



Scaling databases for over 200 million active online users



AVG

Industry: Security
Technologies: Galera, MySQL
Datacenters: Private
Products: ClusterControl

USE CASE

Operating a database system for 200 million active online users

WHY SEVERALNINES

Fast deployment and simplified operations of database clusters

BACKGROUND

AVG Technologies (NYSE: AVG), founded in 1991, offers anti-virus and security software products and online services for home and business users.

AVG has grown its user base to 200 million active users and offers a product portfolio that includes Internet security, PC performance optimization, online backup, mobile security, identity protection, and family safety software.

Headquartered in Amsterdam, AVG is the fourth-largest vendor of anti-virus software and employs close to 500 people worldwide with corporate offices in the US, the UK, The Netherlands, the Czech Republic, and Germany.

CHALLENGE

AVG has a number of online applications for customers, reseller partners, and employees. Applications include support portals, user forums, and paid online services like LiveKive, an online backup service that allows end users to backup all their devices. Applications also include partner portals like Reseller Centre. Reseller Centre allows resellers to access all customer-related actions, so partners can easily transact around license delivery, and more efficiently manage customer orders.

With a user base of over 200 million active users and growing, the company identified some key architectural changes in the infrastructure so it could continue to scale. One of these included centralizing the user information, which included login information, roles, and profiles. The user repository would have an important role in the overall online services architecture since it would authenticate and authorize users every time they logged in. With the number of applications also increasing, this asset would also enable one view of the user from all applications.

The underlying database technology for the user repository needed to be very secure, scalable, and extremely fault-tolerant. Any downtime would have serious consequences since it would directly impact all online services.

We had previously implemented MySQL-based systems using a mixture of MySQL Replication and homegrown functionality for failover and load balancing. However, these were costly to implement, and ongoing operational management was manual and tedious.

Václav Adamec, Senior Systems Engineer for AVG

A quick search for high availability MySQL revealed a number of solutions from different vendors. The ideal solution had to scale in an easy manner since AVG was continuously adding more users and applications to the system. It also had to be self-healing, and automatically recover from failures.

Finally, it had to be easily manageable from an operational standpoint. Any performance issues experienced by applications would have to be investigated and resolved very fast. To keep up with patches and security fixes, upgrades should be automated and performed without any downtime.

SOLUTION

A few solutions were evaluated, including MySQL Cluster, MySQL master-master replication with multiple slave instances, and Galera Cluster with HAProxy for load balancing.

Galera Cluster was chosen for its ability to do master-master replication for InnoDB tables. Using Severalnines tools, the whole evaluation took under a month.

OUTCOME

Deployment of both Galera Cluster and HAProxy was completely automated and done in a matter of minutes. Once deployed, it was easy to see what was going on in the cluster using the ClusterControl web interface. It showed all the database nodes as well as the load balancers. The useful operational tasks that the team was looking for were available out-of-the-box in ClusterControl, including adding more nodes, scheduling of cluster backups, software upgrades, and rolling restarts. The new solution could scale to 20x more capacity as compared to a reference system using MySQL Replication.

We were very impressed because the whole deployment and management of the cluster was made so easy by the Severalnines tools. There were a number of useful hints in the documentation, which we took advantage of in the production environment. So far, it looks like a small miracle and we are now leveraging this solution in other parts of our infrastructure.

Václav Adamec, Senior Systems Engineer for AVG