Success Story

NASA's Columbia Project

One Giant Leap

How NASA, SGI and Intel managed to build and deploy history's most powerful supercomputer in 15 blistering weeks

"The best word is adrenaline...the thought of building a 10,240-processor system in a little over three months was an exhilarating prospect."

– Dick Harkness, Manufacturing Vice President, SGI

No one at SGI's manufacturing facilities in Chippewa Falls, Wisconsin will forget the day they learned that SGI and Intel wanted to build history's most powerful supercomputer for NASA. And they wanted it done in fewer than 120 days.

"The thought of building a 10,240-processor system in a little over three months was an exhilarating prospect, especially since we still had to maintain our normal manufacturing pace," recalls Dick Harkness, SGI's vice president of manufacturing. "We honestly wondered what people were thinking."

What people were thinking, it turns out, was to spectacularly revitalize NASA's computing resources with a single system—one that would put more super-computing power into the Agency's hands than anyone, anywhere had ever seen before. But getting there quickly meant overcoming colossal challenges, from congressional approvals to the breakneck delivery and deployment of new products.

And yet it worked. NASA's "Columbia" supercomputer, so named to honor the crew lost in the 2003 shuttle accident, may have been born of necessity. But it was brought to life by NASA, SGI and Intel in a dramatic sprint to a finish line that at first seemed all but unreachable. Here's how it happened.

A Modest Proposal

As NASA's Advanced Supercomputing (NAS) Division Chief Walt Brooks tells it, the Columbia Project happened so swiftly only because several factors converged at the right time. In March, Brooks began meeting with NASA simulation and super-

