

## Estante Virtual Leverages Percona for State-of-the-Art Solutions

### A Percona Case Study

#### Abstract

**Estante Virtual** has created a revolution in Brazil's second-hand book market. The site is a trading portal for thousands of independent traders and bookstores, with a searchable inventory of more than 7 million books. Nearly a million registered readers use it to purchase thousands of books daily, and the inventory grows by enough books to fill two bookstores every day. But as the site grew, the traffic exceeded the capabilities of the MySQL database and the web servers. At the same time, Estante Virtual's users asked for sophisticated features to make the site easier and faster to use. To meet these challenges, Percona converted the site to higher performance technologies and architecture, and performed full-stack system optimization. The result is an excellent user experience for buyers and sellers, and a high-performance application that can scale significantly with future demand.

## The Application and Architecture Search Performance Optimization

*Estante Virtual* translates to "Virtual Bookshelf" in English, and the site is driving a sea change in the way people buy and sell books in Brazil. Estante Virtual was born in 2005, when CEO André Garcia created an innovative system of subscribing, searching, and selling second-hand books online.

Estante Virtual is more than a website; it is a virtual business development agent for the thousands of sellers who use it. It brings online book selling within the reach of any registered user, making it easy for users not only to find exactly the books they are looking for anywhere in the country, but also to offer their private collections to an audience of hundreds of thousands of buyers. Bookstores and individuals who use Estante Virtual credit the site with more than half of their inventory turnover.

The Estante Virtual application is written in Perl with `mod_perl`, and is hosted at Rackspace. There is an application server, a master database server, and two identical MySQL replicas that also host a Sphinx search daemon and indexes. Despite the small number of servers used, the servers are lightly loaded, due to the performance optimizations applied to the application and database.

Performance is a priority at Estante Virtual, and the site delivers very fast response times. Many similar sites accept slow search performance as inevitable, believing that users will tolerate multi-second search delays. Estante Virtual insists on delivering search results in less than 0.1 second. This is enabled by sophisticated use of the Sphinx full-text search engine.

When CEO André Garcia hired Percona, the site was struggling with performance problems. Search performance optimization was the first priority. The solution was migrating from MyISAM-based full-text search to Sphinx. This solved several serious problems with MyISAM. MyISAM's full-text performance slows as indexes grow, updates are very slow, it suffers from lock contention due to its lack of support for concurrency, the data is prone to corruption, and repair is a lengthy process.

At the time, the MyISAM-based search indexes were so costly to maintain that inserts were queued into a holding table, and then moved to the main table every hour. This process locked the entire site due to MyISAM's table-level locking. To extract as much performance from MyISAM as possible, the `delay_key_writes` option was enabled, which defers flushing to disk upon index updates. Unfortunately, this leaves the data in memory prone to corruption, so index repairs were another cause of frequent locking. Estante Virtual attempted a nightly index update, which avoided some of the locking during peak hours, but was not acceptable to the vendors.

To optimize the search functionality, Percona employed Sphinx and converted the tables to the robust and scalable InnoDB storage engine. Now Sphinx's main indexes are rebuilt hourly, and the vendor indexes are rebuilt every minute, delivering up-to-date search results for both the general public and the vendors who are using the administrative portion of the site.

## Advanced Search Functionality

In addition to performance optimization, Percona implemented advanced search features for Estate Virtual. These features were enabled by Sphinx's support for complex search queries, but at the same time the site's complex functional requirements mandated precise fine-tuning and special techniques to work as desired. Estate Virtual was an early adopter of Sphinx 1.11, which was early in its life cycle and had several critical bugs that caused instability. Percona's close relationship with Sphinx Technologies helped to identify and resolve the problems quickly.

Estate Virtual's users compete for top position in search results, as might be expected, and some of them used "tricks" to be listed earlier when users sorted the results in specific ways. These included manipulating the titles of the books to add numbers and symbols at the front. Estate Virtual worked with Percona to identify high-performance techniques to clean up manipulated data. MySQL's stored procedures and stored functions were not fast enough, so Percona's server development experts wrote custom User-Defined Functions (UDFs) in C for raw speed. These UDFs are used for the queries that Sphinx executes to build its indexes. One of them removes "junk" characters and articles from the beginning of the input string, and another deduplicates characters within the string to help provide relevant search results for misspelled words in queries.

## Full-Stack Optimization

Estate Virtual leveraged Percona's full-stack performance expertise to optimize the entire application, not just the database and search portions. Percona taught the Estate Virtual development team how to quickly and correctly profile certain parts of their application, to identify performance bottlenecks and focus on areas that would yield the most important improvements. Percona helped with everything from the hardware to the user experience at the browser.

To obtain more performance on less hardware and reduce labor costs, Percona migrated the front-end web servers to a more scalable and maintainable architecture. For example, to ease the burden of main-

taining static images, Percona helped move images to a server where they could be delivered more efficiently. Percona also implemented a switch from Apache to the lighttpd web server for static content, which provided greater throughput and reduced overhead.

## Security and High Availability

Percona helped Estate Virtual build redundancy and security into the application, to guard against malicious attackers or accidental loss of data. For example, Percona configured a MySQL replication slave to be intentionally delayed behind the master, a technique not supported directly by MySQL replication.

## Working Relationship

"Percona is the 'family doctor' that I always dreamed of for my servers," said Garcia. "At the technical level, reaching Percona means access to a trusted network of top-level consultants. Before Percona, the only third party I trusted were my data center technicians. It was unthinkable to grant a consulting company permanent access to my servers, because I simply didn't believe a company with the expertise and organizational level I'd require could exist. After working with Percona for a couple of weeks, I realized the consultants (Aurimas, in my case) met the highest standards of organizational and technical skills. From then on my repeated experiences only confirmed that first impression."

Garcia continued, "on the business side, Percona gives me access to state-of-the-art solutions that raise my application's capabilities to another level. That translates to business success. My users can see the system simultaneously gaining new intelligent features, and becoming faster and faster."

For André Garcia, the personal relationship with Percona, and Aurimas in particular, is a delight. "I have a dedicated consultant who not only understands my setup, but he understands my business. I have to translate some Portuguese words for him once in a while, but that is part of the fun! Because he knows both where I want to get and what I care for, he knows exactly what I need before I ask! It's fantastic."

## About Percona, Inc.

Percona is the oldest and largest independent MySQL services company, with an international team of over 40 MySQL experts providing 24x7 support, consulting, training, and engineering services for MySQL to over 1000 customers since 2006. Customers include Cisco Systems, Alcatel-Lucent, Groupon, the BBC, and StumbleUpon.

Percona's founders are world-famous for their expertise in MySQL performance and scaling. They are the authors of *High Performance MySQL, 2nd Edition*. Percona also develops software for MySQL users, including Percona Server and Percona XtraBackup. Percona Server is an enhanced version of the MySQL database server. With Percona Server, users achieve faster and more consistent query execution, better insight into server behavior, and the ability to consolidate servers with commodity hardware and flash storage. Percona XtraBackup is the world's only open-source hot backup utility for MySQL's default transactional storage engine InnoDB, and is praised by social networking giant Facebook. Both Percona Server and Percona XtraBackup are free and open-source.

If you are interested in Percona's products or services, we invite you to contact us through our website at <http://www.percona.com/>, or to call us. In the USA, you can reach us during business hours in Pacific (California) Time, toll-free at 1-888-316-9775. Outside the USA, please dial +1-208-473-2904. You can reach us during business hours in the UK at +44-208-133-0309.

*Percona, XtraDB, and XtraBackup are trademarks of Percona Inc. InnoDB and MySQL are trademarks of Oracle Corp.*