

Oracle Exadata Achieves Fivefold Performance Increase for Critical Product Research Platform



Merck & Co.

Merck is a global research-driven pharmaceutical company. Merck discovers, develops, manufactures and markets a broad range of innovative products to improve human and animal health, directly and through its joint ventures.

CambridgeSoft is a leading provider of software and services for discovery, analysis and collaboration to life sciences and chemical industries, academia and government. Products include Cloud, Enterprise, Workgroup and Desktop versions of Chem & Bio Office, including Chem & Bio Draw and E-Notebook, providing knowledge management, chemical and biological informatics, and scientific database solutions.

Founded: 1891

Headquarters: Whitehouse Station, NJ

Revenue: \$27.4 billion (2009)

EXECUTIVE SUMMARY

A modern version of the traditional laboratory notebook, the electronic laboratory notebook, or ELN, is becoming a critical asset in the pharmaceutical industry, giving scientists new power to collaborate and share experimental data with colleagues from various research divisions. Merck & Co. embraced the CambridgeSoft ELN technology as an early adopter in 2002, and the application was quickly accepted by Merck's global research teams. As a result, the volume of ELN-documented experiments and number of users exploded. Today, with over 4,000 users, the Merck ELN is one of, if not the largest deployed ELN system in the pharmaceutical industry worldwide.

With the rapid growth of users and data volume, Merck and CambridgeSoft worked diligently to keep pace with the expansion and simultaneously maintain performance. As a result, changes were made to the application itself as well as to the infrastructure.

Merck's planning made it clear that rapid user growth was expected to continue and correspondingly, the size of the database would also continue growing very rapidly. This growth was expected to outpace the application of conventional approaches to maintain acceptable performance. In addition to CambridgeSoft's engineering enhancements and traditional infrastructure changes, Merck decided to explore non-traditional options to provide a long-term solution.

Merck found a solution in Oracle Exadata Database Machine, a new computing platform that takes advantage of innovations in smart storage to radically boost system speed and cost performance. The company rolled out the Exadata solution in late 2010, using Oracle GoldenGate data-synching technology to make the switch in a single weekend, as opposed to taking weeks for migration using conventional software. Essentially, there was no business downtime.

Following the move, ELN performance surged, with Merck conservatively estimating a fivefold increase in query response speed as well as substantially faster user login and navigation performance, even as the company simultaneously expanded the user base by 30%. Other improvements have included lower system resource utilization, optimized storage space, and workload consolidation.

All told, Merck is expected to realize benefits totaling an estimated \$8 million from its Oracle Exadata investment.

“We really knocked it out of the park. The chemistry group, which is one of the biggest users of E-lab notebook, is very happy with what we delivered.”

MERCK

ORACLE PRODUCTS

- Oracle Exadata Database Machine
- Oracle GoldenGate
- Oracle Database 11g

KEY BENEFITS

- 159% ROI over 3 years
- Increased speed of ELN searches by 5 times
- Total benefits of \$8M over 3 years
- Reduced memory utilization by 40%, reducing the need for further immediate investment
- Reduced CPU utilization to 10%, providing significant capacity for growth
- Doubled the size of the user base without impact to performance
- Increased scientist efficiency

Beyond the direct infrastructure and productivity savings, the exponential jump in ELN performance enabled by Exadata is empowering scientists to quickly access experimental data and collaborate more effectively with other research teams—a benefit that holds tremendous potential in terms of future product development, innovation, and revenue. With the increase in system responsiveness, Merck’s confidence in the system’s scalability improved, as did the user experience. Scientists now have a tool that keeps up with their way of doing business.

BACKGROUND

One of the world’s largest pharmaceutical companies, Merck & Co. discovers, develops, manufactures and markets a broad range of innovative products to improve human and animal health, directly and through its joint ventures. New Jersey-based Merck operates in 140 countries¹ and in 2009 significantly expanded its market share with the acquisition of Schering-Plough, the maker of Claritin, Coppertone, and many other brands.



