Reuters

Tue Nov 15, 2005 6:59 PM ET

Microsoft aims to enter supercomputer market

SEATTLE (Reuters) - Microsoft Corp. (MSFT.O: Quote, Profile, Research), which built its business by aiming to put a "computer on every desk" is moving into the high-end supercomputing market as hardware becomes more affordable, Chairman Bill Gates said on Tuesday.

The world's largest software maker will offer a version of Windows called "Compute Cluster Server 2003" that offers high performance computing by tying together several personal computers in a high-speed network that allows them to crunch huge amounts of data to simulate weather, analyze DNA and process other research-level tasks.

"Microsoft wants to play a role here," Gates told Supercomputing 2005, an industry gathering being held in Seattle this week.

Gates said that supercomputers, once considered massive mainframe machines that only researchers, a few companies and governments could afford, are becoming smaller and cheaper to the point where a clustered supercomputer system could be built for as little as $10,000.

The market for computer clusters built on Intel Corp.'s (INTC.O: Quote, Profile, Research) "x86" microprocessors used in more than 90 percent of PCs is growing at 15 to 20 percent per year, faster than the 11 percent growth in the PC market, and is expected to swell to as many as 300,000 installed machines by the end of the decade, Microsoft said.

To tap into that growth and compete against Linux -- the freely available operating system that is also being used to create clustered computer systems -- Microsoft is set to launch Computer Cluster Server in the first half of 2006. The latest beta, or test version, of the software was released on Tuesday.

Kyril Faenov, director of Microsoft's high-performance computing unit, said that the software will be designed to be easier to use, so that tasks can be displayed visually and systems can be configured more quickly.

Agence France Presse

Microsoft throws weight behind supercomputer software

15 November 2005

The new software represents Microsoft's first foray into the mainstream technical computing software market, a competitive arena that features low-cost, open-source products from companies such as Linux. Microsoft founder and chairman Bill Gates announced the release at a Supercomputing 2005 conference in the Washington city of Seattle and said the company is funding 10 joint research projects at academic centers worldwide.

Austin Business Journal

UT gets Microsoft research institute

November 15, 2005

Microsoft Corp. will establish a Microsoft Institute for High Performance Computing at the University of Texas and at nine other universities around the world, the software giant said Tuesday.

"This multiyear, multimillion-dollar investment builds on the relationship we started with the Cornell Theory Center years ago, and will allow us to understand much more deeply the requirements and opportunities for applying results of research and innovation in high-performance computing," says Bob Muglia, a senior vice president at Microsoft (NASDAQ: MSFT), based in Richmond, Wash.

Bill Gates, chairman and co-founder of Microsoft, unveiled the investment Tuesday at the Supercomputing 2005 conference in Seattle.

Specific dollar amounts for the project weren't provided. However, a Microsoft spokeswoman says the company has invested more than $100 million annually in academic programs over the past five years.

Aside from UT, the Microsoft institutes will be part of existing research centers at Cornell University, the University of Tennessee, the University of Utah, the University of Virginia, Russia's Nizhni Novgorod State University, China's Shanghai Jiao Tong University, Japan's Tokyo Institute of Technology, England's University of Southampton and Germany's University of Stuttgart.

Bio-IT World

SC05 Keynote: Gates on Reducing Time to Insight and Discovery

By Salvatore Salamone, Bio-IT World

November 15, 2005 | Microsoft Chairman and Chief Software Architect Bill Gates kicked off this year's SC05 conference by sharing his thoughts on how innovations in high-performance computing (HPC) will improve collaboration and speed research developments by connecting scientists, computing resources, and data.

The talk was before a standing-room-only crowd that filled three connected ballrooms in the Washington State Convention and Trade Center in Seattle. And while Gates used the talk to announce the beta 2 release of Windows Compute Cluster Server 2003, the focus of his keynote was on the changes he expected in how researchers work, as new ways to use computing resources become available.

Gates noted that several industry megatrends are coming together to alter the way research is conducted. These include increased processing power of HPC systems, service-oriented architectures that make it easier to share and access data, new computer form factors (Tablet PCs, PDAs, cell phones, etc.) that help collect data, and the digitalization of product design and simulation. "With computer modeling, we're figuring out what is happening in the real world," said Gates.

As these trends give scientists more data and more-complex data to work with, researchers "will need new software to figure out what resources are needed to work with the data," said Gates. One area where he thinks a lot more work must be done is visualization - to provide insight into and get more value out of data. And he believes leveraging the megatrends is a way to reach these goals.

Workflows Rule

At a conceptual level, Gates sees computational workflows that leverage the benefits of megatrends and integrate distinct tasks as the key to speeding insight and discovery.

This workflow includes computational modeling of the physical system being researched; acquisition of experimental data; searches and incorporation of persistent distributed data (such as libraries or public databases); and a software core that includes job scheduling, data mining, and analysis algorithms. All these elements produce information that would be interpreted by a researcher to get new insight into a problem.

While this approach is used in many organizations today, Gates thinks there are several opportunities to use newer technologies to improve the process. For instance, he believes personal supercomputers (systems with supercomputer power that cost less than $10,000, and that sit next to a researcher's desk) give researchers the ability to quickly develop and test new models.

Additionally he sees new resource management tools helping to provide a transparent way for

organizations to selectively use computing resources without the scientists having to be IT managers. For example, when a scientist submits an analysis job, the job could be assigned to run on a personal supercomputer or a larger cluster based on pre-defined workload criteria.

Focus on Life Sciences

Gates' talk focused on issues of importance to anyone working with large amounts of data, computational models, and research and development. But in many parts of the talk, Gates related how life scientists were using the methods he described.

For example, he called out Kyril Faenov, director of HPC at Microsoft, to demonstrate how to use these ideas in a research environment. Faenov's demo showed how a cancer research project was using an iterative data analysis technique to identify biomarkers and proteomics features in mass spectrometry data. In that demo, different analysis jobs were run on a local Linux cluster and a remote Windows HPC cluster.

Gates also spoke of an AIDS vaccine design project effort between Microsoft Research, the University of Washington, and the Fred Hutchinson Cancer Research Center. He noted the project was "the intersection of the advances in science and HPC."

He concluded noting that "the next 10 years will be really exciting, not just in [computing], but in science."

###

Business Week

Of Supercomputers, Microsoft and Sheryl Crow

By Peter Burrows
November 15, 2005

Want to know how bad Microsoft wants to reassert its role in the world of supercomputers? Although it just began participating in the 18-year-old SC/05 Supercomputer conference in 2003, Microsoft seems determined to make a splash at the research-oriented confab, which started today. For starters, chairman **Bill Gates is giving the keynote speech**. The company also has a prominent booth on the show floor--prominent enough to privately rankle some other vendors, who wonder whether the software giant unfairly threw its money around to win the coveted real estate (not true, says conference spokesperson Kathyrn Kelley; no other vendor had bid for the spot, she says, when Microsoft came along).

Why is Microsoft so interested? Because rather than an academic backwater, supercomputers and other high-performance computing technologies are increasingly being used by corporations to do sophisticated data-mining and run other business applications. And more often than not, designers of such systems are looking to Linux, when Windows might do the trick.

Gates isn't the only starpower Microsoft will aim at the assembled techies. The company, with help from Intel, will spend nearly $1 million on a party tonight at the swank **Experience Music Project** building, says one source with knowledge about the event. That includes more than $300,000 to pay Sheryl Crow--easily the biggest act to ever hit the SC/05 post-conference party circuit.

UPDATE: A spokesperson for Microsoft called after this post went up, to say that the bill for the party was well under $1 million, though he declined to provide a more accurate estimate. Hearing that, I tracked down another source that is familiar with the event, and this person also confirmed that the $1 million was significantly too high. Still, the general point stands. This was a top-shelf affair, by Supercomputer conference standards.

CNET News

Microsoft makes its way to 64-bit chips

By Martin LaMonica
Staff Writer, CNET News.com
Published: November 15, 2005, 11:51 AM PST

Microsoft said some upcoming products, including its Exchange 12 e-mail server, will run only on 64-bit processors.

At a conference for its management software customers, company executives detailed its plans to add support 64-bit microprocessors in its server applications and operating systems.

By late next year, Microsoft expects to deliver Exchange 12, which will run only on x86-compatible 64-bit servers, said Bob Kelly, general manager of infrastructure server marketing at Microsoft.

Kelly said 64-bit chips will make the greatest impact on the performance of applications such as Exchange and its SQL Server database.

"IT professionals will be able to consolidate the total number of servers running 64-bit (processors) and users will be able to have bigger mailbox size," he said.

Longhorn Server R2 and a small-business edition of Longhorn Server will be available only for x86-compatible 64-bit chips as well the company's Centro mid-market bundle. Longhorn server is expected to be released in 2007 and the R2 follow-up could come two years after that.

Without providing a specific target date, Kelly said that Microsoft is working on a product called System Center Essentials, which will be a management product aimed specifically at medium-size companies.

He said Microsoft intends to build application-level monitoring into the forthcoming version 3 of Microsoft Operations Manager to complement the present hardware-level monitoring.

Microsoft also said its Microsoft Virtual Server Release 2 will be available in the first month of December, priced at $99 per server for the standard edition and $199 for the enterprise version.

ComputerWorld

Gates outlines supercomputing vision

He sees a merging of 'mass computing' worldwide with supercomputers

News Story by Patrick Thibodeau

NOVEMBER 15, 2005 (COMPUTERWORLD) - SEATTLE -- Microsoft Corp. Chairman Bill Gates appeared before thousands of technical computing users today and forecast supercomputing systems that will cost less than $10,000 and a merging of "mass computing," or Windows PCs, with the world's most powerful systems.

Addressing the Supercomputing 2005 conference here, Gates spoke to a crowd that makes scant use of the Microsoft operating system in high-performance systems. Gates sought to cross that divide by stressing the common problems of what he called mass computing and supercomputing.

