

IBM Triples Performance of World's Fastest Supercomputer

June 27 2007



IBM researcher Shawn Hall inspects a new Blue Gene/P supercomputer. The IBM system will be capable of up to three thousand trillion calculations per second.

IBM's newest supercomputer, Blue Gene/P, is nearly three times as fast as its predecessor and is designed to fit in smaller spaces and use less electricity than other commercially available models.

IBM today announced Blue Gene/P, the second generation of the world's most powerful supercomputer. Blue Gene/P nearly triples the performance of its predecessor, Blue Gene/L -- currently the world's fastest computer -- while remaining the most energy-efficient and space-saving computing package ever built.

The result is a machine that towers over other systems. It enables science and commercial supercomputing to attack vital problems in ways never before possible -- modeling an entire human organ to determine drug interactions, for example. Drug researchers could run simulated clinical trials on 27 million patients in one afternoon using just a sliver of the machine's full power.

The IBM System Blue Gene/P Solution scales to operate continuously at speeds exceeding one "petaflop" -- or one-quadrillion operations per second. The system is 100,000 times more powerful than a home PC and can process more operations in one second than the combined power of a stack of laptop computers nearly 1.5 miles high. The Blue Gene/P supercomputer can be configured to reach speeds in excess of three petaflops, a performance level that many thought unattainable only a few short years ago.

"Blue Gene/P marks the evolution of the most powerful supercomputing platform the world has ever known," said Dave Turek, vice president of deep computing, IBM. "A new group of commercial users will be able to take advantage of its new, simplified programming environment and unrivaled energy efficiency. We see commercial interest in the Blue Gene supercomputer developing now in energy and finance, for example. This is on course with an adoption cycle -- from government labs to leading enterprises -- that we've seen before in the high-performance computing market."

The Blue Gene supercomputer line was born from a visionary IBM initiative to develop a hugely scalable and highly reliable scientific computing platform. With Blue Gene, designers sidestepped two key constraints on state-of-the-art supercomputing -- power usage and space requirements. The Blue Gene supercomputer was purpose-built to fit in smaller spaces and use less electricity compared to other commercially available designs. Today, the Blue Gene/P supercomputer is at least

seven times more energy efficient than any other supercomputer.

The influence of the Blue Gene supercomputer's energy-efficient design and computing model -- once considered exotic -- can be seen everywhere in the industry where people have attempted to lower energy use and get performance without traditional reliance on chip frequency. The breakthrough BlueGene supercomputer design uses many small, low-power embedded chips each connected through five specialized networks inside the system.

Some of the world's leading research laboratories and universities have already placed orders for Blue Gene/P supercomputers. The U.S. Dept. of Energy's Argonne National Laboratory, Argonne, Ill., will deploy the first Blue Gene/P supercomputer in the U.S. beginning later this year. In Germany, the Max Planck Society and Forschungszentrum Julich also plan to begin installing Blue Gene/P systems in late 2007. Additional Blue Gene/P system rollouts are being planned by Stony Brook University and Brookhaven National Laboratory in Upton, N.Y., and the Science and Technology Facilities Council, Daresbury Laboratory in Cheshire, England.

"We view the installation of the Blue Gene/P system as the next phase of a strategic partnership furthering advances in computation in support of breakthrough science," said Robert Rosner, director, Argonne National Laboratory.

At FZ Julich, where researchers have been using a Blue Gene/L machine for two years, a Blue Gene/P system will allow for more breakthrough science -- in such areas as particle physics and nanotech, for example -- while keeping the research facility within acceptable power budgets.

"The big computing power at low electricity rates allows us to boost the performance of very complex and computationally intensive algorithms," said Thomas Lippert, director of the supercomputing center at FZ Julich.

Like its predecessor, the Blue Gene/P supercomputer is a modular design, composed of "racks" that can be added as requirements grow.

Four IBM (850 MHz) PowerPC 450 processors are integrated on a single Blue Gene/P chip. Each chip is capable of 13.6 billion operations per second. A two-foot-by-two-foot board containing 32 of these chips churns out 435 billion operations every second, making it more powerful than a typical, 40-node cluster based on two-core commodity processors. Thirty-two of the compact boards comprise the 6-foot-high racks. Each rack runs at 13.9 trillion operations per second, 1,300 times faster than today's fastest home PC.

The one-petaflop Blue Gene/P supercomputer configuration is a 294,912-processor, 72-rack system harnessed to a high-speed, optical network. The Blue Gene/P system can be scaled to an 884,736-processor, 216-rack cluster to achieve three-petaflop performance. A standard Blue Gene/P supercomputer configuration will house 4,096 processors per rack.

There are some key differences between Blue Gene/L and Blue Gene/P supercomputers. In hardware, the Blue Gene/P supercomputer moves to more (four vs two) and speedier (850 MHz vs 700 MHz) processors per chip; more memory and an SMP mode to support multi-threaded applications. This new SMP mode moves the Blue Gene/P system to a programming environment similar to that found in commercial clusters. The Blue Gene/P supercomputer dramatically scales up collective network performance to minimize common bottlenecks inherent in large parallel-computing systems. Software marks the third key upgrade for the Blue Gene/P solution -- system management, programming environment and applications support have all been refined in Blue Gene/P.

In Germany, a Blue Gene/P supercomputer will become the platform for

new applications scaled for petaflop-level performance at the Max Planck society. "The next-generation Blue Gene system will improve our capacity to prepare, develop and optimize applications from the Max Planck Society for future peta-scale computing," said Hermann Lederer, head of application support at Max Planck's RZG/Garching Computing Center.

The Blue Gene supercomputer operating system is based on the open-source Linux operating system. Applications are written in common languages such as Fortran, C and C++ using standards-based MPI communications protocols. The Blue Gene/P supercomputer is compatible with the diverse applications currently running on the Blue Gene/L supercomputer, including leading research in physics, chemistry, biology, aerospace, astrophysics, genetics, materials science, cosmology and seismology.

A variety of independent software vendors have plans to port existing tools and applications to the Blue Gene/P supercomputer. These include Gene Network Sciences, TotalView Technologies, Inc., Tsunami Development LLC and Visual Numerics, developers of IMSL.

Source: IBM

Citation: IBM Triples Performance of World's Fastest Supercomputer (2007, June 27) retrieved 25 April 2024 from <https://phys.org/news/2007-06-ibm-triples-world-fastest-supercomputer.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--