At the heart of Merck’s operations are the thousands of researchers and scientists who invent, develop and test new products in a network of laboratories that span a wide geography. Collaboration is essential to the success of these research teams: scientists are more productive when they can quickly access the experimental results and techniques of their colleagues. Indeed, increasing collaboration is a major reason why pharmaceutical companies have adopted electronic laboratory notebooks, or ELNs—a secure, highly searchable form of the traditional paper-based lab notebooks.

Merck made the move to online lab notebooks in 2002 when it rolled out CambridgeSoft E-Notebook, a leading ELN solution. Today, ELN is a strategic asset for the company and the primary means for capturing the intellectual property created by Merck’s scientists. Scientists document the discoveries and conclusions of their experiments in the online notebooks and then search the database for information to help them plan the next phase of their experiments.

For example, a chemist working on a promising new drug compound will search the ELN database to learn if anyone else had made a similar or related compound—and then use that information to plan the next stages of his experiment. Searching collective information on experiments enables Merck scientists to be more productive and leads to greater innovation.

The growing use of ELN, however, began to put strains on system capacity and performance. Between 2003 and 2010, the user base doubled and the size of Merck’s ELN database grew to more than 4 terabytes.

¹ From Merck corporate Website

**Figure 1: Active Users of Merck's ELN Application****Figure 2: Number of Experiments Documented in ELN Application**

With system usage and database size soaring, information searches began to slow. "As more and more people plugged into the ELN platform, more and more people wanted to use it," said Merck. "That meant that the platform's performance got progressively worse over time as we struggled to keep up with infrastructure enhancements. CambridgeSoft also worked with us to streamline the search process, which yielded tangible improvements."

Still, slowing performance hurt researchers' productivity, since scientists were hindered from taking full advantage of the system to assist with research.

Merck's initial attempts at fixing the problem with conventional methods—including adding hardware and memory and upgrading CPUs—met with limited success as the expanding user base, experiments and searches continued to outpace efforts to upgrade system hardware. Consequently, in 2010, Merck resolved to permanently solve its ELN performance issues by embracing the next generation of database and storage technologies.

HOW EXADATA MAXIMIZES PERFORMANCE

- Exadata Smart Scan.** The smart storage software in Exadata offloads data-intensive query processing from Oracle Database 11g servers to Exadata's storage layer for parallel data processing. Because there's less data moving through the higher-bandwidth connections, you get dramatically improved performance as well as concurrency for simple and complex data warehousing queries.
- Exadata Smart Flash Cache.** With more than 5 terabytes of flash memory per full rack, Oracle Exadata intelligently caches "hot" data and assigns the rest to disk storage, giving organizations the speed of flash with the cost-effectiveness of disk storage. Exadata Smart Flash Cache can process up to 1.5 million random I/O operations per second and scan up to 50 GB of data per second to deliver ultra-high performance for OLTP applications.

THE ORACLE EXADATA SOLUTION

Oracle Exadata emerged as the logical alternative to conventional database technologies, offering significantly faster performance for both data warehousing and online transaction processing (OLTP) applications, including electronic lab notebooks. Oracle Exadata combines servers, storage, networking, and software in an integrated platform that is massively scalable, secure, and redundant.

Following a successful pilot, which included CambridgeSoft's participation in the Exadata optimization, Merck deployed Oracle Exadata in late 2010 as the standard platform for its ELN solution. Because the switch to Exadata represented a major system change involving both new hardware and new Oracle software, Merck was anticipating more than a week of downtime for the migration. But Merck minimized downtime to a single weekend by deploying Oracle GoldenGate to keep the old and new environments synchronized during the transition.

"We were literally able to flip the switch to launch the new platform," said Merck.

"Scientists left Friday afternoon and returned Monday morning to a new system that was up and running with all the old data. There was no negative impact to the scientists." Moreover, Merck wasn't required to make any changes to the ELN application itself to run on Oracle Exadata.

OPERATIONAL AND STRATEGIC BENEFITS

In the months following the switch to Oracle Exadata, Merck has seen a host of operational and financial benefits ranging from dramatically faster system performance to millions of dollars in avoided solution-replacement costs. From a strategic perspective, the move to Exadata enables Merck's scientists to more fully utilize the ELN platform, boosting their research capabilities and potentially speeding product innovations.