"Many of these challenges that we face in software -- connecting machines together, having parallel algorithms that allow many compute systems to work on a problem and combine their results together -- these problems are very similar to the problems that exist in high-end supercomputing," said Gates.

"Its exciting to think that we can get the best brains from supercomputing and from mass computing and bring those together and make great progress in the decade ahead," he said.

Windows barely registers a pulse in high-end computing. On the latest Top500 supercomputing list, released this week, the breakdown by operating systems has Linux running nearly 75% of the top 500 systems; Unix has 20%, and Mac OS X is listed at 1%. Windows isn't noted at all.

"It is astonishing how they missed out on this altogether," said John Abbott, an analyst at The 451 Group in New York, referring to Microsoft's limited role in high-end technical computing. It's important for Microsoft to be recognized in the high-end market "because that's where half these technologies, like grid and clustering, are all being proven," he said.

Gates said he believes that as chips reach gigahertz speed limits, the need for "parallelism" becomes more important. The environment he sketched out imagines desktop supercomputers linked to more powerful clusters in a heterogeneous environment.

"Microsoft wants to play a role here -- to be a participant and work with partners to see how our software fits in these solutions," said Gates. "These solutions will often be

Gates outlines supercomputing vision

extremely heterogeneous." Making certain that all these systems work together "is just one element on how software can do a better job," he added.

Gates also said that Microsoft is reaching out to more supercomputing centers to understand "what should we be doing with our software, how can it connect up to the other software that they have in a better way?"

Martin Gasthuber, a researcher at DESY (Deutsche Elektronen-Synchrotron), a high-energy physics laboratory for basic science research in Hamburg, Germany, was unimpressed. He said Gates' talk was about marketing Microsoft products. "For me, it's not a vision; it's the next step he wants to do -- which is coherent with the next generation of products he has in mind," said Gasthuber, a researcher at the facility who is working on storage issues.

But Ted Dodds, CIO for the University of British Columbia in Vancouver, said Gates offered "a very accurate description of current trends in technology and likely future possibilities."

Windows clusters are almost inevitable in many respects, said Dodds. This notion "of mass computing and very technical specializing computing teaming together, I think, is pretty evident," especially as researchers turn to low-cost commodity clusters to solve problems, he said.

"Without going into an open-source model, which I would never expect to see Microsoft necessarily do, they would develop products that will interoperate seamlessly" with the dominant open-source platforms, said Dodds.

William Kramer, general chairman of the Supercomputing conference and head of high-performance computing at the National Energy Research Scientific Computing Center at Lawrence Berkeley Laboratory, said Gates' appearance is an indication of the growing awareness of supercomputing's importance. "The output of [high-performance computing] activities are no longer hidden behind a curtain, if you will," he said.

Supercomputing is "being scaled down so more people can make use of these very complicated tools -- and I think that's one of the indications of Microsoft's interest here," Kramer said.

###

Good Morning Silicon Valley

Hi, I'm Clippy, your supercomputing assistant; you look like you're trying to model complex atmospheric phenomena

November 15, 2005

Figuring it's got plain old computing down pat, Microsoft is trying its hand at supercomputing. Bill Gates himself is introducing a test version of a new product, Windows Compute Cluster Server 2003, today at the SC\05, a supercomputing conference held this year in Seattle (get those Windows Cluster jokes out of your system right now). With the move, Microsoft is pushing into territory currently dominated by Linux powered machines run mainly in academic settings, and as John Oates at The Register notes, "Such academic and research units are usually staffed not by Linux enthusiasts but Linux obsessives. Academic ideals of peer-review and openness have further helped Linux gain ground. This is a new and not necessarily friendly market for Microsoft to join." But Microsoft seems less interested in the entrenched market and more interested in parlaying low cost (company engineers put together demo clusters for under $4,000 using store-bought hardware) and familiar administrative features to bring serious computing muscle to businesses. "What we're seeing is, as the price comes down and advanced applications trickle down from academia, they're really being picked up by enterprise [businesses]," said Kyril Faenov, Microsoft director of high-performance computing. At the event today, HP will announce plans to sell high-performance Windows systems alongside its Linux offerings, and Microsoft says it's talking to other manufacturers.

Meanwhile, the latest rankings of the world's fastest computers are out, and the big winner is IBM, holding the top three spots and five of the top nine.

ComputerWorld
InfoWorld

Platform steps up partnership with Microsoft

China Martens, IDG News Service
16/11/2005 08:06:27

Workflow scheduling software from Platfrom will work with Windows Compute Cluster Server when it ships next year

Grid computing software vendor Platform Computing announced support for Microsoft's upcoming cluster server software. The two companies are also developing an open scheduling standard, according to the head of Platform. Additionally, Platform is expanding its existing relationships with both IBM and Hewlett-Packard.

Platform is integrating its workload scheduling software LSF so it will be available when Microsoft's Windows Compute Cluster Server 2003 ships in the first half of 2006, according to Songnian Zhou, Platform's chief executive officer, chairman and co-founder. Platform made the announcement at the Supercomputing conference (SC05) taking place in Seattle through Friday.

Windows Compute Cluster Server 2003 brings together the Compute Cluster Edition version of Microsoft's Windows OS and the vendor's Compute Cluster Pack of interfaces, utilities and management software. Currently in beta testing, Windows Compute Cluster Server 2003 represents a major move by the software giant into the high performance computing market.

Platform and Microsoft have been working together for more than a year and the integration between their two products will be the first time the companies will "go out to market [together] publicly," Zhou said in a phone interview Monday. The integration will enable bi-directional job forwarding so that users can offload jobs to where there's optimal capacity, either on the computing cluster or in the data center, he added.

The vendors are also collaborating on an open standard for Web services job scheduling, according to Zhou. The standard is related to the integration of LSF and Windows Compute Cluster Server 2003 and the plan is to involve other vendors in the standard's creation, he said. "It's not just a proprietary approach," Zhou added. "This is one of the very few times when Microsoft is coming out of bat and working with a third party to drive a standard."

Platform steps up partnership with Microsoft

At the opposite end of the computing scale from Microsoft's low-level clustering, Platform is expanding its relationship with IBM to include LSF integration with Big

Blue's eServer Blue Gene cluster supercomputers, according to Zhou. IBM and Platform have been partners for about ten years, but have worked closer together over the past three years as Big Blue has focused more on a number of grid computing initiatives, he said. The integration of Blue Gene and LSF should be finished in the first quarter of next year, Zhou added.

Platform and IBM are making the integration move to anticipate user demand, according to Zhou. As more organizations bring in Blue Gene machines into their computing environment, they will want to use the same tools to manage the new supercomputer as they've been using with their existing high-performance computers, he said.

In another announcement at Supercomputing, Platform named HP as the third vendor to use its Rocks Standard Edition (SE) Linux cluster management software. The Rocks SE software will be available for the HP Cluster Platform 3000 later this month, according to Zhou. Both Dell and Silicon Graphics already include Rocks SE with their high-performance clusters.

Rocks SE is a commercial version of the NPACI Rocks open-source cluster management toolkit developed by the Rocks Cluster Group at the San Diego Supercomputer Center.

###

Information Week

Gates Outlines Microsoft's High-Performance Computing Plans Nov. 15, 2005

Second beta version of Windows Compute Cluster Server 2003 announced at supercomputing conference.

By Aaron Ricadela

Microsoft's entry into the scientific computing market could yield less expensive, easier-to-use supercomputing systems and pave the way for advances in business computing, chairman Bill Gates said during a keynote speech at a supercomputing conference in Seattle Tuesday.

Microsoft next year plans to release a version of Windows for small supercomputers, and is funding 10 universities in the United States, Europe, and Asia to help develop a set of high-performance computing products that includes operating systems, middleware, and development tools. As computer technology becomes integral to advances in biology, physics, medicine, and earth science, there's demand for simpler tools that can speed up scientists' "time to insight," Gates said during the speech at the SC05 conference.

High-performance computing techniques are also finding their way into industrial applications, including consumer product design and automobile crash testing. Microsoft's upcoming high-performance computing software, its first designed for the market, could help tie together the desktop and powerful clusters of computer servers in new ways, said Gates.

"Computation has become a tool for all the sciences," Gates said. "We need an approach that scales from the smallest supercomputer up to the largest."

Microsoft Tuesday released a second beta version of Windows Compute Cluster Server 2003, a version of its operating system designed to run on clusters of perhaps a few dozen machines. The product is due in the first half of next year, and will compete with the open-source Linux operating system that dominates cluster installations today. Using clustering technology, users can chain together dozens or hundreds of inexpensive PC servers with special cabling to run software that distributes work among the processors. The approach has let universities, supercomputing research centers, and corporate IT departments access supercomputing power at much lower costs than specialized supercomputing architectures have allowed. Microsoft is also designing new capabilities for its graphical development tools that could help scientists program clusters.

As part of its development effort, Microsoft has funded 10 "institutes for high-performance computing," including the universities of Washington, Virginia, Tennessee,

Utah, and Cornell University. Professors there will provide Microsoft with product development assistance as part of the new program. One challenge Microsoft and its partners will face is the need to develop new software algorithms that can run more efficiently across multiple processors. As chip designers become unable to significantly raise microprocessors' clock speeds to boost performance, the computer industry is turning to integrating multiple processing cores on a chip and extracting more parallel processing capability from software code to fuel performance gains. That means techniques used in high-performance computing today could in five or 10 years become commonplace on the desktop.

In an interview after his speech, Gates said Microsoft's research and development of high-performance computing technologies could help usher in speech and computer vision as user interfaces for PCs, improve the performance of software that can prioritize information for users, and run security algorithms more robustly. "We'll be able to use all the extra power," says Gates. "In a certain sense, you could say it's been a Holy Grail of computer science" to distribute code that runs well in one machine across many.

