

Using Open Source Apache Cassandra[®] and Instaclustr to ensure flexibility, scalability, and cost control

Overview

Diginus is a SaaS technology company based in London that specializes in digital marketing and e-commerce solutions. Its solutions aim to improve customers' online activities and provide them with more leads and higher conversion rates. To do this, Diginus' solutions and technologies provide customers with complete insight into all their online activities. These insights are delivered across multiple media channels—all within a single platform.

Protecting customer data is a top priority for Instaclustr. In a time of increasing data volumes, the company therefore requires a high-performance database and data layer technology. For flexibility and scalability, and for cost reasons, Diginus relies on an open source solution, namely Apache Cassandra[®].

Industry:

Digital Marketing and E-Commerce

Instaclustr Solution:

Managed Apache Cassandra[®]

Data Storage:

3 nodes with a storage capacity of 250 GiB

Use Case:

Storage of customer data

Website:

<https://diginus.com/>



In terms of reliability and provider support, we are extremely satisfied with Instaclustr Managed Apache Cassandra. This is confirmed by all employees involved with Cassandra and communication with Instaclustr. Relying on Instaclustr was clearly the right decision and we are now well prepared for the future in terms of database management. Most importantly, Instaclustr's managed platform offering and highly scalable NoSQL data storage will help us grow our business while relieving our own development and operations staff from database management tasks.

Nate Burke
Founder and CEO
of Diginus



The Initial Situation

In the past, Diginus used the services of a hoster for its Cassandra instances. However, there was hardly any support and some problems with failing nodes. In addition, the scaling requirements could not be met. The company's goal was therefore to introduce an easily scalable Cassandra cluster with provider support that could meet its own data storage requirements now and in the future. It was important for Diginus to have comprehensive support for the management and operation of the Apache Cassandra environment, including provisioning, monitoring, and backup, as well as proactive resolution of known and unknown issues that could impact application performance.

Partnering With Instacluster

After an extensive research and evaluation phase, Diginus decided to use Managed Apache Cassandra from IT service provider Instacluster. Diginus also looked at other offerings, but they were out of the question for cost reasons. The decisive factor in the decision was Instacluster's offering and expertise. For example, Instacluster already worked with customers on a large scale in the Cassandra area, also operating mission-critical clusters. Last but not least, the Instacluster solution had the necessary scalability and data storage capabilities to meet Diginus' requirements.

"We chose Instacluster for administrative, functional and also cost reasons. Instacluster was able to provide us with an out-of-the-box solution and made us a very attractive offer in terms of price. In addition, the support offer and the possible migration without downtime were very convincing arguments," explains Nate Burke, founder and CEO of Diginus.

The Implementation

Once the decision was made to use Instacluster, Diginus immediately started the migration without first launching a pilot project. After an initial consulting phase that included a review of the existing cluster, Instacluster recommended changing the structure of some tables to prevent partitions from growing arbitrarily large over time. This required some work on the code and reporting, which took place over a period of about a month.

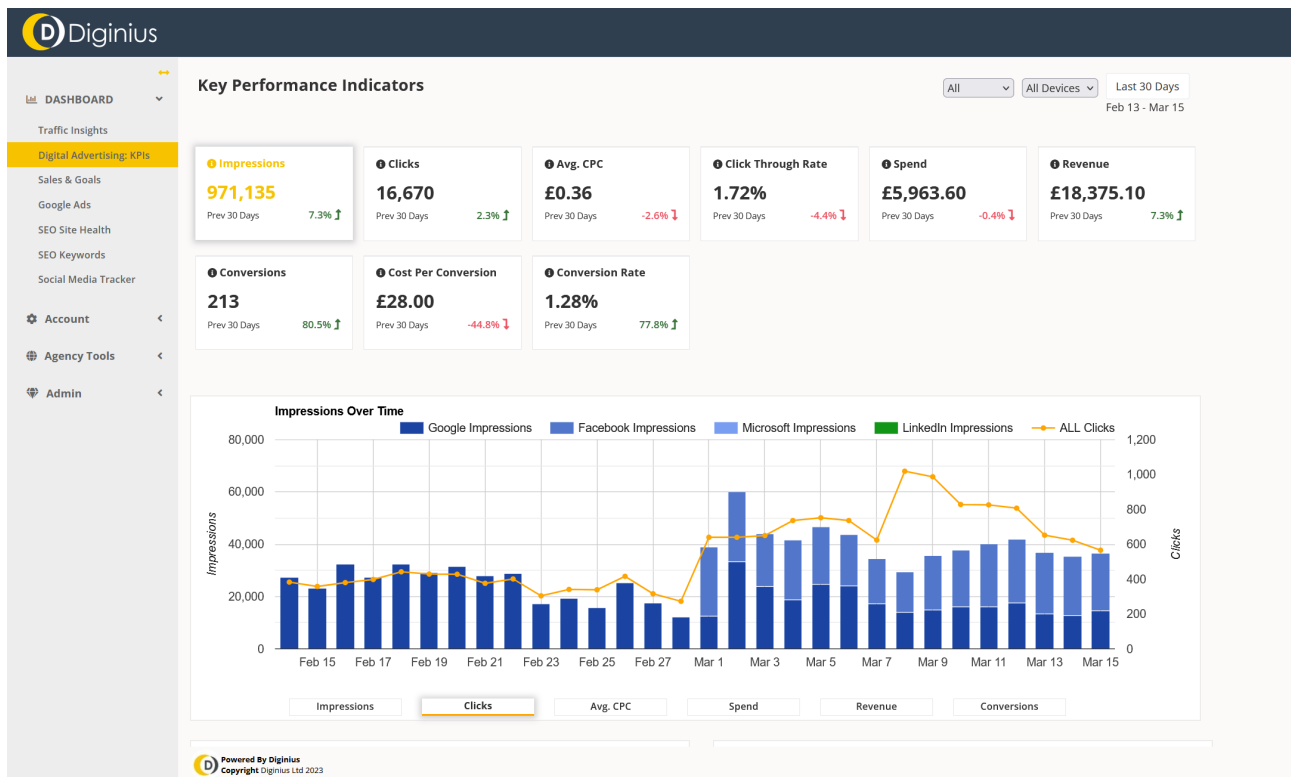
After that, Diginus prepared the database connections for the transition to the new cluster. Instacluster then performed the actual migration, which took only a few hours, and temporarily synchronized the 2 clusters. Diginus then moved all the connections to the new cluster, which was completed within a day. The connections were maintained for about a week to ensure that everything was working correctly. After this verification was successful, Instacluster removed the connection to the old cluster, which was then decommissioned. The migration project was largely carried out by the technical team at Diginus in close coordination with Instacluster.