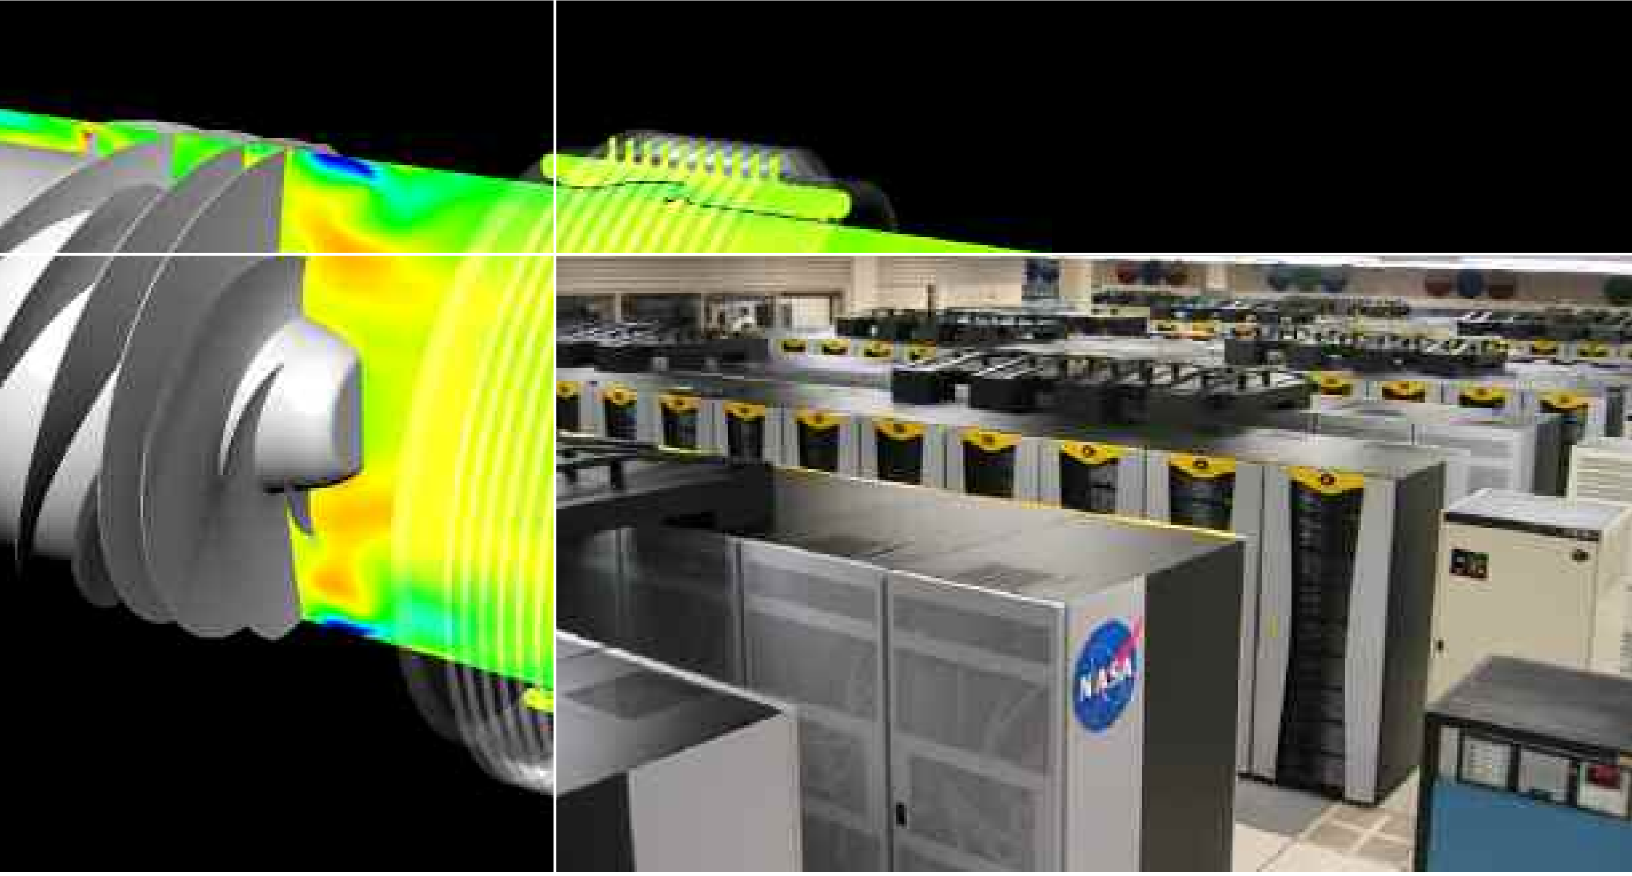

> “We’re trying to solve some of the toughest scientific problems in the world. We needed a system designed to efficiently execute the algorithms used in NASA’s premier science codes, rather than one that would merely do well on artificial benchmarks.”
>
> *– Jim Taft,*
> *Task Leader,*
> *NAS Division’s Terascale Applications Group,*
> *NASA*

computing experts from throughout the nation to develop a program to revitalize high-performance computing at the agency. The effort was guided by a report from the High End Computing Revitalization Task Force of the White House Office of Science and Technology Policy. “Even before the Columbia Project came on the horizon, the NAS team at Ames, along with other Agency high-end computing experts, toiled with the task force to define what it would take to bring the nation’s and NASA’s supercomputing resources on par with even our minimum current requirements,” recalls Brooks.

“What we proposed,” he says, “was a relatively modest investment to stay vital in high-end computing.” The idea: NASA would purchase 15 Teraflops (trillion operations per second) of computing power over three years. Brooks and his counterparts continued to sell their concept through the spring.

Meanwhile, the Agency considered more ambitious approaches. Particularly intriguing was the notion of building a world-class supercomputer by November. One idea was to link thousands of dual-processor commodity servers into a sprawling cluster, but NASA quickly dismissed that approach. “We’re trying to solve some of the toughest scientific problems in the world,” says Jim Taft, task lead for the NAS Division’s Terascale Applications Group. “We needed a system designed to efficiently execute the algorithms used in NASA’s premier science codes, rather than one that would merely do well on artificial benchmarks.”

Brooks and his team instead pointed to Kalpana, an Intel® Itanium® 2-based, 512-processor SGI® Altix® 3000 system in use at NASA Ames since November 2003 and named to honor Kalpana Chawla, a NASA scientist lost in the Columbia accident. In less than six months, Taft says, the Kalpana system - the first 512-processor Linux® system ever to operate under a single Linux kernel - had revolutionized the rate of scientific discovery at NASA for a number of disciplines. On NASA’s previous supercomputers, simulations showing five years worth of changes in ocean temperatures and sea levels were taking 12 months to model. But on the

“What we proposed was a relatively modest investment to stay vital in high-end computing.”