Fivefold Performance Increase

With Oracle Exadata, Merck improved ELN search speeds by at least four-to-five times, easily meeting the IT group's service level agreement (SLA) target of 30 seconds for most searches. Timeouts were eliminated. "Overall, we realized about a 5x improvement in performance, even with complex searches," said Merck.

As shown in Table 1, the move to Exadata led to enhancements across a range of ELN performance measures—from faster user logins and navigation, to speedier text and structured-data searches. Merck's IT team calls the fivefold speed boost a conservative estimate because many of the queries in the legacy environment were terminated before returning results. Other estimates place the performance boost at nearly 10 times (in a controlled environment).




Table 1: Performance Increase Achieved by Merck with Oracle Exadata

	Before: Average time (in seconds)	After: Average time (in seconds)	Average Performance Improvement
Text Search (27 Steps)	99.8	23.2	4.3x
Structured Search (43 Steps)	74.2	14.0	5.3x
Structured Search Combined with Text Search (15 Steps)	232.8	40.0	5.8x

“We really knocked it out of the park,” said Merck. “Our chemistry group, which is one of the biggest users of E-lab notebook, is very happy with what we delivered.” Merck attributes the faster search performance to two technology advancements found in the new Oracle platform: Exadata Smart Scan, which filters out data that is not relevant to a query, and Exadata Smart Flash Cache, which improves performance for OLTP applications.

Lower Utilization Rate

The database efficiency improvements driven by Exadata—which offloads a large portion of data processing to the storage layer—has had a positive impact on system utilization rates at Merck: Memory utilization is now running around at 50–60% of capacity and CPU utilization at around 10–15%—a significant reduction in both cases.

More Productive Scientific Research

The move to Exadata means scientists are more likely to probe every promising idea. “Combined with the CambridgeSoft ELN, Oracle Exadata allows scientists to execute more sophisticated searches and get faster access to information they previously weren’t able to easily access,” said Merck. This also means that there are no concerns about adding thousands more users to the CambridgeSoft E-Notebook platform as is currently planned.

ESTIMATED \$8 MILLION IN TOTAL BENEFITS

Figure 3 illustrates the total benefits of Merck’s Oracle Exadata investment over three years. Additional savings will come from writing off Merck’s legacy database servers and storage equipment, and from productivity improvements related to faster searches and research collaboration.

ABOUT THIS CASE STUDY

Research and analysis for this study was conducted by Mainstay Partners, an independent consulting firm, drawing from interviews with Merck employees, review of planning documents and searches of industry literature. ROI calculations use industry-standard assumptions regarding the time value of money.

Mainstay Partners is the leading provider of independent value assessment and IT strategy services. Our clients include Oracle, SAP, Cisco, Lexmark, Microsoft, Motorola, and HP. For more information, please visit www.mainstaypartners.net. Information contained in this case study has been obtained from sources considered reliable but is not warranted by Mainstay Partners.




901 Mariners Island Blvd, Ste. 105
San Mateo, California 94404-1592

650.638.0575 *ph* 650.638.0578 *fx*

www.mainstaypartners.net

Figure 3: Benefits by Category—Three-Year View



CONCLUSION

Facing an increasingly globalized and competitive marketplace, pharmaceutical companies are being pressured to develop and launch products faster than ever before. That new reality has put a premium on laboratory research efficiency and productivity. More than ever, companies are seeking ways to accelerate the experimental process, collaborate more efficiently, and make rapid decisions regarding research projects.

One way that Merck has responded to the challenge is by deploying Oracle Exadata to transform the data-processing backbone of its ELN application. As a result, Merck scientists are finding and sharing research information five times faster on average, and they are now accessing valuable experimental data that was previously inaccessible.

The business potential unleashed by Merck's ultra-fast ELN is enormous. Today, scientists are significantly more productive and connected to the research of their colleagues across the company, which is strengthening collaboration and building the foundation for rapid, market-beating innovations.