The Solution

Numerous companies already rely on Apache Cassandra, the world's leading NoSQL database, to manage large volumes of data. As a universal transactional database, Cassandra's strengths lie in scalability and availability. Thus, Cassandra is a good option for companies that require 99.9% or higher availability. In addition, Cassandra is ideal for companies that want to flexibly expand their services as load increases, as scaling to multiple nodes can be done in minutes.

A proven operating environment with 24x7 support is especially suited for effective and secure deployment of Cassandra. Instaclustr for Apache Cassandra offers a powerful combination of automated functions within the platform, expert support around the clock, and processes that cover even complex customer requirements. Instaclustr takes over the management of Cassandra, allowing customers to focus on innovations in their other application stack. Diginus therefore also uses Apache Cassandra in a managed platform model from Instaclustr, namely on Amazon Web Services in a 3-node environment with a storage capacity of 250 GiB.

Instaclustr's managed platform for Apache Cassandra is used by Diginus for storing customer data, namely all data that the company collects from external APIs and uses for reporting, bidding, or algorithmic calculations. Diginus' iHub reporting software is also connected to the Cassandra databases.



Diginus Insight software provides a consolidated view of customer data in an easy-to-use interface.”(Source: Diginus)

The Advantages

According to Diginus, the Instaclustr solution is stable and runs completely smoothly. The company particularly emphasizes Instaclustr's comprehensive and continuous support in the operation of Cassandra. Instaclustr also warns Diginus of potential problems and recommends actions to prevent degradation of cluster availability and performance. Thanks to the permanent system monitoring by Instaclustr, Diginus only has to deal with a small amount of administration. It also ensures that timely action is taken in the event of problems or disruptions.

Diginus intends to continue to dynamically drive the company's growth. In this context, the requirements for data storage will also continue to increase. With Apache Cassandra and the support of Instaclustr, Diginus feels it is in an excellent position to meet this challenge. This also makes it easy to add additional nodes or scale existing nodes.

Key Advantages

- High scalability, availability, and performance with low latency
- No vendor lock-in thanks to 100% open source
- Comprehensive provider support
- Automatic health checks
- Zero downtime migration

About Instaclustr

Instaclustr helps organizations deliver applications at scale through its managed platform for open source technologies such as [Apache Cassandra®](#), [Apache Kafka®](#), [Apache Spark™](#), [Redis™](#), [OpenSearch®](#), [Cadence®](#), and [PostgreSQL®](#).

Instaclustr combines a complete data infrastructure environment with hands-on technology expertise to ensure ongoing performance and optimization. By removing the infrastructure complexity, we enable companies to focus internal development and operational resources on building cutting edge customer-facing applications at lower cost. Instaclustr customers include some of the largest and most innovative Fortune 500 companies.

© 2023 Instaclustr Copyright | Apache®, Apache Cassandra®, Apache Kafka®, Apache Spark™, and Apache ZooKeeper™ are trademarks of The Apache Software Foundation. Elasticsearch™ and Kibana™ are trademarks for Elasticsearch BV. Kubernetes® is a registered trademark of the Linux Foundation. OpenSearch is a registered trademark of Amazon Web Services. Postgres®, PostgreSQL® and the Sionik Logo are trademarks or registered trademarks of the PostgreSQL Community Association of Canada, and used with their permission. Redis™ is a trademark of Redis Labs Ltd. *Any rights therein are reserved to Redis Labs Ltd. Cadence is a trademark of Uber Technologies, Inc. Any use by Instaclustr Pty Limited is for referential purposes only and does not indicate any sponsorship, endorsement or affiliation between Redis and Instaclustr Pty Limited. All product and service names used in this website are for identification purposes only and do not imply endorsement.