During Gates' speech, Microsoft high-performance computing director Kyril Faenov demonstrated a prototype "personal cluster" of four Windows computers running at 25 billion computations per second, on which departments or workgroups could run preliminary calculations before submitting them to a larger system for more detailed modeling of data. The prototype could be priced at less than $10,000, he said. During the demonstration, Faenov ran a version of The MathWorks' Matlab software on the personal cluster to analyze proteomics data, and then transferred the job to a 64-node cluster run by Intel Corp. at a different location in Washington.

Gates said that as part of its high-performance computing effort, Microsoft researchers are entering into new types of collaborations with scientists and medical researchers.

###

InfoStor

Isilon accelerates clustered storage

By Ann Silverthorn

November 15, 2005—Isilon's announcements at this week's SuperComputing 2005 conference included a partnership with Microsoft on clustered computing and clustered storage, as well as an acceleration product for clustered storage.

At the show today, Microsoft introduced the public beta of Windows Compute Cluster Server 2003 software, the company's first official entry into the realm of high-performance computing (HPC). The software is designed to bring HPC to departmental and workgroup levels. The product's core platform is Server 2003 with Ethernet and InfiniBand support. (For more information on InfiniBand-related announcements at SuperComputing 2005, see **InfiniBand makes a comeback at SuperComputing show**)

Also today, Isilon announced that it will team with Microsoft to demonstrate the use of Isilon's IQ clustered storage systems with Windows Compute Cluster Server 2003. The two companies demonstrated the role clustered computing and storage can play in accelerating digital workflow, fostering enterprise collaboration, and reducing IT and infrastructure costs.

Isilon's clustered storage systems include the IQ 1920, 3000, 4800, and 6000, which are powered by the OneFS distributed file system and offer an InfiniBand option with Cisco's InfiniBand Server Fabric Switches for use in clusters. Isilon offers the InfiniBand option at the same price as the Gigabit Ethernet version. About 90% of Isilon's customers have chosen the InfiniBand option, according to company officials.

Isilon also announced a new product to boost the performance of its clustered storage systems. The IQ Accelerator can scale to provide 6GBps of total aggregate throughput from a single file system. Using Isilon's Dynamic Performance Acceleration (DPA) software, administrators can install the IQ Accelerator without interrupting users.

The IQ Accelerator will be available in the first quarter of 2006. Pricing has not been finalized, but it is expected to cost about one-third of the company's existing nodes, which are priced at $7 per gigabyte for high-capacity configurations (IQ 4800 and 6000) and $12 per gigabyte for high-performance configurations (IQ 1920 and 3000).

NewsFactor Magazine Online
CIO Today
Sci-Tech Today
Top Tech News

Microsoft Updates Supercomputer OS

By Jay Wrolstad
November 15, 2005 10:06AM

"There has been a huge increase in high-performance computing, dominated by cluster systems, and Microsoft wants to gain a foothold," said Illuminata's Gordon Haff. "But it could be a tough sell for them because of the preference for Unix over Windows."

In an effort to gain a stronger foothold in the high-performance computing (HPC) realm, Microsoft has released the Beta 2 version of Windows Compute Cluster Server 2003.

The Beta 2 release includes all the features developed for the first beta of Windows Compute Cluster Server. It also features performance tuning for InfiniBand and gigabit Ethernet networks. The final version of the operating system will be released early next year.

For mainstream customers to take advantage of HPC systems, these computers must be easier to deploy, use, and manage, Microsoft contends.

Microsoft is therefore collaborating with some 20 software partners across several industries to create a range of commercial applications for Windows Compute Cluster Server 2003.

The software giant also is collaborating with the academic community. Along with University of Stuttgart in Germany, the University of Southampton in the UK, and Nizhni Novgorod State University in Russia, Redmond plans to help establish 10 institutes for high-performance computing.

Unix, Linux Dominate

Although the supercomputer race is a feisty one, even with companies like Apple trying to garner space in research labs and universities, Microsoft has not been in the running for a top spot.
Traditionally, supercomputers have been Unix-based. More recently, Linux-based machines have begun to take on this role at research institutions.

“There has been a huge increase in high-performance computing, dominated by cluster systems, and Microsoft wants to gain a foothold,” said Illuminata’s Gordon Haff. “But it could be a tough sell for them because of the preference for Unix over Windows.”

Microsoft’s ultimate goal is not necessarily the research labs. The company has noted that it is interested in providing enterprise customers with the benefits of high-performance computing, especially in light of the trend toward commercializing the grid-computing strategies developed in the labs.

Enterprise Eyes Clusters

The advantage of cluster-based supercomputing systems is that they typically use off-the-shelf components, making them less expensive than custom-built, standalone supercomputers.

Haff suggested that Microsoft could compete in this market by offering better prices, and noted that the company already has a well-established set of development tools at its disposal to court researchers.

Supercomputers currently are used primarily for simulation applications and other scientific tasks, although these days, the analyst noted, they are taking on more I.T.-related chores, such as data mining and credit-risk analysis.

Microsoft sees Windows Compute Cluster Server 2003 as an OS that makes high-end computation easily accessible to I.T. workers and easily integrated into the enterprise infrastructure

###

Puget Sound Business Journal

Microsoft plans university-based research institutes

November 15, 2005

Microsoft Corp. will establish a Microsoft Institute for High Performance Computing at 10 universities around the world, the software giant said Tuesday.

Bill Gates, chairman and co-founder of Redmond-based Microsoft (NASDAQ: MSFT), unveiled the investment Tuesday at the Supercomputing 2005 conference in Seattle.

Specific dollar amounts for the project weren't provided. However, a Microsoft spokeswoman says the company has invested more than $100 million annually in academic programs over the past five years.

The Microsoft institutes will be part of existing research centers at the University of Texas, Cornell University, the University of Tennessee, the University of Utah, the University of Virginia, Russia's Nizhni Novgorod State University, China's Shanghai Jiao Tong University, Japan's Tokyo Institute of Technology, England's University of Southampton and Germany's University of Stuttgart.

The Register

HP gets visual with workstation clusters

Blades, chips and software too

By Ashlee Vance in Seattle
Published Tuesday 15th November 2005 19:09 GMT

SC05 HP continues to refine its server cluster arsenal, announcing this week a new visualization package, more software and fresh processors for its pre-packaged systems.

The HP Scalable Visualization Array (SVA) combines numerous workstations, Linux and graphics cards to give customers a graphics powerhouse. The new system will be used to handle clustered parallel visualization software, display high-resolution images and perform real-time rendering. HP will slot the SVA product into its Unified Cluster Portfolio line that includes systems which hold more than 1,000 processors.

With SVA, HP will support Version 3.0 of the XC System Software for managing clusters and application workloads. In addition, it will ship the Scalable Visualization Array software for configuring the systems and sending jobs to the boxes. The HP SFS and Remote Graphics packages can be purchased as options.

Customers will also find Version 2.4 of the HP ClusterPack for HP-UX available. The fresh release should make Unix clusters easier to manage by allowing new systems to be added to a cluster with minimal overall disruption. The software now works with HP Systems Insight Manager.

On the hardware front, HP has added Intel's dual-core "Paxville" version of Xeon as an option with its pre-packaged clusters. Customers could already buy systems with the dual-core Opteron chip from AMD. In addition, customers can now tap HP's full line of blade servers with the clusters.

HP has also made it easier for customers to buy its cluster boxes. You can now cruise to HP's web site and configure systems with up to 32 servers, varying processors, memory, interconnects and software. HP will then process the order and deliver a quote to the user.

Last but not least, HP vowed to support Microsoft's upcoming Windows Compute Cluster Server 2003 operating system, which was also discussed today at the Supercomputing conference here. In the first half of next year, HP will ship the OS an option with its Cluster Platform 3000 and 4000 systems and the Platform 4000 blade boxes. HP also plans to port its message passing interface (HP-MPI) to Microsoft's specialized version of Windows. ®

The Register

Gates promises a supercomputer under every desk

By Ashlee Vance in Seattle
15 Nov 2005 20:48

Er, unless security is a problem

SC05 Not satisfied with owning your PC, Microsoft would like to own your personal cluster too. So said Chairman Bill Gates today at the Supercomputing 2005 conference here in Seattle, where he laid out a vision that includes inexpensive super-powered machines available to average users - not just government labs and universities.

"What we see as a key trend here is that we will have supercomputers of all sizes, including ones that will cost less than $10,000 and be able to sit at your desk or in a department," Gates said.

A rather humble Gates admitted that Unix and Linux dominate the supercomputing landscape, a fact that only Redmondian logic could deny. To that end, Microsoft plans to create Windows cluster software that can work in tandem with Linux-based kit. Microsoft showed a demo where a Windows cluster and Linux cluster cranked away on calculations at the same time and delivered up a shared set of data to the end user. (What kind of demo? Cancer research of course. How subtle.)

"It's all about heterogeneity," Gates said.

While some reports suggeseted that Gates talked up a new, publicly available beta of its Windows Compute Cluster Server 2003 operating system, Gates mentioned the product's name but gave no more information. The OS is due in production form in the first half of next year. Microsoft, however, did reveal the beta 2 news in a press release.

Microsoft also opened up about its cluster OS at an event in Spain.

Overall, Gates' message proved a bit boring. He mentioned that supercomputers do important things and will be more important in the coming years because we face important issues. Microsoft, of course, will have a tough time cracking into this important work since Linux is the OS of choice for the supercomputer crowd.

Gates does seem to be on the right track though as far as the personal clusters are concerned. Orion Multisystems already makes a super-charged 96-processor deskside cluster than can plug into a standard wall outlet. Researchers have been seeking such machines which will let them crank away at projects right in their labs and transfer bigger jobs to the giant clusters that require special cooling systems and have higher power requirements.

