severalnines

Data Scalability – eCommerce's Achilles Heel

Customer Case Study: Africa Internet Accelerator

March 2014

1. Africa Internet Accelerator (AIA)

AIA builds and operates South Africa's winning e-Commerce companies.

AIA's hands-on approach coupled with the collective experience of the passionate team behind it make for best practice implementation across all key business areas: Marketing, Supply Chain Management, Sourcing, UX/UI, CRM, Business Intelligence, Finance, Engineering and product management.

AIA's individual team members have built some of the fastest growing ecommerce companies in the world: lamoda.ru (Russia), zalora.com & lazada.com (SEA), theiconic.com (Australia), wimdu.com (Germany & China), zando.co.za (South Africa), jumia.com (Nigeria & Morocco), linio.com (LATAM), jabong.com (India), Groupon.co.za (Groupon South Africa).

AIA At a Glance

Industry: eCommerce

Location: Cape Town, South Africa

Data Center: Private Data Center in Germany

Use Case: Scalability of eCommerce infrastructure

Why Severalnines: Real-time performance insight, high availability operations and competency in database clustering

2. Challenge

The main challenge AIA was trying to solve was rapid growth, while trying to do everything in-house. The company had moved from Master-Slave MySQL setups to Galera Cluster in order to scale. The team manually set up Galera, and it was running nicely until they hit some incidents. These made the ops team realise the need for a management tool.

One incident the team ran into was a Split Brain scenario during a weekend. They had lost a MySQL node after a configuration change that was driven by a complex Puppet setup. The cluster was started with 1 node having an empty `wsrep_cluster_address = gcomm://` and a split brain ensued.

The impact to the business when this Split Brain occurred was loss of resources, loss of revenue and delayed business processes. The database cluster was a bit like a black box, and the team really lacked insight into what was happening under the hood. Also, since the database cluster was a mission-critical resource, the team also needed tools to assist them properly manage the cluster. It was important to find a vendor with good credentials in database clustering.

> "There were too many animals in the zoo. We had a whole eco-system of applications, all contending for database resources. We were outgrowing our Master-Slave database infrastructure and badly needed to scale."
>
> - Riaan Nolan, Senior Technology Manager who has worked on a number of large-scale eCommerce properties

3. The Solution

The ops team investigated possible alternatives for the database, and even looked at companies like Amazon and RightScale to help them guarantee consistency and uptime of the database layer. During the evaluation, the main criteria were cost, stability and functionality. Rapid deployment was also important as it was key to agility, and the existing infrastructure was already well automated with Puppet.

> "In a startup environment, budget has "almost" the last say so cost was a major factor. Stability, rapid deployment and scaling were other factors."
> - Riaan Nolan

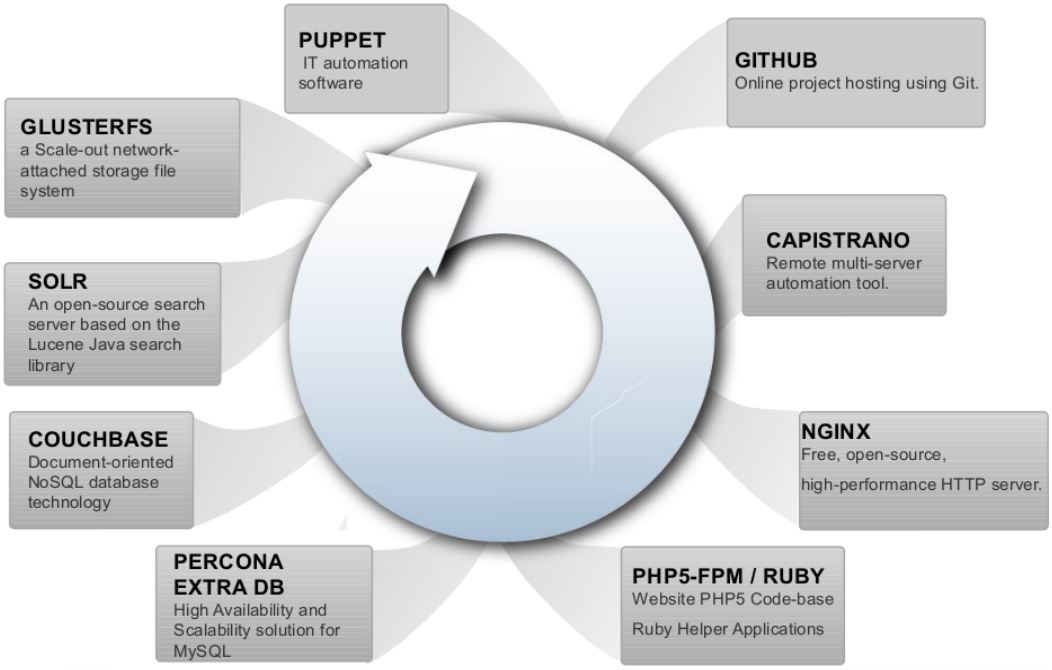

After researching different database technologies, the team concluded that Galera was still the right solution. It offered a multi-master setup, and architecturally, it was a good fit with the rest of the infrastructure. The missing piece was management and operational insight. Puppet could still be used to manage the database hosts, while automation of configuration changes, node recovery and updates to Galera would now be done by ClusterControl.