– Jim Taft,
NAS Chief,
NASA

SGI® Altix® system, scientists could simulate decades of ocean circulation in just days, while producing simulations in greater detail than ever before. And the time required to assess flight characteristics of an aircraft design, which involves thousands of complex calculations, dropped from years to a single day. “That kind of leap is incredible,” says Taft. “What took a year on the best computing technology previously available, we could now accomplish in days on the Altix system.”

NASA scientists began to imagine what an SGI® supercomputer built from 20 nodes, each with the power of Kalpana, could offer. “We could easily do all the benchmarking anyone could want,” Taft says, “but we’re more interested in a system capable of doing useful science.” The entire NASA team envisioned the science that would be possible on such a system: detailed hurricane predictions, global warming studies, electronic wind tunnel simulations, galaxy formation and supernova analysis, and experiments leading to safer space exploration.

Thirty Days to Yes

So the race was on. Once SGI and Intel determined they could meet a November installation deadline, Brooks and his team at NASA began a 30-day dash to change the agency’s operating plan and seek approval from NASA administrators, the Office of Management and Budget, Congress, and the White House.

Another crucial challenge was to prove NASA could pull it off without spending another dime over its approved budget. “Once Congress saw how we could acquire four times the computing resources for the same money,” says Brooks, “it was hard to refuse.”

In just 30 days, NASA received the green light on Project Columbia. “People both inside and outside the agency were inspired,” says Brooks. In light of Return to Flight initiatives following the loss of the Columbia shuttle and crew, the need was increasingly urgent. Lawmakers also took note of the impact the new supercomputer would have on other national science projects. “The Columbia Accident Investigation Board spent three months conducting analysis to seek the root cause of the accident. If they had this new system then, it would have been possible to do this in a matter of days.”

‘40 Days and 40 Nights’

Back in Chippewa Falls, SGI’s Dick Harkness and his team of 200 were ready. While NASA briefed Congress and the Office of Management and Budget on the Columbia concept, SGI’s manufacturing facility prepared workers to adapt to new processes. “Manufacturing flows were completely transformed to accommodate faster, more efficient builds. SGI’s factory personnel worked ‘40 days and 40 nights’ to meet production demands,” says Harkness. Assembly and QA of 512-processor Altix systems - until then a rare and involved event - quickly became a streamlined and easily repeatable manufacturing process.

Another challenge for SGI: Squeezing more than 10,000 processors into NASA’s supercomputing room in Mountain View, Calif., meant Columbia had to incorporate eight 512-processor nodes made new high-density, high-bandwidth version of the SGI Altix 3000 system.

“There simply wasn’t room on the floor for 20 traditional Altix nodes,” says Bill Thigpen, NASA’s Columbia project manager. “We needed eight nodes to be half the size of the original Altix 3000 systems for us to get all the hardware in the room.”