"We need an approach here that scales from the smallest supercomputer up to the very largest," Gates said.

Microsoft plans to tie its other productivity, database and workflow applications into the cluster OS over time. This would make it possible for researchers to set up projects and then easily see relevant information about the data they were collecting and how other scientists were using the information. Apparently, RSS feeds are very important in this vision, which gives you some idea of how Microsoft may fail in the supercomputing space.

The 9,000-person Supercomputing crowd appeared enthralled with Gates and peppered him with probing questions following the keynote. Gates was forced to address his opinions on global warming and supercomputing security.

"I am not really an expert on the energy problem," Gates said to the first question.

"Security is one of the few things that, if we don't do it right, could take this vision I have talked about and really hold it back," Gates said.

Hold on to your hats, friends. Your DNA could be one worm away from entering the global computing grid. But at least you'll have a really fast PC for playing Doom 78. ®

###

The Register

Windows Cluster goes after Linux heartland
There'll be blood on the dancefloor...

By John Oates at Microsoft's IT Forum in Barcelona
Published Tuesday 15th November 2005 13:55 GMT

Microsoft's entry into the world of supercomputers will help push the technology beyond government and academic departments and towards business users, the company's head of server and tools claimed today.

Bob Muglia, senior vice president of server and tools for Microsoft, announced the public availability of the beta version of Windows Compute Cluster 2003 at the vendor's conference in Barcelona today, saying it will take on Linux and Unix in their traditional homeland - very high-end machines and groups of machines. The company is working with more than 20 companies to create applications to run on Windows Cluster.

Speaking after his keynote speech, Muglia said: "This is for any workload which needs high-power computing. We are seeing a transition from government and academic use to a broader market - to bring it into the mainstream. The sweet spot is not for really big machines but in the range of 4 to 64 way machines."

Muglia said the main markets targeted would be oil and gas, pharmaceuticals, manufacturing, financial services and universities. Such machines run intensive tasks like modelling

Speaking without apparent irony Muglia said: "We've spent some time talking to Independent Software Vendors recently and the software community welcomes the arrival of a consistent environment to this area."

He might be right but he might also have a fight on his hands - such academic and research units are usually staffed not by Linux enthusiasts but Linux obsessives. Academic ideals of peer-review and openess have further helped Linux gain ground. This is a new and not necessarily friendly market for Microsoft to join.

To support its move Microsoft announced "a multiyear, multimillion-dollar investment in the academic community". It is bankrolling the establishment of ten institutes for High-Performance Computing with universities including Stuttgart(Germany), Southampton(UK) and Nizhini Novgorod State University(Russia).

Is the arrival of Windows really what cluster computing needs? Let us know what you think at the usual address.®

TechNewsWorld
LinuxInsider

Microsoft Continues Push Into Supercomputing Realm

By Jennifer LeClaire
11/15/05 2:10 PM PT

"Most sciences are becoming computational sciences, which is why advanced computing capabilities need to be seamlessly integrated into the end-to-end scientific process," said Bill Gates. "We see many opportunities to collaborate with the scientific community on innovative solutions that will accelerate the pace of insight and discovery."

Microsoft (Nasdaq: MSFT) Chairman Bill Gates today shared his vision of how the software industry can contribute to accelerating scientific research and engineering innovation during his keynote speech at Supercomputing 2005.

Of course, part of his vision includes the software giant's Windows Compute Cluster Server 2003, Microsoft's first product offering for high-performance computing. With the delivery of Microsoft Windows Compute Cluster Server 2003 Beta 2 today and general product availability scheduled for the first half of 2006, Microsoft said it will offer an HPC platform that accelerates the time to insight.

HPC Investment

Gates also announced an investment in 10 Institutes for High-Performance Computing (HPC) worldwide. Microsoft said this multiyear, multimillion-dollar investment in joint research projects at some of the world's leading research institutes will help guide the company's ongoing software research and product innovation to address the most challenging technical computing problems.

"Technical computing is crucial to the many discoveries that impact our quality of life -- from making safer, more efficient cars and airplanes to addressing global health issues and environmental changes," Gates said.

"Moreover, most sciences are becoming computational sciences, which is why advanced computing capabilities need to be seamlessly integrated into the end-to-end scientific process. We see many opportunities to collaborate with the scientific community on innovative solutions that will accelerate the pace of insight and discovery."

Integrating With Existing Tools

To help ensure simple integration of Windows Compute Cluster Server 2003 with a broad set of applications and tools, Microsoft said it is working closely with software, application and hardware partners across the industry.

Microsoft is being joined in its booth at Supercomputing 2005 by 17 software partners and a number of hardware partners, including Dell (Nasdaq: DELL), HP, IBM (NYSE: IBM), NEC (Nasdaq: NIPNY) Corp., Orion Multisystems Inc., Tyan Computer Corp. and Verari Systems Inc., to demonstrate 20 industry-specific HPC applications.

As part of this effort, Microsoft said it is working with Intel (Nasdaq: INTC) to provide software vendors with a robust set of resources -- including 64-bit hardware, software and tools from both companies, and joint on-site engineering assistance -- to enable broader application support on Windows Compute Cluster Server 2003 when it is released to the market.

Microsoft is also working with software partners across the manufacturing, life sciences and geosciences, and other industries to make a robust set of commercial applications available on Microsoft Windows Compute Cluster Server 2003.

Money-Making Opportunity

Chris Alliegro, Lead Analyst at Directions on Microsoft, told TechNewsWorld that there is a clear financial opportunity for Microsoft in the super computing arena. The company, he said, has core foundation technologies that, if positioned correctly, could be of significant interest to the scientific community.

"It's more than just an opportunity to modify the Windows operating system to support a fairly niche computing architecture," Alliegro said. "It's an opportunity to leverage its investment in the operating system, in data formats, XML, database, data aggregation, data analysis, and put together a compelling story for the high performance computing universe."

Should the Linux camp be worried? Alliegro thinks so. That, he said, is because Microsoft paid respect to the fact that the super computing environment is heterogeneous and the company is going to have to interoperate with Linux machines.

"Microsoft acknowledged the fact that the scientific, academic and research community to a large extent are leaning toward low-cost solutions like Linux," Alliegro said. "So Microsoft knows it can't come in with a pitch for its proprietary operating system and its proprietary protocols versus Linux. They know that would turn a lot of people off."

###

Washington Technology

Newsbytes

Microsoft eyes supercomputing market through software

By Joab Jackson
11/15/05

With today's beta release of Windows Compute Cluster Server, Microsoft Corp. is hoping to enter the high-performance computer market—or at least the lower reaches of that market.

The Redmond, Wash.-based company is pitching the operating system to engineers and researchers "who would like some power, but don't want to have to take their computer jobs to a supercomputer center," said John Borozan, group product manager for the Windows server division. By running this operating system on groups of servers, offices and departments could run small supercomputing clusters.

A departmental implementation of a high-performance computer (HPC) cluster may run about $50,000 to $250,000, Microsoft estimates. The cluster sizes could range from four to 64 nodes. According to Borozan, a typical four-node system may yield about 10 billion floating-point operations per second (GFLOPS) of computational power, depending on the hardware used.

The Windows Cluster Server is based on the Windows Server 2003 server software, but has additional features that allow computers to be yoked together to work in parallel on computationally intensive tasks.

The package includes an open-source message passing interface (MPI) library for passing data among multiple machines. The MPI library, called MPI Chameleon or MPICH2, was developed by Argonne National Laboratory. Microsoft added performance and security enhancements.

The server software will also come with cluster management software and a job scheduler to coordinate workloads.

Traditionally, Unix and Linux-based machines have dominated the supercomputer cluster space. As such, Microsoft is working with independent software vendors and universities to generate more HPC tools for its own platform. A number of HPC software vendors have ported their applications over to this HPC version of Windows, including MatLab from The MathWorks Inc., and Mathematica from Wolfram Research Inc.

Microsoft will offer a version of its Excel spreadsheet that can be run across a multi-node cluster, according to Borozan. Users can also build their own applications using

Microsoft Visual Studio and the .NET framework.

Beyond the .NET languages, Microsoft is also working with third parties to build compilers for other languages favored in the HPC space, such as Fortran.

To pull data from Unix file systems, the Compute Cluster Server comes with a copy of Microsoft Services for Network File System, which allows Microsoft Windows machines to access NFS-based drives.

On the hardware side of things, Windows Cluster Server can run on x86 64-bit microprocessors—specifically Intel's EM64T Xeon and Pentium chips and AMD's Athlon and Opteron. It will not run on Intel Itanium processors, Borozan said. Out of the box, the software will work Gigabit Ethernet and Infiniband interconnects.

The full production release of the server software is scheduled for the first half of 2006.

###

Seattle Post-Intelligencer
Hispanic Business

Windows Gets Super Powers

Todd Bishop
November 15, 2005

Microsoft Corp. is moving into the world of high-performance computing with a new version of Windows that aims to make the ultrapowerful systems far more common in businesses and research institutions.

This morning at SC/05, an international supercomputing conference in downtown Seattle, Bill Gates will lay out Microsoft's strategy for entering the high-performance computing market with a program called Windows Compute Cluster Server 2003.

The program increases the competition between Windows and open-source rival Linux, which currently leads the high-performance computing market. But rather than simply trying to take market share from Linux, Microsoft says it will try to expand its potential market by increasing the use of high-performance computing overall.

"It should just be available to a lot more people," said Craig Mundie, Microsoft senior vice president and chief technical officer, describing the company's strategy in an interview.

High-performance computers are fast and powerful machines used by researchers and companies to work with huge amounts of data. The fastest are known as supercomputers.

Microsoft's Windows Compute Cluster Server is designed to run integrated collections of computers operating together as high-performance systems. The company will release a second test version of the program today. A final version is due out in the first half of next year.

