

Case Study.

Delivering High
Performance
Computing
to the Masses



Compete.

Council on
Competitiveness

Delivering High Performance Computing to the Masses

Microsoft is working with the NSF-funded National Center for Supercomputing Applications at the University of Illinois to develop high performance computing software products that can be used by hundreds of thousands of users. As high performance computing has moved out of the government laboratories and universities and is being adopted more broadly throughout the public and private sectors, the need for out-of-the box supercomputing software is growing rapidly. The supercomputer center and Microsoft engineers work as a team – engaging in a two-way knowledge transfer – that allows the software company to develop new software and tools to “deliver HPC to the masses.”

Like other highly successful companies, Microsoft has a succinct, simple business model – even when it comes to developing the complex software necessary for high performance computing (HPC).

As Kyril Faenov, the company’s general manager of high performance computing, says, “Microsoft is in the business of empowering the largest possible number of users by providing general purpose software.”

Microsoft Corp., headquartered in Redmond, WA, develops, manufactures, licenses and supports a wide range of software products, including the Microsoft Windows operating system and Microsoft Office suite of productivity software. Global annual revenues are around \$51 billion, and the company employs 79,000 employees in 102 countries. Unquestionably, it has been the most successful software company, making personal computers so easy to use that they are ubiquitous the world over. Could Microsoft apply those same skills and focused business model to the world of HPC, making HPC software similarly easy to use so that more companies would be encouraged to adopt the technology?

Coping with Complexity

There was a time when computers – huge mainframes that sat in glass-enclosed data centers – were the purview of the university and government research com-

munities and a small group of companies who had the resources and special talent required to operate them. Although the hardware has shrunk, HPC shares some of the characteristics that typified the early days of general purpose computing.

It is true that HPC is behind many of science’s most advanced discoveries. In addition, HPC has been used highly successfully in a number of commercial applications, such as automotive, aerospace, energy and financial modeling. But the fact is that even entry-level HPC systems are still difficult to use for people who are not a part of the scientific and technical computing ecosystem. As a result, large communities within both the public and private sectors only use desktop workstations and have complex problems that remain unsolved. And the relatively small group of companies that have successfully applied HPC technology to drive innovation and productivity often struggle to realize its full potential.

“Our industry, academic and government customers are primarily looking for turnkey HPC solutions,” Faenov says. “These organizations do not want to be concerned with the complexities of supercomputing. Instead, they would rather focus on helping make the unique and irreplaceable resources represented by their researchers and engineers more productive. These are the people

that ultimately come up with the creative breakthroughs that allow the company to be more competitive. Our goal is to help make these invaluable people as effective as possible by taking the complexity out of implementing and using supercomputing systems.”

Partnering with Experts to Accelerate Solutions

For Microsoft, lowering the barrier of HPC adoption means providing HPC software that can be used by hundreds of thousands of people to more effectively do their work. This could be a daunting task for a company that is not itself a customer of HPC systems in the traditional sense – Microsoft does not use supercomputers to design new drugs or test the safety of automobiles. This is where Microsoft’s relationship with the National Center for Supercomputing Applications (NCSA) comes in. NCSA, located at the University of Illinois, is one of several university supercomputing centers supported through the National Science Foundation to provide high performance computing capabilities to organizations across the country.

NCSA not only has one of the world’s top supercomputer systems, they have the experts who are able to solve some of the thorniest HPC problems and advance the technology in areas commercial vendors like Microsoft have yet to explore. The experts at the centers, Faenov says, have the engineering know-how, patience and foresight to be at the forefront of HPC and pioneer the exploration of new configurations. Faenov’s team members work side-by-side with them in an iterative developmental process.

The Microsoft engineers build the first version of the Windows-based HPC operating system. They then run the software on NCSA supercomputers to determine if

the code is in fact viable, discover where the problems are, and gain the opinions and advice of the center’s engineers. Then it’s back to Redmond for another round of coding, followed by a return to NCSA and more testing – and the cycle continues until the software is fully developed.

Working with an HPC center like NCSA and creating operating system software for its supercomputers is a very different experience from working with a data center that, for instance, handles the needs of millions of users accessing a Web site. Often, all that’s needed by the data center’s Web server are a few simple commands that respond to user queries such as “what items do you have in stock, how much do they cost, what colors do they come in, etc.” Supercomputer centers, however, tend to address one or two very complex problems at a time – for example, creating a model of an advanced airplane wing, exploring the microscopic world of protein folding or generating a detailed model of the sun’s corona; computations that may use all the supercomputer’s resources and run for days.