> "We already knew that we liked Galera Clusters, but we needed a way to simplify the management and operational aspects. We also needed a tool that could provide us a deep level of insight into runtime operations and performance."
> - Riaan Nolan

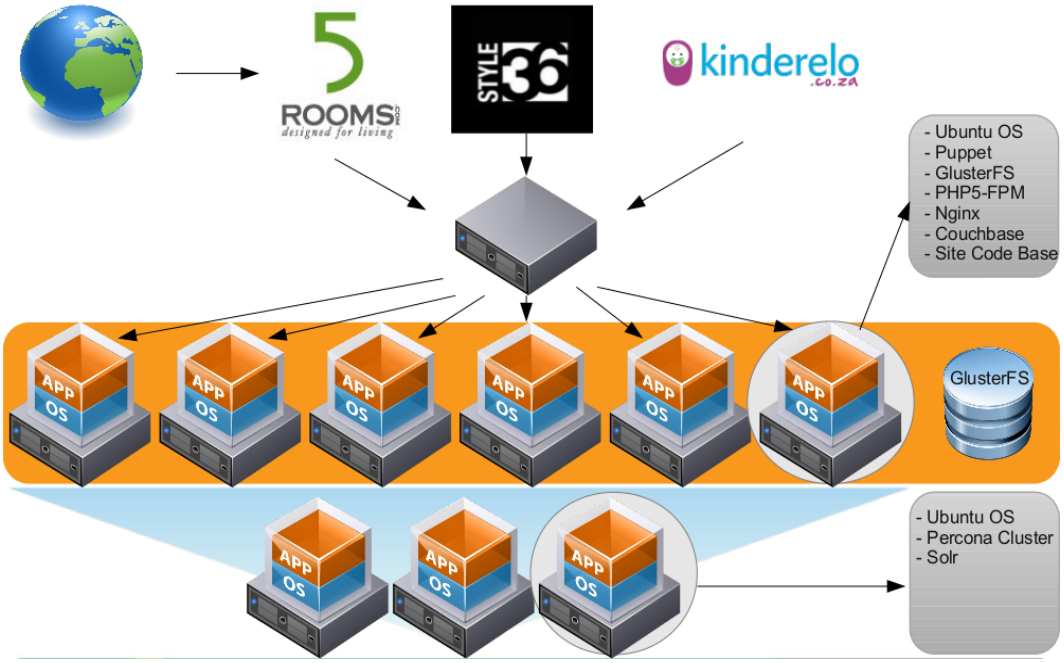

Having insight into the cluster also helped tremendously in managing the platform and increasing performance. The ops team is now able to manage bad queries in real time. With the Health Reports, the team is also able to proactively work on their database schema and queries. As the application is constantly evolving, with new code potentially causing malfunctions, the team is able to react to in a more targeted way than before.

Before

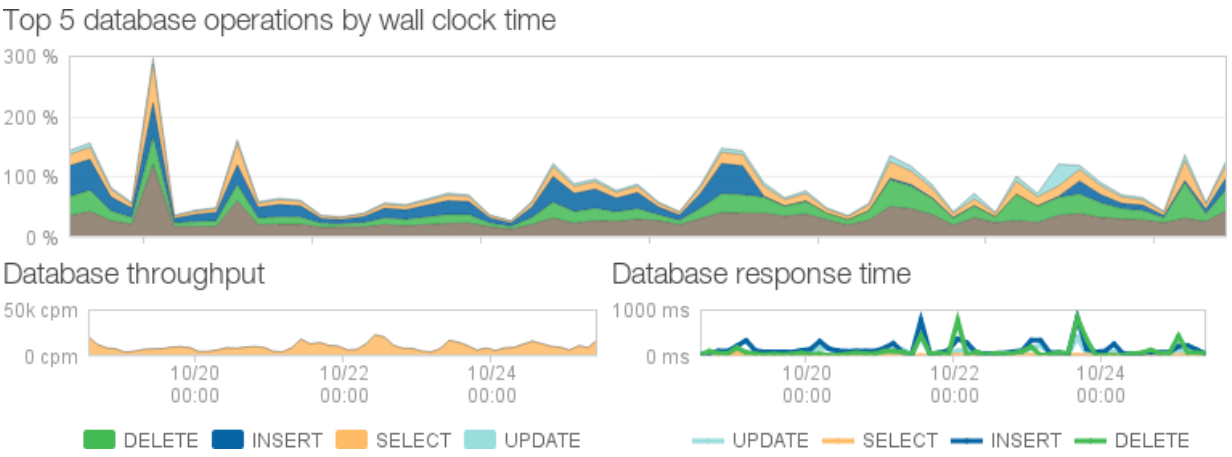

After