Cluster computing makes up about half the high-performance computing market and is "clearly one of the higher-growth areas," said Bruce Toal, marketing and solutions director for Hewlett-Packard's high-performance and technical computing group.

One of Microsoft's goals, Mundie said, is to make such systems easier to operate by offering an integrated package that includes and connects to many of the same administrative features that many companies already use to run Windows PCs and computer servers. Examples include security, automatic updates and user authentication.

Microsoft says it envisions companies using its high-performance computing software in such areas as manufacturing, geosciences, life sciences, oil and gas and financial services.

At the conference today, Hewlett-Packard plans to announce that it will offer high-performance computing systems running Windows as an option for customers alongside Linux-based systems. Microsoft says it is also in talks with other major manufacturers of high-performance computing systems.

Microsoft also is showing a series of major high-performance computing applications that already have been adapted for the new Windows version.

But a separate announcement by HP underscores the position of Linux as a major force in the market. The company will base a new high-performance computing product, the HP Scalable Visualization Array, on the open-source operating system.

Linux is "one of the things that Microsoft has to contend with more in high-performance computing, because it's so well-entrenched," said analyst Carl Claunch, a Gartner Inc. research vice president. However, he said, Microsoft will benefit from the fact that many of its target customers already are running other Windows systems.

Microsoft hasn't said how much it expects to charge for Windows Compute Cluster Server. The Linux operating system is available for free. As with its main Windows Server software, however, Microsoft says it is aiming for a lower "total cost of ownership," taking into account the cost of such things as operation and maintenance.

About 100 people at the company are working on Windows Compute Cluster Server 2003. But that number could grow, said Zane Adam, a director in the Windows Server division. "We'll invest what's needed to make this happen," Adam said.

Microsoft's Mundie led a supercomputing company called Alliant Computer Systems before joining Microsoft in the early 1990s.

About two years ago, Mundie said, he began encouraging Gates to move the company into high-performance computing, based on what he was seeing in the market.

High-performance computing systems, once confined to high-end research labs and large businesses, have become less expensive and more common in recent years, as computer hardware costs have dropped.

And although proprietary systems still exist, more are being built on the standard AMD and Intel microprocessors that Microsoft is accustomed to working with.

Microsoft cites data from the IDC research firm showing a 70 percent increase in high-performance systems last year.

"The volumes have gone up in absolute terms to the point where it's an interesting business," Mundie said. "Many of the attributes that have been a part of Microsoft's historical strength, in distribution and support, now have relevancy to this market."

Kyril Faenov, Microsoft's director of high-performance computing, compared the current growth in the field to past trends in markets where Microsoft has experienced success.

"We're going through an inflection point in this market very similar to what business and personal computing have gone through over the last couple of decades," Faenov said.

Microsoft's challenge, he said, is now to make high-performance computing "digestible and usable" by the much larger number of users that can now afford it.

###

CNET
Builder.com Australia
New York Times
ZDNet Australia

Gates touts Microsoft's supercomputing move

Ina Fried, CNET News.com
November 16, 2005

Microsoft Chairman Bill Gates on Tuesday announced the company's foray into the world of supercomputing, though its first operating system for computer clusters remains in beta testing.

Speaking at a supercomputing conference in Seattle, Gates announced that the company has reached the Beta 2 stage for its Windows Compute Cluster Server 2003. The product consists of both a cluster-optimised version of Windows Server 2003 as well as software for job scheduling and other tasks. It is scheduled for release in final form in the first half of next year.

"Technical computing is crucial to the many discoveries that impact our quality of life -- from making safer, more efficient cars and airplanes to addressing global health issues and environmental changes," Gates said in a statement. "Moreover, most sciences are becoming computational sciences, which is why advanced computing capabilities need to be seamlessly integrated into the end-to-end scientific process."

Separately, Microsoft also announced that the Compute Cluster Server and several other upcoming server software releases will work only with 64-bit processors. Such chips, which include Intel's 64-bit Xeons and Advanced Micro Devices' Opteron, are becoming the norm on servers, and 64-bit processors are making their way onto desktop machines.

"We are making big bets on 64-bit technology and working closely with our industry partners to enable a smooth transition for customers, so they can begin to realise the benefits of mainstream 64-bit computing," Microsoft Senior Vice President Bob Muglia said in a statement.

Other titles that will be only 64-bit include the next version of Exchange, as well as the upcoming "Centro" midmarket server and the Longhorn version of Small Business Server. Microsoft is planning a version of Longhorn Server that will work on 32-bit chips, though it expects the first major update to Longhorn Server, Longhorn Server R2, to be exclusively 64-bit capable.

Academic institutions and some industrial customers have been combining clusters of standard Intel- and AMD-based servers for some time. But Microsoft says it has seen a shift where such products expand beyond a niche market into more and more businesses.

Microsoft is pitching its tools as on par with the performance of Linux. The company also claims its tools are easier to manage and integrate with the rest of a corporate computing environment.

"HPC (high performance computing) is starting to broaden out," said Kyril Faenov, director of Microsoft's HPC unit. "What that leads to is demand, on behalf of customers, to really provide this raw power in a way that is easier to consume and easier to integrate into what they are already doing."

ZDNet Australia's sister site CNET News.com reported in May 2004 that Microsoft was planning a high-performance computing version of Windows, a move later confirmed by Microsoft.

In March, a Microsoft engineer said the company hoped to have a product out by autumn in the US, though that turned out to be a beta. The new beta version will be public, unlike the one released at September's Professional Developer Conference in Los Angeles, which was limited to about 1,600 testers.

In addition to announcing the new beta, Microsoft is touting the support it is receiving from hardware makers. The company will outline 19 key applications, from other software makers, which will run on the new version of Windows.

And, Microsoft also announced investments in 10 high-performance computing institutes that will serve as early customers and help the company determine where to go with the cluster-software effort.

###

DantaMonitor ComputerWire

Linux Clusters Continue to Expand in Top 500 Supers Ranking

By Timothy Prickett Morgan
November 15, 2005

The number of parallel supercomputing clusters and the aggregate performance of these machines both continued their expansion in the latest Top 500 supercomputer rankings.

The Top 500 list was announced at the SuperComputing 2005 tradeshow in Seattle, which is not exactly a coincidence. Microsoft Corp, having seen Linux take off as the core platform on cluster supercomputers in the past five years, will debut its own Windows HPC variant at SC2005 this week. The question now is: Can Windows storm the Linux stronghold?

If the Top 500 is any indication, it would seem not. In a supercomputing environment, support for Fortran, C, and sometimes Java is critical, and so is the reliability of the core operating system platform deployed on a cluster. Excepting the very high-end, exotic machines that tend to dominate the Top 500 rankings, which are basically hand-tooled, very expensive supercomputers that only government agencies and academic research institutions can afford, most big HPC server buyers are interested in using as many commodity components as possible.

That means X86 and now X64 processors and usually a Linux operating system. Windows can run on the same iron, of course, which is why it has a much better shot than Unix. But all of the compilers, workload management, job scheduling, cluster management, and resiliency software that has been either ported from Unix to Linux or created out of thin air by the open source HPC community has to be moved to Windows for Windows to have a chance in the HPC market.

Bill Gates, Microsoft's founder, chief software architect, and a former compiler maker in his own right, might be a highly respected individual in the IT market, but research labs are going to expect more from a Windows HPC variant than they get on their desktops. It will be very interesting to see what penetration Microsoft gets in the Top 500 list in the next two years. There is certainly potential for Microsoft to win a lot of deals, as Linux distributors have done.

But for now, the Top 500 is about Linux and Unix, and the $7-plus billion HPC server market is one where Microsoft has almost no presence except for a few token accounts. The current Top 500 list makes this very clear. For instance, there is not a single Windows machine on the list, even though there are a few experimental Windows clusters out there in the world. The reason is simple: to get on the Top 500 list, the ante is

now 1.64 teraflops, which is almost double the performance of the 500th system on the list.

If Microsoft wants Windows to get a toe-hold in the Top 500 list in the coming years, it is going to have to find some very big institutions to use its software. Eating its own dog food is probably the best place to start, and it seems likely that Microsoft Research, the software giant's research arm, will probably start moving up the Top 500 list after the company installs its own cluster; it is possible that Microsoft pumps the money for a giant Windows cluster into the Cornell Theory Center or some other academic institution as well.

Right now, Linux rules the Top 500, after entering the list in 1998. In the latest Top 500 rankings, Linux machines account for 372 out of 500 machines (74 percent), 1.19 petaflops, or thousands of teraflops, out of a total of 2.3 petaflops of aggregate sustained computing capacity (52 percent), and 333,373 processors out of a total of 731,069 (46 percent) across those 500 machines. Various Unixes (including AIX, Solaris, HP-UX, the open source BSDs, and the BSD-derived Mac OS X) account for 109 machines, and the remaining 19 machines on the list are mixed clusters that use both Linux and Unix.

By processor architecture, there are only 14 vector machines in the whole Top 500, which shows how far clusters based on scalar, commodity processors have come in the past 15 years. In June 1993, there were

334 vector machines on this list accounting for about a third of the aggregate 1.1 teraflops of computing power in the Top 500 list (the first one ever published, in fact), and the 131 other machines were exotic parallel clusters built by companies such as Kendall Square, Thinking Machines, MasPar, and a bunch of smaller vendors that no longer exist. And tellingly, IBM Corp did not have a single machine on the list. Most of these machines had proprietary interconnects, proprietary processors, and a basically proprietary implementation of Unix.

Fast forward to November 2005's list, and most of the machines are built on commodity Intel Corp or AMD processors: There are 206 machines built from 32-bit Xeons, 46 built from Itaniums, 81 built from 64-bit Xeons, and 55 built from Opterons; IBM's Power family of processors have a nice showing, with 73 machines.

