



Case Study

Database

Intel® Xeon® processor
7300 series

eBusiness / Technology

Yahoo! Search Turns Up Quad-Core

Yahoo! looks for ways to enhance the performance of its mission-critical Oracle* database and chooses the Intel® Xeon® processor 7300 series

As a leading Internet search provider, Yahoo! is constantly looking for ways to improve the user experience. The company's business units analyze user click-through data to determine which pages, information, and ads are of most interest to readers, and then apply that information internally to fine-tune page layout and ad design. To enable this research, the company maintains a 1.4-petabyte database containing 90 days' worth of click-through information. The data is continually refreshed with the most recent day's activity.

Loading each new day's data was taking too long—as much as 15 hours—on the company's aging database servers. "We needed to reduce the time to 3 hours," says Tim Heath, Yahoo!'s chief database architect. "That way, the data would come in at 6 a.m. and finish loading by 9 a.m., making it available to our analysts throughout the day."

Enabling faster 1.4-petabyte queries

The Yahoo! team also needed to shorten the time required to run the most demanding research queries, those involving the full 1.4 petabytes of data. The goal was to cut the average time from 20 hours to just 8 hours, so that analysts could complete the huge queries in one workday. Only commercially available hardware and software could be used to achieve this goal, because the team needed to ensure full compatibility with the Structured Query Language (SQL) standard. SQL-based tools would enable non-engineering business personnel to easily write queries.

Heath's team tested servers and processors from multiple vendors and decided to standardize on 4-socket IBM* x3850 servers with the Intel® Xeon® processor 7300 series. "The Intel quad-core processors completely eliminated the performance limitations we were experiencing with our dual-core processors from another vendor," says Heath. "We are no longer CPU-bound, no matter how large the query."

Measures of Success

- Yahoo! needed to speed up database queries to meet the demands of the company's business units
- Employees had to be able to run the most demanding queries within one workday, instead of the 20-plus hours the queries had been requiring
- Loading time for each new day's data had to be decreased from 15 hours to less than 3 hours so that the data would be available throughout the day
- The Yahoo! IT team needed to meet these goals with commercial hardware and software to keep costs down and enable users to easily run queries

"Now people can get to work and hit the ground running, performing time-sensitive research without waiting hours for the system to finish loading."

Tim Heath
Chief Database Architect
Yahoo!

Yahoo! uses its 1.4-petabyte Oracle* database powered by Intel® quad-core technology as a tool for continually analyzing and improving the user experience.



Full-database queries run 93 percent faster

The Yahoo! team was able to meet or exceed its performance goals using Intel quad-core technology-based servers in a 16-node cluster running the Red Hat Enterprise* Linux* operating system and Oracle* 10g database software. The team cut the time required to run demanding, full-database queries from 20 hours to 73 minutes, exceeding the original goal of reducing the query time to 8 hours. "That's 93 percent faster, which greatly increases the amount of research employees can do in a day," says Heath.

Data loading time reduced from 15 hours to 3 hours

Loading time for each new day's data has decreased dramatically as well. "We went from 15 hours to less than 3 hours to load an average of 10 terabytes of data," says Heath. "The database is completely refreshed by the time employees arrive for the day. Now people can get to work and hit the ground running, performing time-sensitive research without waiting hours for the system to finish loading."

No additional operating costs

The IT group has been able to run the new systems without any additional operating costs. "Our staff was already familiar with IBM and Intel technology, so they didn't need any special training," says Heath. "And with the energy efficiency of the Intel processors, we haven't had to install any special power or cooling gear. Intel technology is a big part of our planning—we're already looking forward to six-core Intel processors to power our next advance in analyzing and improving the user experience."

Return on Investment

- Yahoo! was able to accelerate database queries with a cluster of 16 commercially available servers based on the Intel® Xeon® processor 7300 series
- The Yahoo! IT team has cut the time required to run the most demanding queries by 93 percent, exceeding their reduction goal
- Loading time for each new day's data has decreased from 15 hours to less than 3 hours, making the data more available for analysis
- The Yahoo! IT team is able to run, maintain, and manage the new systems without any additional training or operating costs



Find a business solution that is right for your company. Contact your Intel representative or visit the Reference Room at www.intel.com/references

This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

Intel may make changes to specifications, product descriptions and plans at any time, without notice.

Intel, the Intel logo, and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2008 Intel Corporation

0908/YMB/TDA/XX/PDF 320163-001US

