

severalnines

spree

DRIVING ECOMMERCE

GROWTH FOR SOUTH

AFRICA'S ONLINE FASHION LEADER

SPREE

- Industry: eCommerce
- Location: South Africa
- Data Center: Private

USE CASE

High Availability and Scalability for Magento eCommerce platform

WHY SEVERALNINES

Easy deployment, management and scaling of database cluster environment

INTRODUCTION

In little over 18 months, Spree.co.za has become South Africa's number 1 online fashion destination. The website offers curated fashion choices for Women, Men and Children, lately adding further categories namely Beauty, Home, Décor and Gifting. Stocking both local and international designers, Spree is making a name for itself in the fashion industry as a savvy and forward thinking brand.

Spree was launched in April 2013 and has a staff of just over 100 employees. The company falls within Naspers, a global platform operator with principal operations in internet services, especially ecommerce (i.e. classifieds, online retail, marketplaces, online comparison shopping, payments and online services); pay television (directto-home satellite services, digital terrestrial television services and online services) and print media. The group provides services in 130 countries.

CHALLENGE

Spree had launched their Magento-based eCommerce platform running on an active-

“Our single Master DB could no longer handle the number of connections and queries we needed to serve during peak periods. This obviously resulted in us dropping traffic and losing revenue. With traffic expected to grow, we had to look for a solid solution.”

Trent Goulden, System Administrator

passive MySQL database setup. The main challenge was the rapid growth that the company experienced, driven by massive local media exposure within the first months of launch. The ops team quickly found out that the main master database could no longer handle the number of connections they needed to serve during peak periods, this resulted in dropped traffic and revenue loss. Even though online users could browse the website, orders could not be placed if the database was unavailable.

For an online shop, the website is the cash register and it has to be open for business 24 hours a day. So high

availability was a key requirement, under no circumstances should the system go down.

Since Spree did not have a dedicated DBA, they needed a cost-effective solution that was easy to manage. The traffic pattern for an online retailer can vary greatly, with peaks during the holiday rush. Beefing up for peak traffic requires scaling up systems to allow customers to seamlessly transact on the site. Therefore, scaling the database needs cannot be a tedious procedure that takes time away from other activities.

SOLUTION

The number of large transactional queries that the team was seeing pointed them towards a highly available, multi-master setup. NoSQL databases were ruled out, the database was highly transactional and Magento was based on MySQL. The team briefly looked at Continuent Tungsten, but were not comfortable with a Java-based solution.

The team spent 2 months evaluating Galera Cluster. The initial deployment was quickly done with the ClusterControl. The cluster was subjected to a battery of availability tests. Some of the important characteristics tested included synchronous multi-master reads and writes, and rolling restarts with zero downtime. ClusterControl provided a wealth of insight and visibility into the cluster, which helped the team understand what was going on.

“The solution was incredibly fast to deploy, the ClusterControl tool was simple to use and we had our basic server setup fully operational within 40 minutes of completing the configuration”, says Trent Goulden. “The ClusterControl health checks and process monitoring make it very hard to actually kill the cluster. In our evaluation, the solution was able to automatically recovery from every simulated failure we could throw at it – with zero data loss.”

The team went live after the evaluation. The actual migration was straightforward, and required no schema changes. At the time of writing (October 2014), the system has been running live for over a year without issues, which has brought the peace-of-mind that Spree now has a data layer that is highly available. No traffic is dropped due to a database limitation, even when subjected to extreme traffic conditions.

“Some of the benefits we’ve seen is powerful DB monitoring at a glance. The Health Reports have really aided us in identifying simple fixes which have shown fairly significant performance gains”, says Trent Goulden. “We feel more confident about our DB layer these days – things almost take care of themselves. We spend zero time sorting out DB uptime issues anymore, and this has allowed us to really focus our efforts on other areas of the business. Severalnines will remain the bedrock of our technology stack for the years to come.”

Spree’s technology stack continues to evolve, but at the time of writing, they used a mix of Nginx, Varnish, Node.js, PHP5, ElasticSearch and Percona XtraDB Cluster (Galera).

“The solution was **incredibly fast to deploy**, the ClusterControl tool was **simple** to use and we had our basic server setup **fully operational** within **40 minutes** of completing the configuration”

Trent Goulden, System Administrator



The image shows a screenshot of the ClusterControl web interface. At the top, there's a navigation bar with various menu items. The main content area features a 'Cluster Load' graph with a blue line representing the load over time. Below the graph, there are four circular icons representing different functions: 'Deploy' (a blue circle with a white arrow pointing up and down), 'Monitor' (a blue circle with a white heart rate line), 'Manage' (a blue circle with a white server rack icon), and 'Scale' (a blue circle with a white plus sign and a server rack icon). The text 'ClusterControl' is prominently displayed in the center, with the tagline 'The Single Console for Your Entire Database Infrastructure' underneath it.