the ultimate value of deploying multiple data processing technologies, and is likely to attract a new set of potential customers. As before, we believe that iguazio will need to work hard to build credibility as an alternative to both traditional on-premises data processing and public cloud providers, and the company's ability to establish a list of reference customers is likely to be key to building momentum.

Context

We were first introduced to iguazio in the summer of 2016, shortly before the company (then known as Iguaz.io) dropped the dot. Iguazio takes its name from the Iguazu waterfalls on the Argentina-Brazil border, and was cofounded in 2014 by former executives of companies such as XtremIO, Mellanox and Voltaire with a mission to simplify the delivery of multiple approaches to data processing.

Iguazio's offering – known as the Enterprise Data Cloud – is designed to deliver multiple data processing services via stateless applications, along with a combined distributed in-memory data processing and storage engine. It is built on a combination of Docker-based microservices, the Kubernetes orchestration engine and the etcd distributed database (for cluster coordination and state management), and is designed to avoid the cost and complexity of deploying multiple clusters for Hadoop, NoSQL, streaming processing and object storage, for example.

Specifically, the iguazio Enterprise Data Cloud offers a unified data model that combines an in-memory distributed database with nonvolatile storage and cache capabilities to provide a real-time data catalog of data stored within the Enterprise Data Cloud, along with indexing, search, data optimization and data lifecycle capabilities, and support for ACID transactions. Above the unified data model is iguazio's integrated data firewall, which provides deep packet inspection and security policy enforcement capabilities.

Above that sits the V3IO messaging layer, which provides communication between the stateful storage and the stateless APIs, which are designed to be compatible with popular data storage and processing products and services, including Amazon S3, the Hadoop Distributed File System and the Linux file system for file and object storage; Amazon Kinesis, Apache Kafka and Apache Spark Streaming for stream processing; and Amazon DynamoDB, Spark DataFrames and time series for key-value-based data processing.

While Enterprise Data Cloud remains in beta testing, the company already has a number of paying customers and is seeing demand driven by the desire to deliver continuous analytics via a four-stage pipeline of data ingestion (real-time, rather than traditional batch ETL), data enrichment (aggregation and context), data analysis (using MapReduce, Spark and increasingly machine learning libraries such as TensorFlow) and data serving for visualization and insight.

The company cites specific interest in the financial services space for use cases such as intra-day risk analysis, portfolio analysis, market surveillance, IT operations analysis and numerous other real-time

decision-making projects – including those related to the Internet of Things. General availability is now expected in the second half of 2017, which is a slight delay from what was expected to be August 2016, due to (according to the company) the need to respond to growing demand from the company's beta testers. Iguazio now has 60 employees, up from 50 in August 2016, the majority of which are based in Herzliya, Israel, although the company is growing its US-based sales presence.

Iguazio raised a $15m series A funding round in late 2015 led by Magma Venture Partners, with Jerusalem Venture Partners and unnamed strategic investors. Equinix is a key partner that supports hosted deployments of Enterprise Data Cloud via Equinix Data Hub.

Competition

We still have yet to see anything that compares directly to iguazio's combination of data processing, storage and cloud architecture, and it would appear that the greatest competition for Enterprise Data Cloud will continue to come from enterprises assuming that they need to deploy multiple clusters dedicated to Hadoop/Spark, NoSQL and stream processing (for example). The closest competition is likely to come from MapR with its Converged Data Platform, while Cloudera, Hortonworks and IBM also offer competing functionality, perhaps without the same focus on convergence.

Iguazio also notes that MemSQL is increasingly talking about tackling the same problems with its ability to support transactional, analytic and streaming workloads, as well as its new cloud version. Iguazio also sees multiple potential customers considering a combination of cloud services. Amazon Web Services is the most popular choice, but Microsoft and Google also offer a variety of services that could be considered as alternative options to iguazio's Enterprise Data Cloud. Iguazio offers compatibility with Amazon's S3, DynamoDB and Kinesis, as noted above, while Enterprise Data Cloud is also available on AWS, so it isn't a simple competitor relationship, although iguazio is proving attractive to service providers to help them compete with AWS.

SWOT Analysis

Strengths

The company has a differentiated product and a seasoned management team with experience in the storage, data processing and networking requirements of cloud providers and large enterprises.

Weaknesses

The company is something of an unknown quantity, and will need to establish a good list of early adopter reference customers if it is to be considered a serious contender for mission-critical workloads.

Opportunities

We see wisdom in the company's positioning to reduce the complexity of supporting multiple data processing and storage engines.

Threats

The incumbent data platform providers hold the cards in terms of established customer relationships, while cloud service providers are the obvious default alternative.

Matt Aslett

Research Director

M&A ACTIVITY BY SECTOR
M&A ACTIVITY BY ACQUIRER
FIGURES SHOWN INDICATE NUMBER OF TRANSACTIONS

COMPANY MENTIONS (PRIMARY)

iguazio

COMPANY MENTIONS (OTHER)

Amazon Amazon Web Services Cloudera Docker Equinix Google Hortonworks IBM Jerusalem Venture Partners Magma Venture Partners MapR Technologies Mellanox Technologies MemSQL Microsoft TensorFlow Voltaire XtremIO

CHANNELS

Data Platforms & Analytics

SECTORS

All / Information management / Data management / Non-relational databases

All / Information management / Data management / Hadoop

All / Information management / Data management / Relational databases