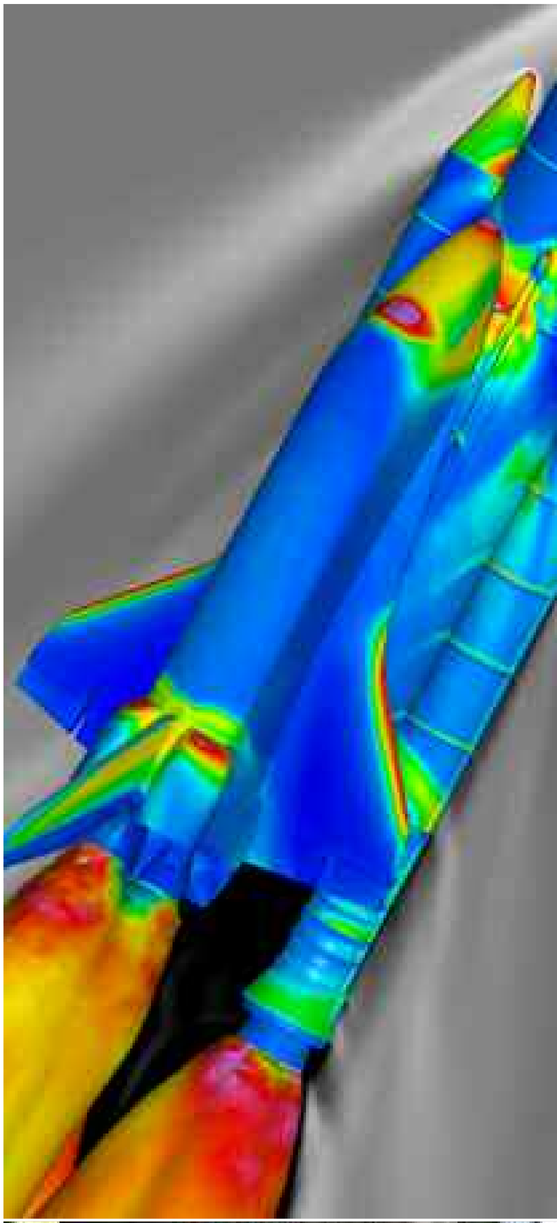

"This effort created a powerful national resource. This is a story about opportunity and drive and a willingness to stand up to the seemingly impossible–and make it happen. With the building of this great system completed, the work that will be performed will literally make our world and our universe safer for mankind. What could be more important than that?"

– Bob Bishop, CEO, SGI

For SGI, that spelled a challenge, since the Bx2 hadn't even achieved engineering release when Columbia plans were cemented. Indeed, based on typical parts delivery schedules, SGI only was to receive raw Bx2 parts by the time the finished systems were due at NASA.

But SGI's engineering and manufacturing group joined forces to deliver eight Altix Bx2 systems weeks ahead of schedule. The team met an even greater challenge, accelerating by four months the delivery of optional water-cooled doors—the first ever to be offered from other than a Cray product—that allow the denser Bx2 nodes to avoid overheating as they operate amid 12 air-cooled Altix nodes. "The water-cooled doors were crucial to this installation," recalls Thigpen. "This wouldn't have worked without them."

'Some Kind of Record'

The 19 new Altix nodes joined the Kalpana system at NASA Ames beginning in late June, and with them came a 440-terabyte SGI® InfiniteStorage solution to help NASA store and manage terabytes of new data generated every day.

For those on site, the rally continued, says Bill Thigpen. "Here we were, pulling out old systems and installing new ones, replumbing our water cooling system, and literally reconfiguring the floor on the fly, and meanwhile we had a large community of users who needed access to our systems every day."

According to NASA, the installation of the Altix nodes themselves was surprisingly easy. "It's phenomenal how quickly this combined team was bringing the nodes up and providing them to users to do real science," says Thigpen. "We had people from throughout NASA and several universities using the first installations within a week of having them hit the floor."

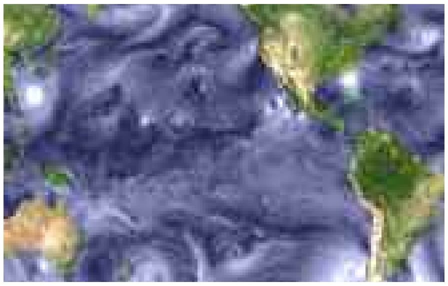

Jim Taft agrees. "In some cases, a new Altix was in production in as little as 48 hours. This is starkly different from implementations of architectures not based on the SGI architecture, which can take many months to bring to a reliable state and ready for science."

Japan's 5,120-processor Earth Simulator, for instance, wasn't fully usable for more than four years after inception of the project. "Imagine what you've lost in that time," says Taft, "not only in productivity, but in processor obsolescence as well. You're generations behind the curve before you even get started."

For those who drove the Columbia effort, however, the achievement symbolizes much more than mere Teraflops, or even supercomputing superiority. "This effort created a powerful national resource," notes SGI's CEO Bob Bishop. "This is a story about opportunity and drive and a willingness to stand up to the seemingly impossible - and make it happen. With the building of this great system completed, the work that will be performed will literally make our world and our universe safer for mankind. What could be more important than that?"

sgi

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.925.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3739 [12/15/2004] J14784