

Platform™

Incyte Genomics Shortens Production Cycles and Doubles Productivity with Platform Solutions

Overview

With one of the industry's largest Linux-based data centers, Incyte needed to optimize its heterogeneous compute resources to accelerate discovery.

Challenge

- Incyte had developed its own Linux clustering technology, which made it difficult to reallocate unused processing power.
- Incyte needed visibility into how IT resources were being used to integrate them into the cluster.

Solution

Incyte selected Platform LSF® and Platform Analyzer on a Linux cluster of 1,200 CPUs for a dedicated production pipeline, where several million jobs will be handled automatically.

Results

- The fully optimized allocation of resources achieved by Platform LSF increased Incyte's overall productivity by 50% in the first month.
- The substantial improvement in utilization of compute resources is having a significant impact on time-to-market, enabling Incyte to launch new products faster than ever before.
- Platform LSF and Platform Analyzer are quickly becoming some of Incyte's most critical pieces of infrastructure – equally as important as its disk farms for databases and environmental equipment.

As one of the leading drug discovery companies, Incyte Genomics has the largest commercial portfolio of issued U.S. patents covering full-length human genes and the proteins they encode, information that has become increasingly valuable to the discovery and development of new drugs and diagnostics.

Incyte's data center houses 4,500 generic Intel, Compaq Alpha, Sun Microsystems and SGI CPUs used for high throughput processing. Approximately half of those CPUs are Linux®-based. Because they are cost-effective and reliable, the Linux boxes handle most of the expensive, 'heavy-lifting' computation that Incyte performs.

To optimize the computing capacity of its existing Linux clusters, Incyte needed to identify and gather idle computing resources throughout the enterprise, automatically allocating them on a priority basis to shorten production cycles and meet increasing time-to-market demands.

Incyte's search for a robust, reliable and mature solution that met its demanding enterprise requirements led it to Platform LSF and Platform Analyzer. The Platform solution meets Incyte's compute-intensive needs by providing full visibility and flexible, seamless, on-demand access to unutilized computing resources and automatically matching supply with demand across Linux clusters.

Business Challenge

As a strong proponent of distributed computing versus very large hosts, Incyte developed its own Linux clustering technology three years ago. However, the main drawback of the custom code was that a Linux cluster had to be dedicated to one specific application, making it difficult, if not impossible, to collect the unused processing power and reallocate it to accommodate additional applications or one-time ad-hoc projects. If a job was running behind schedule, necessitating an expansion of processing power, the CPU had to be reconfigured manually, which was time-consuming and an inefficient use of human resources, contributing to significant lost productivity.

As Incyte's customer base grew and new applications were introduced, the data center expanded and additional CPUs were added. By 2000, the situation had become critical. Incyte needed proven, industrial-strength software to harness the growing amount of available computing resources across Linux clusters. The software would need to integrate seamlessly with other applications supported by an individual cluster, and coexist with Incyte's proprietary cluster technology.

"Platform LSF software definitely delivers from a performance standpoint, and has made greater computing capability available to more personnel at Incyte. Our testing ranks Platform LSF as a well-built, extremely reliable, top-quality software solution. Even the programmer on my team who did most of the software development for Incyte's proprietary software is a big fan of Platform LSF, and thinks that it's far superior."

Stu Jackson

Director of Bioinformatics



Solution

To gain visibility into how its IT resources were being used, Incyte selected Platform Analyzer to collect baseline utilization statistics for the cluster, such as system load, CPU usage, pending jobs by user and job pending time. With these results, Incyte was able to make changes to maximize its computing power, which it continues to do on an ongoing basis. The results clearly indicate a sharp increase in the overall utilization of the environment after Platform LSF was installed and enabled.

As part of its evaluation process, Incyte ran a full-size pilot installation on its Linux farm with Platform LSF software demonstration licenses on approximately 1,000 CPUs for several months to assess its behavior in production. To date, two million jobs have run through Platform LSF. Incyte is now implementing Platform LSF and Platform Analyzer on a Linux cluster of 1,200 CPUs for a dedicated production pipeline, where several million jobs will be handled automatically.

Platform LSF coexists well with Incyte's in-house technology, which enables researchers to use proprietary code while previously untapped CPU cycles are allocated to other jobs. Platform LSF's monitoring and control capabilities provide Incyte's managers with up-to-date information about the company's computing environment, and ensure fair sharing of resources through job queuing and prioritization. As new projects launch and throughput increases, Platform LSF accepts additional jobs and incorporates them into the existing framework of available computing power.

The 100 CPU Platform LSF cluster installed in the full-length (FL) sequencing production facility optimizes the use of computing resources by scientists to mine the data needed to build full-length genes. Harnessing and efficiently using computing resources is helping the FL team to meet tight production schedules and deliver a requisite number of full-length genes to customers on time.

The fully optimized allocation of resources and availability achieved by the Platform LSF software implementation saw a 50% increase in productivity in the first month after implementation. This optimized utilization is making a significant impact on time-to-market, enabling Incyte to launch new products faster than ever before.

Customer Site

Incyte Genomics (www.incyte.com)

Incyte is a drug discovery company that is using its expertise in genomics, medicinal chemistry and molecular, cellular and in vivo biology to discover and develop novel therapeutics. The company has three drug discovery programs underway focused primarily on the identification of new small molecule drugs for cancer and inflammation.



About Platform Computing

Platform Computing's intelligent, practical enterprise grid software solutions help organizations optimize IT resources to Accelerate Intelligence™. We plan, build, run and manage grids that link IT to core business objectives, and help our customers improve service levels, reduce costs and enhance business performance. With industry-leading partnerships and a strong commitment to standards, we are at the forefront of grid software development, propelling over 1,600 clients toward powerful insights that create real, tangible business value. For more please visit www.platform.com.



World Headquarters

Platform Computing Inc.
3760 14th Avenue
Markham, Ontario
L3R 3T7 Canada
Tel: 905 948 8448
Fax: 905 948 9975
Toll-free tel: 877 528 3676
info@platform.com

United States

Boston: 781 685 4966
Detroit: 248 359 7820
Reston: 703 251 4850
Newport Beach: 949 798 5654
New York: 212 672 1770
San Jose: 408 392 4900

Asia-Pacific

Beijing: +86 10 62381125
Hong Kong: +852 2869 5687
Tokyo: +81 3 5326 3105

Europe

Düsseldorf: +49 (0) 2102 610390
London: +44 (0) 1256 370500
Paris: +33 (0) 1 41 10 09 20

For more information, visit www.platform.com/contactus

Copyright © 2005 Platform Computing Corporation. ® TM Trademarks of Platform Computing Corporation.
All other logo and product names are the trademarks of their respective owners, errors and omissions excepted. Printed in Canada. Platform and Platform Computing refer to Platform Computing Inc. and each of its subsidiaries.