

58 MILLION MERCADOLIBRE CONSUMERS GENERATE PLENTY OF DATA; GALERA CLUSTER MANAGES IT

ABOUT

MercadoLibre is the on-line trading platform in Latin America in which 58 million consumers in Argentina, Brazil, Chile, Colombia, Ecuador, Mexico, Peru, Uruguay, and Venezuela use to buy and sell a variety of consumer items and much more.

It's the largest e-commerce ecosystem in Latin America. Mercado, the company behind MercadoLibre, offers a wide range of services to its sellers and buyers throughout the region including marketplace, payments, advertising and e-building solutions; it operates in thirteen countries across Latin America. Based on unique visitors and page views, it is a market leader in the major countries where it operates.

With a market of over 550 million users and in a region with one of the world's fastest-growing Internet penetration rates, it provides buyers and sellers a robust on-line trading environment that fosters the development of a large and growing e-commerce community.

Mercado offers a complete technological and commercial solution that addresses the distinctive cultural and geographic challenges of operating an on-line trading platform in Latin America.

Among the world's e-commerce market leaders, MercadoLibre was recognized in 2010 as one of the top 30 world's hottest brands.

QUICK FACTS

Customer name: Mercado

Industry: Retail

Size: 1,000 to 5,000 employees

Consumers: 58,000,000

Headquarters: Buenos Aires, Argentina

More info: <http://investor.mercadolibre.com/>



BACKGROUND

With its success and accompanying growth, MercadoLibre's reliance on technology accelerated and its data center was filling rapidly with physical servers as well as with data. To speed things up, MercadoLibre's IT team began to virtualize much of their data center.

Virtualization saved them time in managing physical servers, reduced the data center space and electrical power that was needed, and enabled them to create more quickly new servers.

However, by moving to cloud computing, MercadoLibre's IT team knew there had to be more efficiencies and an improved end-user experience to be gained. They wanted their internal users to be able to create their own servers, virtual machines, storage, and networks.

And with the right cloud computing platform, the IT team knew they could give their users that ability. Users would be able to instantly provision themselves with what they needed without having to wait for a system administrator to create an instance.

To get there, MercadoLibre investigated OpenStack and the Debian operating system, which they chose for several reasons, including their diverse and active developer community, since it's an industry open standard, as well as flexibility and its ability to scale.

The ability to scale is also what led the IT team to investigate Galera Cluster for its massive increase in data volumes.

CHALLENGES

Mercado's main challenges in terms of their data & database strategy were of a technical nature.

The active-standby schema of MySQL Cluster or DRBD didn't scale properly for them. They needed an actual active-active solution that would scale simply in processing power when adding more nodes without having to deal with failing failovers. They wanted to scale without limits.

Their private Openstack Cloud is constantly growing with more and more services being added to it; they needed a solution that would scale with the same speed and ease as their cloud does in order to support the data growth.

In more specific terms, they needed to find a solution to replace their existing database setup that consisted of MySQL with Heartbeat, as well as DRBD (active – backup) solution schema.

Case study: MercadoLibre chooses Galera Cluster

© NASA

SOLUTION

The Mercado team knew from the outset that they would want to implement an open source solution for their database strategy. They considered at first PostgreSQL XL, but eventually opted for Galera Cluster. There were several reasons for choosing Galera Cluster, one of them being that their database of choice also needed to be supported by Openstack.

The three most important criteria for their evaluation were:

- Scalability
- Avoidance of dealing with failing fail-overs in a no-go high availability schema
- Guaranteed data persistence, consistency and availability.

Finally, they decided to implement Galera Cluster for the following reasons:

- Active replication;
- Active – active schema;
- Failoverless design;
- Cost (no licence fees); and
- Scalability.

