



PRA QUEM BUSCA MAIS



CASE STUDY

# Catho and SanDisk®: A Perfect Match for Brazil's Top Job Search and Employment Site

## Solution Focus

- Online job search and matching
- MySQL database

## Summary of Benefits

- Up to 77% faster response system response times
- 14 to 4 server consolidation

*"We were able to reduce this number of servers to one server running the Fusion ioMemory PCIe card. It was amazing because we had a lot of traffic. Throughput was very high, while response time was kept very low. It was a very interesting solution for us."*

**Eber Duarte, Head of IT for Data Infrastructure, Catho**

## Background

Catho is the largest job offer web site in Latin America, hosting more than 300,000 open positions. The company is very well known in the region and is extremely popular in Brazil specifically. Job seekers visit Catho to find employment opportunities quickly, while employers post positions to locate highly skilled employees. The company's mission is to assist job seekers and employers to find one another in this human resources exchange that ultimately results in a more productive and effective economy.

On an average day, Catho welcomes 4,000 new users signing in for the service. Six million unique visitors were present on the site in the month of August 2015 alone.

## The Challenge

Response time for the Catho application is of prime importance. "Our challenge is to handle a lot of information from job seekers and employers, as we try to make this match between people. We need to be very fast and precise. Therefore responding in a short time is mandatory," said Eber Duarte, Head of IT for Data Infrastructure at Catho. The Catho database system is based on MySQL 5.5 and is continuously busy due to web site traffic and required replication processes.

However, the architecture team began to detect some latency in the MySQL replication processes. "This was a very difficult issue to deal with," Duarte told us. "The replication lag was increasing and there was a much greater workload on the web site." The hard disk drive architecture simply could not keep up with the I/O-intensive workloads. "In this regard, we were splitting the traffic through some slave machines and we started having too many servers to handle this traffic. We were concerned that, at some point, a user might write something on the web site and not see the updates immediately. In addition, we had some performance bottlenecks on the database side. These were more problematic internally, but were also very important to address."

The administration of the environment was also becoming very difficult because there were so many machines. The team felt that too many servers had been deployed and they were overwhelmed with maintenance tasks. "As we split the traffic among many servers, the read and write operations were happening on different servers," Duarte explained. "Data was replicating over the wire and this data needed to be re-written in the slave machines, which was, at times, introducing some latency into the process." The team was concerned that these issues, if left unchecked, would begin to directly affect user experience. Catho wanted to ensure that recruiters could continue to access profiles and CVs quickly, without having to wait for data to refresh.



a Western Digital brand

## The Solution

The Catho database team began to consider alternatives to address these performance issues. “We considered buying more servers to handle the traffic, but this would have adversely affected our business financially,” said Duarte. This solution would have required that Catho increase the number of servers, acquire additional space in the data center, and spend additional dollars on power and cooling.

“I was at a MySQL conference in Santa Clara and saw a benchmark where the guy was comparing regular disks, hard drives, and SSDs, and they were using Fusion ioMemory PCIe cards. The results were very impressive—specifically the amount of traffic and response time that the server could handle. At that time I was really curious about using these cards just to see how they would behave in our environment.”

Duarte determined that Catho should test the Fusion ioMemory™ ioDrive®2 PCIe cards in their environment. He reached out to SanDisk staff in Brazil, who provided an opportunity to conduct a Proof of Concept evaluation of the Fusion ioMemory PCIe card. “We just set up an environment with one server that was in production and replaced the old hard disk drives with Fusion ioMemory PCIe cards. Then we returned the server to handle the traffic back into production. The results were impressive.”

Catho initially had 14 servers to handle one slave cluster of traffic. “We were able to reduce this number of servers to one server running the Fusion ioMemory PCIe card,” said Duarte. “It was amazing because we had a lot of traffic. Throughput was very high, while response time was kept very low. It was a very interesting solution for us.”

Because Catho has geographical redundancies, the company operates two data centers. “We now have just two servers in each data center because we cannot have a single point of failure and we kept two servers in each of the two data centers to handle the traffic,” explained Duarte. “We would be able to process all the requests using only one server. The same server that we were using with the regular hard drives is the same as the one that we are using with the Fusion ioMemory cards and the performance is remarkable.”

The SanDisk deployment team spent some time onsite to assist Duarte and his team with tuning at the file system level, tweaking some configurations in MySQL Server, and tuning parameters to obtain all the benefits of the Fusion ioMemory solution. “We were pleased with our experience with SanDisk. They have very good, highly skilled staff who were responsive to our needs. We discussed things at a very deep level. The success of this project was due to our ability to work together.”

