



# Expanding the Resources for Next-Generation Seed Discovery

Monsanto increases performance for seed research by up to 28 percent with an expanded HPC cluster based on Intel® Xeon® processors



MONSANTO



“When we realized that we could achieve substantially better floating-point performance and larger memory capacity with the Intel® Xeon® processor E5 family, it was an easy decision to select these processors for the HPC expansion project.”

– John Rothermel,  
Senior UNIX®/Linux® Engineer,  
Monsanto

For Monsanto, high-performance computing (HPC) plays a critical role in the discovery of next-generation agricultural products that can help farmers produce more food with fewer resources. To accommodate increasing demand for HPC resources from company researchers, the IT group decided to expand its existing HPC cluster by adding Cisco Unified Computing Systems\* (UCS\*) B-Series\* blade servers equipped with the Intel® Xeon® processor E5 family. The new processors help to accelerate new product development by increasing workload performance up to 28 percent. At the same time, using the dense, energy-efficient processing architecture in blade servers helps the company add capacity while controlling costs.

## CHALLENGES

- **Increase capacity and support critical science.** Expand the primary HPC cluster to accommodate increasing demand for HPC resources.
- **Enhance performance.** Capitalize on the latest processing technologies to improve performance of research applications.
- **Control costs.** Provide greater capacity to more users on the primary HPC cluster while controlling expansion costs and reducing the need for auxiliary environments.

## SOLUTION

- **Cisco UCS B-Series blade servers with the Intel® Xeon® processor E5 family.** Monsanto's IT group is adding 608 cores to its 1,000-core cluster by incorporating Cisco blade servers based on the Intel Xeon processor E5 family. Workload scheduling is managed by IBM Platform LSF\* software.

## TECHNOLOGY RESULTS

- **Increased application performance.** The new processors are helping boost application performance by up to 28 percent compared with previous-generation processors.
- **Greater processing density.** Blade servers equipped with the Intel Xeon processor E5 family helped increase core capacity by nearly 61 percent while expanding rack space by only 22 percent.

## BUSINESS VALUE

- **Controlled costs.** A dense infrastructure helps support growth while controlling costs. Running more workloads on the primary cluster also reduces the need to buy and manage standalone servers.
- **Faster product development.** Providing researchers with greater capacity and improved performance ultimately helps the company deliver next-generation seeds to farmers faster than before.

Monsanto researchers are constantly at work developing new seeds that can help farmers increase yields and improve their crops' resistance to disease and pests. The researchers rely on HPC resources to handle the compute-intensive process of genome assembly, which helps produce the genetic maps that guide new product development.



## Intel® Xeon® processors provide a robust, reliable architecture for complex HPC workloads

In 2011, the company migrated these and other research workloads from a competing processing architecture to Intel Xeon processors to increase HPC performance. The new 1,000-core cluster improved performance by more than 30 percent and accelerated genome assembly by 80 percent compared with the previous environment, all while controlling power, cooling, and real estate costs.

By 2012, it was already time to expand that cluster. "As the company grows, so does demand for HPC resources," says John Rothermel, senior UNIX/Linux engineer at Monsanto. "We needed to significantly expand our primary cluster's capacity to keep up with that demand."

The IT group decided to pursue the expansion project while simultaneously implementing a new workload-scheduling solution. "Some of our users had difficulty scheduling their workloads on our primary cluster, so they consequently moved the workloads to distinct, standalone systems," says Rothermel. "We wanted ways to better serve our users with the primary cluster while also improving overall cluster utilization and simplifying management."

### Expanding the Cluster with the Intel Xeon Processor E5 Family

The first key decision in the expansion project was to adopt blade servers. "We had been using Cisco UCS B-Series blade servers internally for other IT needs," says Rothermel. "We saw an opportunity to bring some of the density, flexibility, and simplified cabling benefits of blades to our HPC environment"

As the IT group evaluated the latest Cisco blades, they discovered some of the advantages of the Intel Xeon processor E5

family. "When we realized that we could achieve substantially better floating-point performance and larger memory capacity with the Intel Xeon processor E5 family, it was an easy decision to select these processors for the HPC expansion project," says Rothermel.

The IT group also knew that the Intel processors would help keep workloads running without problems. "Some workloads run for several days. If something fails in the middle, researchers have to start again from the beginning," says Keith Turnbull, infrastructure engineering lead at Monsanto. "The Intel Xeon processors provide the mission-critical reliability our researchers need to stay productive."

The IT group decided to deploy Cisco UCS B-Series blade servers, adding a total of 608 new cores to the cluster. Researchers use homegrown applications as well as utilities for genome analysis and bioinformatics such as NCBI Basic Local Alignment Search Tool\* (BLAST\*) and hyaluronan-mediated motility receptor (HMMR) utilities, all running on Red Hat Enterprise Linux\*. To improve workload scheduling, the IT group implemented IBM Platform LSF software for the newly added servers.

### Expanding Capacity While Controlling Costs

The expanded cluster will enable the IT group to support more—and larger—workloads on the primary cluster. At the same time, adopting dense, energy-efficient blade servers with the Intel Xeon processor E5 family is helping keep costs under control. "We increased the capacity of our cluster by 61 percent while increasing rack space by only 22 percent," says Rothermel. "We can support more research and help researchers solve more complex problems without dramatically increasing power, cooling, real estate, or management costs."

By providing adequate capacity for more HPC workloads and improving workload scheduling on the primary cluster, the IT group is reducing researcher dependence

### LESSONS LEARNED

The Monsanto IT team deployed the new servers as a separate farm initially, migrating only some of the existing applications. "It was a much more seamless experience to migrate applications from one farm to another than having to activate an entire 1,608-core cluster all at once," says Rothermel. "Fortunately, the new Intel® Xeon® processor E5 family used with the added servers can provide more than enough performance to handle the migrated workloads until we are ready to start running all the servers as an integrated cluster."

on those discrete, standalone systems. "We are able to free up existing standalone server resources for other projects and avoid the costs of buying and managing new ones," says Turnbull.

### Boosting Application Performance by Up to 28 Percent

The new processing architecture is also helping to deliver substantial performance improvements compared with the previous-generation processors. "Running BLAST, the Smith-Waterman\* algorithm, and an MPI\*-based protein design application in early testing, we recorded performance boosts of 15, 18, and 28 percent, respectively, with the Intel Xeon processor E5 family," says Rothermel. "In actual use, some of our users have seen even greater improvements for their workloads compared with the performance they were achieving on those standalone systems."

### Accelerating Results and Speeding Time to Market

Those performance improvements are helping speed up research considerably. "One research team moved a workload from a standalone server to our expanded HPC cluster and accelerated results from two weeks to just six days," says Rothermel. "With faster results, our company can ultimately bring next-generation products to market faster than before and better address a range of critical challenges facing farmers today."

**Find a solution that is right for your organization. Contact your Intel representative, visit [Business Success Stories for IT Managers](#), or explore the [Intel IT Center](#).**



Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

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, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

© 2012, Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

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