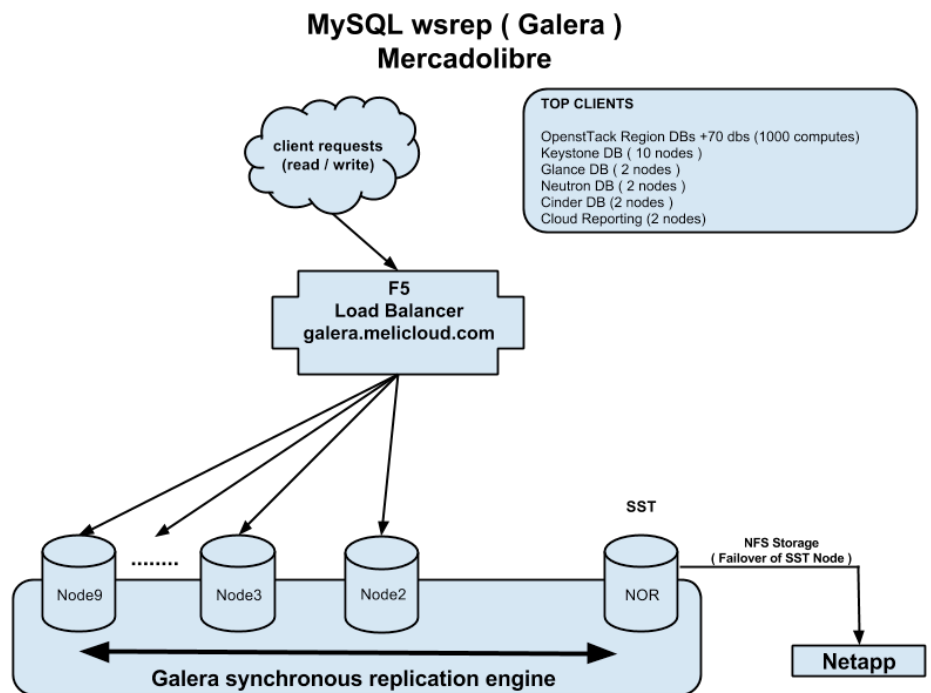
The overall solution is a well-thought-through combination of technologies.

Ubuntu is used predominantly as the operating system for host and guest of their Openstack private cloud; in fact, their Galera cluster is running on top of Ubuntu 12.04 LTS.

Openstack is being used as the main piece of software for cloud orchestration and deployment with approximately 1,200 physical servers and more than 10,000 virtual machines.

Their clients are in most cases web servers with RestAPIs that serve their sites.

The Galera clients are mainly the Openstack nodes and their orchestrator for managing APIs uses Python and SQLAlchemy to interact with MySQL.



GALERA CLUSTER SOLUTION

The Galera Cluster itself consists of 9 nodes (8 actives – 1 NOR) balanced by an F5 bigIP load balancer with more than 80 databases across the cluster and over 700 queries per second per node.

Implementing Galera Cluster took little time, mostly because it's very easy to use and to understand-- installing and configuring Galera Cluster the way that the team had planned to took them less than a month. They tested initially the solution for about three months and then went into production, as there was some urgency to replace their existing MySQL solution. They have just one endpoint configured on their automated deployments, which contributed greatly a faster implementation time. Now, if they need to scale they just add nodes without downtime or any similar challenges, which they faced previously.

The team took advantage of Codership's support, training and consulting services, which proved very beneficial. They used several tickets in order to work with Codership's support staff, who were readily available to check if the team had implemented Galera Cluster in the best way possible. This contributed to the great performance improvements the team recorded, and helped them significantly throughout the course of the project.

RESULTS

Now that their new database solution is deployed, MercadoLibre is happy to report great cost savings, scalability that will lead to further cost savings and an ease of deployment that wasn't possible before.

Just as crucially, they now enjoy real availability; they needed to be on-line with no downtime 99.99% of the time, throughout the year. System downtime is something that their business cannot afford to happen.

As their team noted themselves: "Imagine if the cloud database crashes: we lose the ability to deploy new machines, scale our infrastructure, or even control it. No downtime leads to revenue and the best user experience; and for an e-commerce site that means everything."

Their new solution provides them with many advantages they didn't have before:

- Dynamic scaling with no downtime
- 100% of availability – all year long
- Active / active query handling.



NEXT STEPS

Looking ahead, MercadoLibre expects its business to grow even faster, but the IT team no longer will have to spend so much time deploying and managing their infrastructure. Their developers will be able to provision automatically themselves with the resources they need.

“We’re opening two more data centers in the U.S. and our cloud solutions are going to be hosted there. Naturally, the data is going to be handled by MySQL Galera clusters. We’re also thinking of expanding the one that we have in Virginia in case we need it.”



Codership’s Galera Cluster community and support are fantastic. It’s difficult to find a dedicated support team such as theirs these days. The product is great, it just works as we expected it to with no surprises. It scales and we happily promote it ourselves as the defacto database for Openstack and for everyone that needs an active-active schema that really works.

- Leandro Reox & Alejandro Comisario.

SUMMARY

In Leandro Reox’s own words:

“We could not be happier with the product and the support provided by the Codership team. It’s just great and we encourage everyone to try it and use it.

Of course, one needs to consider things such as locking in a shared active-active schema in the application or abstraction layer, but we don’t see that as a ‘down-side’ to the solution, it’s simply the nature of it. We love Galera Cluster and we invite anyone who is interested to check out our implementation.”

Quoted sources: interview with Leandro Reox & Alejandro Comisario, [Openstack case study](#), [mercadolibre.com](#)

