itWorldCanada.com

Print

Abebooks finds clusters affordable, scalable
By: Rebecca Reid
(05 Mar 2004)

After exponential growth in its rare book business, the Advanced Book Exchange Inc. (Abebooks) desperately needed an advanced IT system.

When it launched in 1995, Abebooks maintained a listing of 4,000 books from 12 vendors. Just nine years later, it lists 55 million books from 12,000 book sellers, located in 42 countries. Abebooks has over two million customer records in its database, half a million of which change daily, said John Snider, manager of operations at Abebooks in Victoria, who was an Oracle consultant until 1999.

To increase scalability after its tremendous growth spurt, the company upgraded from Microsoft Corp.'s FoxPro to a database from Oracle Corp. in 1998, Snider said. Abebooks has also upgraded its hardware almost every year since its inception and currently runs Oracle 9.

In Sept. 2003, the cost of buying more powerful hardware finally caught up with Abebooks and the company needed more processing power at a reasonable price. That's why it turned to Oracle's Real Application Cluster (RAC) application software, allowing Abebooks to spread out processing power between several smaller, less expensive servers, rather than run it on one larger server.

"The long-term goal was...some sort of capability of managing our growth in a cost-effective manner," Snider said.

With RAC, as users run out of capacity they just need to add more nodes, said Sam Samhouri, regional sales manager for Western Canada at Oracle in Vancouver. Plus they can buy lower-end boxes rather than having to purchase high-end servers, he added.

"We'd always been migrating our database from one server to another server and using these monolithic Symmetric Multi-Processing (SMP) boxes," Snider said. "From the time I started until today we've seen one order of magnitude increase in size of our database. When I started there were about five million books, now there are 55 million books — ten times the number and we don't see an end to the growth."

Fast forward to Sept. 2003.

"We were juggling balls with the servers — everything came crashing to a halt because the environment couldn't support the number of transactions and the number of books," Snider said. "The environment started slowing down and it started crashing."

Now they're running the cluster on two Intel Corp. Xeon-based servers on Red Hat Inc.'s Linux. Snider said the Intel boxes run twice as fast as their old Sparc-based servers. In addition, the Abebooks purchased a third server, which is already configured for when the company needs more capacity.

With the cluster, Snider said batch processing is separated from online functions, increasing the speed of the Web site. With the company's previous architecture, the batch processing would steal CPUs from online processing. Abebooks has also benefited from the redundancy clustering provides. Snider said a hardware problem before RAC would have meant about one or two hours of downtime.

It took three days for Abebooks to install the RAC software on the nodes and five days to migrate the 300GB of data over to the new systems.

When Snider joined Abebooks in June 1999 , the company had progressed from running the Oracle database on a Microsoft Corp.'s Windows NT server, to running it on an IBM Corp. pSeries RS/6000 F40 server with AIX and then upgraded again to an IBM pSeries RS/6000 M80 server.

Snider added some RAM to the M80 but eventually Abebooks maxed out the two CPUs on the server. The company looked into adding two more CPUs but IBM quoted them a price of more than $100,000, Snider said, which was too costly for the company.

Then Abebooks took a server it had been using at its failover site and put it into production — a V480 from Sun Microsystems Inc. because it only cost between $35,000 and $40,000 to add two more CPUs to it — and swapped the IBM M80.

This is when Abebooks had to consider some options. It couldn't upgrade the V480 because it was maxed out. So it came down to either upgrading the IBM M80 or buying a bigger, faster server. Updgrading the IBM M80 would have been even more expensive than the initial quote of $100,000 because Abebooks would have had to upgrade disk arrays and memory along with the CPUs.

Until a decision was made, Abebooks ended up running its entire database on the IBM M80 and V480, meaning it had no failover in case the system crashed.

This is when Snider started seriously considering RAC but the company was still hesitant because migrating to a clustered environment would meaning buying all new hardware.

"We costed it out and the price to build a two-node cluster with Oracle RAC was basically the same price as upgrading our existing environment," Snider client. Although the company didn't see immediate cost benefits, Snider said Abebooks would receive long-term ROI from it.