Before acquiring the Fusion ioMemory PCIe cards, Catho had 14 servers in a general purpose cluster with the following configuration:

- x86 server
- Processor : Intel® Xeon® CPU E5-2665 0 @ 2.40GHz (32 cores)
- 126GB of DRAM
- Four 600GB SSD drives
- Two 300GB SAS drives

Catho replaced the 14 servers with only four servers with the following configuration:

- x86 server
- Processor : Intel® Xeon® CPU E5-2643 v2 @ 3.50GHz (24 cores)
- 126GB of DRAM
- Two 300GB SAS drives
- One Fusion ioMemory ioDrive2 PCIe card

---

*“The main result was with respect to server density. We reduced the number of servers from 14 to four. We used to have 14 servers to deal with this and we were able to handle this same traffic with one server. In addition, during the test in the production environment, we were able to handle around 3000 queries per second.”*

**Eber Duarte, Head of IT for Data Infrastructure, Catho**



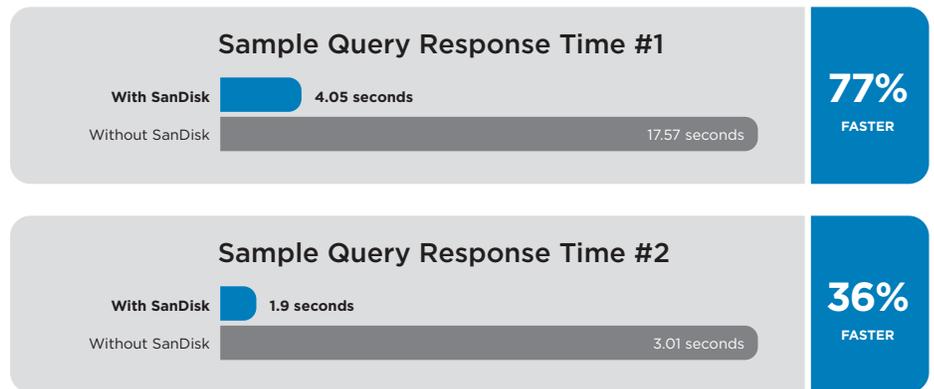
Fusion ioMemory ioDrive2 PCIe card

## The Result

Catho felt that it was very important to reduce the growth of the footprint in the data center. “The main result was with respect to server density,” Duarte confirmed. “We reduced the number of servers from 14 to four. We used to have 14 servers to deal with this and now we are able to handle this same traffic with one server. In addition, during the test in the production environment, we were able to handle around 3000 queries per second.”

The Catho team collected some slow queries in the MySQL log and compared the execution of these queries with the Fusion ioMemory PCIe card. The team saw a 30 percent faster response time when compared with traditional hard disk drives. “In the tests that we performed, one server with the Fusion ioMemory PCIe card was able to handle all the website traffic. However, we bought four for availability reasons—two servers for each datacenter. In addition, the average queries per second in this new server are up to 3000. In the old server we almost were able to hit this number, but we needed seven servers to perform close to this. We also perceived a great query response time improvement.”

A sample query was tested against both configurations, with and without the Fusion ioMemory PCIe card.



### Contact information

fusion-sales@sandisk.com

#### Western Digital Technologies, Inc.

951 SanDisk Drive  
Milpitas, CA 95035-7933, USA  
T: 1-800-578-6007

Western Digital Technologies, Inc. is the seller of record and licensee in the Americas of SanDisk® products.

#### SanDisk Europe, Middle East, Africa

Unit 100, Airside Business Park  
Swords, County Dublin, Ireland  
T: 1-800-578-6007

#### SanDisk Asia Pacific

Suite C, D, E, 23/F, No. 918 Middle  
Huahai Road, Jiu Shi Renaissance Building  
Shanghai, 20031, P.R. China  
T: 1-800-578-6007

For more information, please visit:

[www.sandisk.com/enterprise](http://www.sandisk.com/enterprise)

**SanDisk®**  
a Western Digital brand

At SanDisk, we’re expanding the possibilities of data storage. For more than 25 years, SanDisk’s ideas have helped transform the industry, delivering next generation storage solutions for consumers and businesses around the globe.

When Catho tested a few queries obtained from the MySQL slow-log, the servers without Fusion ioMemory PCIe cards responded 30 percent more slowly on average.

Catho customers will likely also benefit from this deployment. Although, at one time, the IT team was concerned how a performance issue could potentially affect the customer experience and customer service, with the deployment of the new configuration, the team no longer has this concern.

### Overall experience

To Duarte and his team, the performance gains were very clear. The project as a whole, including deploying the Fusion ioMemory PCIe cards proved to be very simple. “We have had the server up and running for almost seven months now and we have no stopping or problems or anything that has jeopardized the productivity of the servers,” confirmed Duarte. “It is really easy to use. It is very easy to set up and get things working.”

The Catho team also plans to continue deploying Fusion ioMemory PCIe cards when needed. “The main goal of our database infrastructure is to expand the use of Fusion ioMemory PCIe cards. We just deployed this equipment in one cluster of slaves and we have many others where we are going to expand the use of these cards to reduce the number of servers and handle our traffic more efficiently.”

The performance results and cost savings discussed herein are based on internal testing and use of Fusion ioMemory products. Results and performance may vary according to configurations and systems, including drive capacity, system architecture and applications.

©2016 Western Digital Corporation or its affiliates. All rights reserved. SanDisk is a trademark of Western Digital Corporation or its affiliates, registered in the United States and other countries. Fusion ioMemory, ioDrive, and others are trademarks of Western Digital Corporation or its affiliates. All other trademarks are used for identification purposes only and may be the property of their respective owners. Catho\_CS\_SanDisk\_v4 06/03/16 5023EN