Comments Faenov, “Our collaboration with NCSA helps us to understand the HPC environment and the unique needs of the high-end HPC user. Microsoft then encapsulates that knowledge into replicable, standardized, commercially supported products that thousands of customers can use at one-tenth or one hundredth of the original operating cost. The value we add is not just to make these capabilities available to a broader range of users, but to lower the cost and complexity of the applications and make them available to organizations that don’t have multimillion dollar IT budgets. Partnering with NCSA and their experts is a critical part of this process.”

HPC Competitive Impact: Priceless New Knowledge to Create World-Class Products

One of the primary benefits that Microsoft realizes from its relationship with NCSA is a comprehensive understanding of the challenges associated with operating high performance computing systems.

“One of the things that we’re learning from working with NCSA is where the complexities lie,” Faenov explains. “How do users operate these high performance computing systems? What are some of the emerging technologies and applications that we need to be aware of – for example, the complex biology and chemistry applications associated with pharmaceutical investigations or the massive geologic simulations used by oil and gas companies? How do we manage these applications to ensure that the entire system doesn’t crash – or if it does, how do we recover?”

Armed with a new understanding of these challenges and complexities, Microsoft leverages its experience in crafting software systems to automate many of the human intensive jobs that can slow down HPC solutions. Included are tasks such as system administration, troubleshooting, managing the hundreds, even thousands of computers that power today’s data centers, submitting jobs, analyzing data, and packaging data for communicating with other researchers. “By collaborating with NCSA, we are realizing our goal of delivering cost effective HPC to the masses,” Faenov says.

He emphasizes that their work with the centers is a two-way street – knowledge transfer goes in both directions. For example, the Microsoft engineers not only learn about the complexities of HPC, but also about the more mundane but increasingly important daily realities

of operating supercomputer systems – such as power management and what it takes to keep these machines running on a 24/7 basis. In turn, the Microsoft engineers educate their colleagues at the center on innovations that the company is driving, such as the latest in developer tools, and database and systems management software that could help streamline center operations.

“I would be hard-pressed to put a numerical value on the dollars and time saved by our work with the NCSA and their staff,” Faenov concludes. “In reality, it’s priceless – we are learning from experts in the field as we work together to create world-class supercomputer software systems. I can’t imagine developing these high performance computing products without a hands-on understanding of the actual problems that need solving. Our collaboration with NCSA is a fundamental enabler that allows us to build the products that solve our users’ most pressing, challenging and exciting HPC issues.”

“I would be hard-pressed to put a numerical value on the dollars and time saved by our work with the NCSA and their staff... In reality, it’s priceless.”

Kyril Faenov, general manager, high performance computing, Microsoft Corp.

In Brief

Key Challenges

- Bring a new level of simplicity and productivity to high performance computing system implementation and use
- Improve overall productivity by developing tightly integrated infrastructure solutions with operating systems, middleware and tools
- Lower both the operating cost and total cost of ownership
- Advance widespread adoption of HPC throughout industry, academia and government by achieving the above

Solutions

- Enter into a joint relationship with the National Center for Supercomputing Applications at the University of Illinois that enlists the expertise of their faculty, staff and computational resources
- Use the results of the joint team's research to create high performance computing software that automates many of the human intensive jobs, such as system administration, troubleshooting, data analysis, and packaging in order to lower the overhead cost of supercomputing

Key HPC Benefits

- Allows Microsoft to collaborate with an ecosystem of experts to create new, easier to use, less costly high performance computing software, middleware and tools that can be readily adopted by a broader range of business, government and academic users
- Enables Microsoft to move toward its vision of "delivering HPC to the masses," to help them drive innovation and productivity

Web Site

- www.microsoft.com



Instead of using 100% virgin paper, we used paper that has been 30% Post-Consumer Recycled and made with 100% wind-generated electricity. We saved:

- 5 trees** preserved for the future
- 1667 gal** of water flow saved
- 276 lbs** of solid waste not generated
- 509 lbs** of greenhouse gasses prevented
- 3 million BTUs** of energy not consumed

Environmental impact statements were made using the Environmental Defense Fund Paper Calculator.

Council on Competitiveness
1500 K Street NW, Suite 850, Washington, D.C. 20005 T 202 682 4292
Compete.org



Compete.

Council on
Competitiveness