The dominant interconnection method is off-the-shelf Gigabit Ethernet, with 249 machines in the November 2005 list, followed by 70 machines using Myricom's Myrinet interconnect, 26 using InfiniBand, 42 using one or another variant of IBM's SP switching technology from its pSeries Unix family, 19 using Silicon Graphics Inc's NUMALink, 31 using Hewlett-Packard Co's HyperPlex.

There's a mix of crossbars and other interconnects in the list, but clearly people are perfectly happy to build X64 and GigE clusters, and are doing so with wild enthusiasm.

On the current list, 360 of the 500 machines are clusters, with 104 being classified as massively parallel processors (which means they have a high-bandwidth interconnect that more tightly couples the machines together) and another 36 being so-called constellations (which are clustered SMP architectures that offer bigger computing nodes and bigger main memories than clusters of tiny servers).

These days, IBM systems utterly dominate the Top 500 rankings, and that is so by design. Big Blue has been very aggressive in the supercomputing market since the advent of its "Deep Blue" RS/6000 PowerParallel SP machines in 1993. (By the way, the Cornell Theory Center got the first SP box in 1994, and that might make this center, which is not the most prestigious in the world in terms of aggregate computing power, a very interesting leading indicator nonetheless. Watch the CTC, which was funded by IBM in the 1980s to build its first parallel supercomputer prototypes and which has been courted by Microsoft, very carefully.)

The biggest IBM computer on the list is also the most powerful in the world, the Blue Gene/L machine at Lawrence Livermore National Laboratory, built by the U.S. Department of Energy to simulate nuclear explosions and manage the nuclear arsenal. This machine has 131,072 32-bit PowerPC 440 processors running at 700 MHz, which runs a cut-down Linux kernel on its processing nodes. It is rated at 280.6 teraflops, which is a bit higher than IBM expected. The number two machine on the list is also a Blue Gene machine, but one at IBM's T.J. Watson Research Center with a rating of 91.3 teraflops.

Blue Gene was a research project launched in 1999 that was aimed at creating a box to do protein folding simulations and reaching a 1 petaflops performance level by 2010, but it has come out as a commercialized product that could end up making IBM money. Dave Turek, IBM's vice president of deep computing, says that 19 of the 219 systems IBM has on the list are Blue Gene boxes.

The reason is simple. "Blue Gene is not just about performance, but about the cost of electricity and cooling," he said, adding that the Blue Gene project has now more than paid for itself and that interest in the HPC marketplace for the BlueGene box, which was just commercialized late last year, remains quite high.

IBM also has another big RISC/Unix box on the Top 500 list, the ASCI Purple parallel pSeries server, also at LLNL. This box has 10,240 1.9 GHz Power5 processors (housed in p5 575 chasses) running AIX and using the "Federation" high performance switch. It is rated at 63.4 teraflops. The "MareNostrum" cluster at the Barcelona Supercomputer Center in Spain, which is a cluster of PowerPC blade servers running Linux and using Myrinet interconnect, is number eight on the list with 27.9 teraflops pf power; a baby Blue Gene box at the University of Groningen in the Netherlands rated at 27.45 teraflops rounds out IBM's entries in the top 10 rankings.

Those 219 IBM supers in the list have a combined 1.2 petaflops of computing power (52 percent of the total installed capacity) and over 444,654 processors--many of them being Xeons and Opterons as well as Powers and comprising an impressive 61 percent of total processors installed.

Rounding out the top 10 super sites in the world is the "Columbia" system at NASA-Ames built by SGI using its Altix Linux-Itanium servers and InfiniBand interconnect from Voltaire, which is rated at 51.9 teraflops. The "Thunderbird" Linux cluster built by Dell for Sandia National Laboratories using Xeon processors and InfiniBand interconnect comes in at number five with 38.27 teraflops, and it sits on the same site as the 36.2 teraflops "Red Storm" Linux-Opteron cluster with Cray's own interconnect.

Red Storm is ranked number six. The Japanese giant, "Earth Simulator," ruled the Top 500 list for some time, but at 35.86 teraflops, this parallel vector machine built by NEC has dropped to number seven on the list. Number 10 on the list is the Crazy XT3 Linux-Opteron cluster at the Oak Ridge National Laboratory. The "Thunder" Linux-Itanium cluster--built for LLNL by California Digital and using Quadrics interconnect--fell to number 11 on the list after being the number two machine only 18 months ago.

By vendor rankings, HP is the second-most prolific vendor on the current Top 500 list, with 169 systems that collectively account for 431.9 teraflops. While HP has 33.8 percent of systems, it has only 19 percent of the aggregate processing capacity in the Top 500 list. While it is hard to say what HP was counting on when it acquired Compaq many years ago, one of the things that seemed obvious was that the marriage of Compaq and HP, with a merged Itanium roadmap, would be able to take their collective leadership position and get a bigger chunk of the HPC server market.

There is little doubt that had Intel more aggressively delivered Itanium processors with better thermal properties and multiple cores per socket, HP would have a much more impressive showing in the Top 500 list. Itanium chips are excellent at exactly this kind of number-crunching work, but they run too hot and are too expensive compared to alternatives in HP's own product line.

Cray has a mix of Linux-Opteron and vector-based X1 systems on the Top 500 list totally up 18 machines, and SGI has 18 machines on the list as well (all of them Altix Linux-Itanium boxes). Linux Networx, which has been coming on strong in the HPC market in recent years, has almost the same slice of the pie, with 16 machines.

Dell has also grown its HPC business substantially, with 17 machines on the list, but it is limited in that it does not support Itanium or Opteron processors, which offer better performance on many applications than Xeons, and it does not have a high-end RISC/Unix play on which to chase the exotic clusters that IBM, HP, Fujitsu-Siemens, and Sun Microsystems Inc sell. Sun has fallen the furthest in recent years on the Top 500, with only four machines with a total of 9.1 teraflops on the November 2005 list.

Considering that Sun bought the supercomputing carcasses of one unit of Cray as well Kendall Square and Thinking Machines and then Gridware (now Grid Engine) for grid software, you would think that Sun would have a fairly large supercomputer presence, which would be reflected in the Top 500 list. Perhaps the power-conscious "Galaxy" Opteron servers, which support Linux and Windows as well as Solaris, will give Sun another run at the HPC server market. Time will tell, but Sun is clearly hoping for this to happen.

The Top 500 supercomputer rankings are compiled by Hans Meuer, of the University of Mannheim in Germany; Jack Dongarra of the University of Tennessee in Knoxville; and Erich Strohmaier and Horst Simon of NERSC/Lawrence Berkeley National Laboratory. The list is a leading indicator of the kinds of technologies that often get deployed in the wider HPC server market, which is obviously comprised of many more than 500 sites.

###

Agence France Presse

Microsoft throws weight behind supercomputer software

15 November 2005

US software titan Microsoft released super-computing server software on Tuesday, with the company's founder declaring such analytical power vital to innovation and invention.

The world's largest software maker released a "beta 2", or second test, version of Windows Compute Cluster Server 2003 software designed to combine the analytical power of varied computers in a network that acts as a single powerful server.

The new software represents Microsoft's first foray into the mainstream technical computing software market, a competitive arena that features low-cost, open-source products from companies such as Linux.

Microsoft founder and chairman Bill Gates announced the release at a Supercomputing 2005 conference in the Washington city of Seattle and said the company is funding 10 joint research projects at academic centers worldwide.

The "multi-million-dollar investment", he said, includes projects at Tokyo Institute of Technology, in Japan; University of Stuttgart, in Germany; Shanghai Jiao Tong University, in China; University of Southhampton, in England; and Nizhni Novgorod State University, in Russia.

"Technical computing is crucial to the many discoveries that impact our quality of life; from making safer, more efficient cars and airplanes to addressing global health issues and environmental changes," Gates said.

"We see many opportunities to collaborate with the scientific community on innovative solutions that will accelerate the pace of insight and discovery," the head of the Redmond, Washington-based group said.

Gates urged software designers to collaborate with academics and government agencies to make technical computing easier and more productive.

Among the ways to enhance scientific research is to use "modeling and cluster" techniques that enable groups of computers linked by the Internet to operate as unified systems to solve problems or analyze data, Gates said.

For decades, supercomputing has been the domain of "the few with specialized custom hardware and software systems," Microsoft senior vice president Bob Muglia said recently at a Microsoft IT Forum 2005 in Barcelona, Spain.

"What's needed now are software technologies that interconnect computing systems, people and data to produce more rapid answers to the questions of science, and to help researchers use computation in the most effective manner," Muglia said.

Red Herring

Microsoft Targets HPC Market

Gates and company see demand for affordable high-performance computing.

November 15, 2005

Microsoft attended to some long-unfinished business on Tuesday when it attempted to redefine the tough-to-penetrate world of high-performance computing with a direct appeal to key universities and an upgrade to its server cluster technology.

In keynote speeches on two continents, Microsoft attempted to storm the barrier that has kept its Windows technology out of the clubby environment of high-performance computing (HPC). That world has long been ruled by IBM and a small number of specialized system makers such as Hitachi and Fujitsu.

In that world, Unix has been one of the most commonly used operating systems. What has to really gall Microsoft is that Linux, a relative newcomer and open-source entry, has been readily accepted in that world.

A high percentage of the adoption opportunities for Linux in the business, scientific, and academic disciplines have come at the expense of Unix.

“This has been the area where Linux has made inroads,” said Roger Kay, president of Endpoint Technologies Associates. “Microsoft wants to head off the growth of Linux and protect its own weak flank. The popularity of Linux has also blunted Windows growth in the server environment.”

Analyst Charles King, principal analyst with PundIT Research, said that Linux not only has been accepted in the supercomputing world but is very popular in that space because of its flexibility.

“Linux allows users to go in and tweak the operating system, to customize it for specific purposes and design the applications to work in customized environments,” said Mr. King.

