



severalnines

PUTTING THE BRAKES ON DATABASE DOWNTIME

KNORR-BREMSE

- Industry: Manufacturing
- Location: Austria
- Data Center: Unknown

USE CASE

High availability databases used during production

WHY SEVERALNINES

Simplified operations of DB clusters

INTRODUCTION

The Knorr-Bremse Group is the world's leading manufacturer of braking systems for rail and commercial vehicles. For more than 110 years now, the company has pioneered the development, production and marketing of state-of-the-art braking systems. Other lines of business include intelligent entrance systems, HVAC systems, power conversion systems and driver assistance systems.

Knorr-Bremse is headquartered in Munich, Germany. The group has a global workforce of 24,000 employees in 27 countries, with a revenue of 5.2 billion euros in 2014.

CHALLENGE

Knorr-Bremse's locations in Austria have a number of internally-developed applications that are used during production of their products. Work is organized in three shifts per day to ensure that manufacturing processes run 24 hours a day, without any interruption. Production processes rely on continuous availability of data, and any disruption in the databases would impact the output of the plants.

“Our production processes relied on **continuous availability of data**, and any database downtime would **directly impact** the output of our production plants. Our MySQL replication setup was **not very good** at handling planned or **unplanned downtime** of the individual servers.”

Juergen Mayer, Manager IT District Austria

Strong data consistency is another imperative, as these processes need to have data that is realtime and accurate – even if the data is distributed between servers.

The Database Situation

Knorr-Bremse's MySQL databases made use of master-slave replication to achieve redundancy, but that was not satisfactory. The operating system and other software components on the database servers were updated on average twice a month with updates and security patches. It was important to do so without database downtime. Failover of the regular master-slave setups was manual. Since replication was asynchronous, there was risk for data loss.

The company already had a virtualized infrastructure in place, with a mixture of Ubuntu 12 and 14 running on VMWare hosts. Any new solution had to fit the existing environment.

Finally, the systems were operated by a team of system administrators (sysadmins). There was no fulltime database administrator (DBA) specialized in managing and maintaining the database systems. So it was important to have a system that could be managed by the existing team, rather than having to recruit a DBA.

SOLUTION

Galera Cluster with HAProxy looked like a very good alternative to the standard asynchronous MySQL replication. It automatically managed failures of single nodes, and would resync them when they came back online. At least on paper!

The team decided to evaluate the technology. A test cluster was quickly deployed using the online Severalnines cluster configurator. The existing database was easily migrated to the Galera cluster, as the tables were mostly InnoDB. The process

to evaluate, test and then go live took about 3 months. As the team was getting up to speed, there were a number of questions that came up. They leveraged the Severalnines support team to get these resolved, so that the team could quickly get to a fully working solution.

The solution went live in December 2014, and there has not been any downtime incident since. The sysadmin team is also able take out individual servers in order to do their regular maintenance work without affecting applications that need access to the data.

Since replication is synchronous, all database servers have exactly the same data. So all applications are able to access any database server via a redundant HAProxy load balancer, with the certainty that the data is consistent across the whole cluster.

Finally, the management is automated using Severalnines ClusterControl - from scheduling backups to removing nodes out of the cluster for maintenance, and then re-introducing them back in production. Having a visual representation of the Galera cluster, from database load, SQL queries and host metrics, allows the team to get a good view into what is happening.

“We are **very glad** to have moved to Galera Cluster with ClusterControl, as we now have a **highly available and stable** database solution for our applications. This would not have been possible without the **tools from Severalnines**, which helped us get productive in a very short time. The support team was also **fast and competent**, so we could quickly resolve any issues that arose.”

Juergen Mayer, Manager IT District Austria



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 monitor icon), '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.