Microsoft’s attempts at industrial-grade operating systems have met with little more than modest success. Windows NT, for example, has not been readily accepted in the HPC environment. But elements of NT have been integrated into subsequent operating systems such as Windows 2000 and XP.

Microsoft Chairman Bill Gates, in a keynote speech at Supercomputing 2005 in Seattle, described a new world where mountains of information are being collected by a wider

variety of institutions and individuals than before. He said there is a greater need for affordable high-performance computing to cut down the time between insight and discovery.

Moving High Performance to Mainstream
In another keynote speech—this one given at Microsoft's IT Forum 2005 in Barcelona, Spain—Bob Muglia, senior vice president of Microsoft Server and Tools Business, said it is time to move high-performance computing into the mainstream.

"Scientific discoveries in the last decade have signaled a need for increased collaboration among scientists, as well as the ability to manage an increasing amount of data being generated by automated tools, such as sensors, satellites, and others," said Mr. Muglia.

"What's needed now are software technologies that will interconnect computing systems, people, and data to produce more rapid answers to the questions of science, and to help researchers use computation in the most effective manner," he added.

But Microsoft's international assault on the world of high-performance computing was not entirely about speech-making. The company also announced the Beta 2 release of Windows Compute Cluster Server 2003, the company's PC-based answer to traditional high-performance platforms.

Microsoft hopes this platform becomes the populist answer to Unix, and a rival to Linux for HPC environments.

"The supercomputing space is a logical place for Microsoft to go," said Mr. King. "Clustering x86 systems have brought supercomputing applications to the PC world. The question is whether anybody in the space wants Microsoft there."

University HPC Deployments
The Redmond giant also announced that it will invest funds to help establish 10 institutes for high-performance computing at key universities worldwide.

"This multi-year, multimillion-dollar investment builds on the relationship we started with the Cornell Theory Center years ago and will allow us to understand much more deeply the requirements and opportunities for applying results of research and innovation in high-performance computing," said Mr. Muglia.

Other universities included in Microsoft's HPV deployment are the University of Southampton, Tokyo Institute of Technology, the University of Tennessee, and the University of Texas.

"The research centers at the institutes have been running HPC platforms for years in critical research environments, so their scientists are experienced not only in advanced computing but also in a broad range of HPC platforms serving high numbers of users," said Mr. Muglia.

###

ABC News
Addict3D
Akron Beacon Journal
Beloit Daily News
Biloxi Sun Herald
Boonville Daily News
Boston Herald
Bradenton Herald
Brenham Daily Press
BusinessWeek
Canton Daily Ledger
Carthage Press
Centre Daily Times
Charlotte Observer
CIO Today
Columbus Ledger-Enquirer
Durant Daily Democrat
Duluth News Tribune
Forbes.com
Fort Wayne Journal Gazette
Fort Wayne News Sentinel
Fort Worth Star Telegram
Gadsden Times
Globe and Mail
Grand Forks Herald
Hartford Courant
Kansas City Star
Kentucky.com,
Kerala Online (India)
KOMO
Leading the Charge (Australia)
Macon Telegraph
Mexico Ledger
Miami Herald
mLive
Monterey County Herald
MSN Money
MSNBC
News and Observer
Newsday
Newsfactor
Norman Transcript
Pioneer Press
Rocky Mountain News
Sacramento Bee
San Jose Mercury News
San Luis Obispo Tribune
Sci-Tech Today
Seattle PI
SiliconValley.com
The State
Tacoma News Tribune
Tahlequah Daily Press
Tallahassee.com
TechWhack
Top-Tech News
Tribune Democrat
Washington Post
WJLA
Worcester Telegram
Yahoo!

Microsoft Creating Supercomputer Software

Microsoft Developing Software for High-Performance Computers

By ALLISON LINN AP Business Writer

SEATTLE Nov 15, 2005 — Microsoft Corp. is developing software for high-performance computers often used in engineering and science research, in a move that puts the company in another head-to-head battle with open-source developers.

Microsoft Chairman Bill Gates plans to announce the release of another test version of the company's software at a supercomputing industry conference Tuesday in Seattle.

Craig Mundie, one of Microsoft's chief technical officers, said the company is entering the market as high-powered computers are becoming more affordable, allowing the market to expand to a point where it makes sense for Microsoft to get into it. Such computers are used for computing, storing and sharing data for a wide variety of uses, ranging from creating better medications to engineering automobiles.

While such computers used to easily cost as much as $1 million, Mundie said some are now available for as little as $4,000 or $5,000.

The Microsoft product, called Windows Server 2003 Compute Cluster Edition, is due out in the first few months of 2006. The company has not yet set pricing.

With its new software, the Redmond company is aiming for small-to mid-sized versions of the powerful computers, mainly the type that are used by academic researchers and businesses. Mundie said the company also is hoping that its software will make it easier for less technically savvy people to set up such systems.

Currently, most such computers are run on custom-rigged software using Unix or open-source Linux technology. Unlike Microsoft, which believes that software blueprints should be closely guarded, people who develop using open-source technology hold that the underlying code and any developments should be freely shared.

On the Net:

http://sc05.supercomputing.org/

###

The News Tribune

Microsoft launches desktop search tool

ALLISON LINN; The Associated Press
Published: November 15th, 2005 02:30 AM

Microsoft Corp. today released a business version of software that aims to help people more quickly find documents, e-mail and other data stored on Windows-based computers.

The free new desktop search product comes after years of complaints over how hard it can be to locate Microsoft Word documents, sort through long e-mail lists and find other data people use during the workday.

The software will feel similar to Microsoft's consumer offering for searching files on desktop computers. But the product is designed so corporate technology executives can easily install it on many computers simultaneously and better control how it is used.

The product also will work with products that search corporate networks, including those offered by competitors.

Workers have for years complained about how difficult and time-consuming it can be to do things like search through e-mail or Microsoft Word documents. Rivals, including Yahoo Inc. and Google Inc., also have responded to that problem with their own desktop search offerings.

Microsoft sees the business- focused desktop search product as a step before the company releases its next version of Windows, called Vista, which promises better search capability.

However, the new version of Windows, due out next year, isn't going to include an even more advanced way to store and organize information, called WinFS. That will now come later. Microsoft also said that its consumer desktop search product will eventually let people search information found in Microsoft's new Windows Live online offerings.

Windows Live, another Microsoft effort to compete with Google and Yahoo, seeks to expand the Windows franchise online with better Web-based e-mail and other products. The search capability for those services, many of which are also still in development, is due out next year.

Super computer Microsoft Corp. is developing software for high-performance computers often used in engineering and science research, in a move that puts the company in another head-to-head battle with open-source developers.

Microsoft launches desktop search tool

Today's announcement: Microsoft Chairman Bill Gates plans to announce the release of another test version of the company's software at a supercomputing industry conference today in Seattle.

The product: Windows Server 2003 Compute Cluster Edition

Price: While such computers used to easily cost as much as $1 million, some are now available for as little as $4,000 or $5,000. Microsoft has not set a price on its product.

On sale: Early 2006

Use: Such computers are used for computing, storing and sharing data for a wide variety of uses.

###

Bloomberg

Microsoft Plans High-Performance Computer Software, NYT Says

By Dan Hart

2005-11-15 05:54 (New York)

Microsoft Corp. plans to offer software that will give scientists and engineers a way to gain high-performance computing from their existing desktop computers. Microsoft is set to release its Windows Computer Cluster Server 2003 software in the first half of next year, which is intended for systems that have as many as 64 processors and can be extended to larger machines.

Byte and Switch

Cluster Clamor

November 15, 2005

Clusters are taking center stage in the news this week, as Microsoft plans an announcement tomorrow, and NAS vendors Isilon, OnStor, and Panasas today disclosed new products.

Microsoft CEO Bill Gates is expected to play up Windows's role in large computer clusters in his Supercomputing 2005 keynote Tuesday in Seattle. And Microsoft will unveil "Beta 2" of its Compute Clusters Server (CCS) 2003 in an effort to demonstrate easier-to-implement high-performance clusters in Windows.

Until now, Unix and Linux have been the major operating systems of choice for clustered computing. Microsoft hopes to counter that by tweaking Windows CCS with cluster job scheduling, increased InfiniBand support, and a Microsoft Message Passing Interface (MPI) compatible with the Linux MPI.

We won't know for yet how successful Microsoft may be in its clustering initiative. Windows CCS won't be available until the first half of next year.

Meanwhile, Isilon, OnStor, and Panasas are looking to build on inroads they've already made with clustering in the competitive NAS space. At this week's conference, Isilon is unveiling an accelerator node for its clustered storage systems, OnStor is rolling out a new NAS gateway with improved clustered performance, and Panasas is adding blades for its ActiveScale Storage Cluster (See Isilon Intros IQ Accelerator.)

Isilon's IQ Accelerator is a device aimed at improving the performance of its other products. It looks like Isilon IQ nodes, except the accelerator nodes have no disk, and their function is to accelerate performance on the disk-based nodes -- for systems that don't require extra capacity, that is. (See Isilon Shows Off Its IQ .) IQ Acclerator is scheduled to ship in the first quarter of 2006.

OnStor calls its Bobcat 2280 a streaming media server. The vendor claims it is optimized for high-performance computing verticals that deal with large files, such as life sciences, medical imaging, and content delivery. The new Bobcat is different from OnStor's previous gateways because it has an extra accelerator chip, called a FastPath processor, placed between a front-end storage processor and a back-end file processor. (See ONStor Hits NAS Gas.) OnStor claims the third processor makes the system 40 percent faster.

One customer claims even better performance gains. John Welter, VP of operations for North West Geomatics, says he's running a four-way cluster with Bobcat 2260 gateways. He recently added two 2280s for testing. (See North West Geomatics Selects ONStor.)

"We deal with huge image files, two to three gigabytes in size," he says. "We're seeing close to a 50 percent increase in speed. With clustering, performance is very important."

Welter says he originally chose OnStor over Isilon because OnStor connects to any vendor's SAN, while Isilon only works with its own storage.

Then again, Isilon has one thing ONStor doesn't offer yet -- InfiniBand connectivity to speed performance. (See WilTel Wins $35M Contract.)

Another customer who emphasizes performance overlooksn chose Isilon's hardware lock-in because of the specialized work he does. Parag Mallick, director of proteomics at Cedar-Sinai Medical Center, plans to add an IQ Accelerator node to his Isilon systems.

Proteomics consists of studying protein in a person's blood to determine if that person has a disease. Mallick says each drop of blood he analyzes represents 50 Gbytes of data and his work generates more than 1 Tbyte a day.

"When you do pattern discovery, you want to pull all the data you've analyzed quickly," Mallick says. "There are lots and lots of reads required over your entire cluster. One of the advantages of clustered computing is having multiple read heads."

North West uses Linux and Cedars Sinai uses Linux and Solaris Unix for most of their computing. North West's Welter says that could change after Microsoft gets CCS out. He compares it to Microsoft's foray into IP SANs with its iSCSI software initiator a few years back. (See Microsoft Delivers iSCSI for Windows.)

"Microsoft's adding clustering capabilities that have been available on Linux," Welter says. "For instance, we have one application that runs on Windows and we're forced to run Gigabit Ethernet, while on Linux you can run InfiniBand or whatever you want. This is like iSCSI. You could do iSCSI on Windows before Microsoft came out with its initiator, but it was ugly."

Panasas' system is built primarily for Linux clusters, although it does support Unix NFS and Windows' CIFS protocols through a gateway between its storage blades and the NFS or CIFS clients. Its new StorageBlades hold 8 TBytes of capacity per 4U shelf and up to 80 TBytes per cabinet. Customers can cluster cabinets.

— Dave Raffo, Senior Editor, *Byte and Switch*

Organizations mentioned in this article:

· Microsoft Corp. (Nasdaq: MSFT - message board)
· Isilon Systems
· ONStor Inc.
· Panasas, Inc.
· North West Geomatics
· Cedars-Sinai Medical Center

###

ComputerWeekly

Gates to unveil supercomputing strategy

by **Antony Savvas**
Tuesday 15 November 2005

Microsoft chief software architect Bill Gates will use this week's supercomputing show in Seattle, Washington, to explain how the company will enter the supercomputing market.

Although Microsoft's first operating system for supercomputing clustering remains in beta, the company is keen to take advantage of a growing market for clustered server platforms aimed at the enterprise market.

Sun Microsystems, Hewlett-Packard, IBM and others are all looking to capture a share of the market that is moving out of the science lab and university and into the enterprise, as the need for ever-higher computer processing increases.

At the same show for instance, Linux Networx, which makes supercomputers made from clusters of Linux machines, unveiled its new LS-1 platform.

The company is aiming the LS-1 at potential customers who have never owned a clustered product before and have so far been put off by the belief that such a platform is too complicated to manage.

Gates is expected to announce that Microsoft has started to test the Beta 2 of its Windows Compute Cluster Server 2003. The product includes software for job scheduling and other tasks.

Windows Compute Cluster Server is expected to ship in the first half of 2006.

Although Linux is enjoying increasing popularity in the academic world when it comes to high-end clustering, Gates will tell his audience that a clustered Windows platform will integrate much more easily with existing Windows-based enterprise platforms.

ComputerWorld
CIO

Smaller companies eye supercomputing

The use of x86-based chips helps boost performance while keeping costs down

News Story by Patrick Thibodeau

NOVEMBER 14, 2005 (COMPUTERWORLD) - SEATTLE -- Golf club maker Ping Inc. is not a typical supercomputing user. It's not making airplanes, looking for oil or investigating proteins. It is, instead, a company with 1,000 employees that designs golf clubs.

But this midsize company may represent a big part of the future of supercomputing. Ping is using a Cray Inc. supercomputer built with Advanced Micro Devices Inc. Opteron chips to run simulations of golf club designs. And Eric Morales, staff engineer at Ping, said supercomputing has allowed the Phoenix-based golf club maker to drastically reduce development time. Cray earlier this month announced Ping's use of its XD1 system, which can support 12 to 800 processing cores.

"It takes the development from weeks down to days, and it helps us get to market faster," said Morales. Simulations of product changes that once took a full day to run can now be processed in 20 minutes or less.

Morales believes many other midsize companies will turn to supercomputing -- and if the attendance at this year's Supercomputing 2005 is any indication, high-performance computing is drawing a lot more interest. As of yesterday, some 8,600 people had registered for the conference here, with Microsoft Corp. Chairman Bill Gates expected to give the conference keynote address tomorrow. Conference attendance is up about 10% from last year.

Gates' speech tomorrow "signals that Microsoft is interested in the technical computing market," said Earl Joseph, an analyst at IDC. Coinciding with that speech, Microsoft is expected to announce an update of its compute cluster server operating system.

The fact that high performance software is of interest to Microsoft is recognition of a

market that has reached $7.25 billion this year, a net increase of 49% over the last two years, said Joseph. The market is expected to continue to gain ground with technical and commercial computing users, "primarily due to the attractive pricing of clusters, combined with their growing capabilities," said Joseph.

"We are seeing many new users enter into technical computing due to the greatly lower cost of entry with clusters," he said, noting that many users are finding it faster, more accurate and cheaper work on science and engineering projects on such systems.

Price/performance improvements, due in part to systems using x86-based chips, are making supercomputing more accessible to businesses, according to vendors. The Top500 supercomputer list, released earlier today at the show, showed large gains in the use of commodity chips (see "IBM holds on to Top500 supercomputer lead").
Dave Turek, vice president of IBM Deep Computing, said demand for supercomputing is being driven by a number of forces, including the creation of new businesses that rely heavily on high-performance systems, such as digital animation and bioinformatics. Supercomputing “is also being driven by vast amounts of data that demand rapid analysis for real-time decision-making,” he said.

Turek pointed in particular to RFID-enabled devices, which can be used to track product shipments, as capable of “generating huge amounts of data.”

While enterprise vendors are making capacity-on-demand systems available for processing large amounts of data, Turek doesn’t believe those systems will account for more than 10% of the market.

Ed Turkel, manager of product marketing and high performance computing at Hewlett-Packard Co., said high-performance clusters, ranging from 32 to 64 nodes (a node is typically a two-processor system), are being increasingly adopted in industrial markets. Microsoft offers its own cluster product, which Turkel believes will help HP sell high-performance clusters to more commercial users.

###

ComputerWorld

Microsoft offers Beta 2 of Compute Cluster Server 2003

The final version of the software is due out by June 2006

News Story by Eric Lai

NOVEMBER 15, 2005 (COMPUTERWORLD) - Microsoft Corp. continues to edge toward the high-performance computing space, releasing today the second beta of its Windows Server 2003 operating system for clustered systems.

Beta 2 of Windows Compute Cluster Server 2003 is being tested by Microsoft users in clusters as large as 128 nodes, all connected in a single system aimed at superfast problem solving.

Built on the 64-bit version of Windows Server 2003, the clustering edition is expected to be released by mid-2006, according to Kyril Faenov, director of high-performance computing with the Windows Server Group.

As it has done in the past with other software markets such as databases and CRM, Microsoft is expected to first target lower-end departmental and workgroup clustering by offering easier-to-use features and integration for companies already running Windows, said Laura DiDio, an analyst at Boston-based Yankee Group.

The first beta version of the software was released in September and is being tested by 1,600 companies and institutions, including the Seattle-based genetics research lab of Merck & Co. The drug company is just beginning to set up a cluster of 20 64-bit dual-processor machines, according to Eric Schadt, senior scientific director for genetics research.

Schadt's group now uses a 500-processor cluster running Linux on a variety of hardware to create simulations of gene networks and their behavior during drug treatments. Migrating to Windows Compute Cluster, he said, will save time because Merck does most of its prototyping and data mining in Windows.

“What a great day it will be when I don’t have to have a dual-boot computing environment where I am always switching between Linux and Windows, depending on the problem of the day,” Schadt said.

Installing the Windows software “has not been too painful,” he said. “I think Windows will catch up quickly [to Linux] if the environment works well, which does remain to be seen.”

Microsoft’s announcement was made during the annual SuperComputing conference in Seattle. Microsoft Chairman Bill Gates was slated to deliver the keynote address today at the event, which runs through Friday.

High-performance computing is a new market for Microsoft. The latest list of supercomputers released yesterday by Top500.org showed that 74% of the 500 fastest supercomputers in the world run Linux, 20% use various flavors of Unix, and the rest run other operating systems, including Mac OS X. None run Windows (see "IBM holds onto Top500 supercomputer lead").

But the door to the market is now opening for Microsoft with the rapid uptake in 32- or 64-bit systems. More than three-fourths of the 500 fastest supercomputers used Intel Corp. or Advanced Micro Devices Inc. processors, which both support Windows. By contrast, 41% of supercomputers relied on Intel or AMD chips in November 2003, and just 12% did so in November 2002, according to Top500.org.
The Cluster Server’s current 128-node limit may hinder the software’s appeal at large scientific or medical institutions for now, but it may fit well with the more modest performance needs of mainstream enterprises, DiDio said. “There is such a thing as too much firepower,” she said.

Microsoft has not announced a price for Windows Compute Cluster Server 2003. It has previously said that it expects to price each copy less than the Standard edition of Windows Server 2003, which starts at $999.

###

eWeek.com

Supercomputers: Windows' New Best Friend

By Peter Galli
November 15, 2005

Microsoft Corp. chairman and chief software architect Bill Gates will use his keynote address at the **Supercomputing show,** being held in Seattle this week, to announce the second beta for Windows Compute Cluster Server 2003, a 64-bit operating system for industry-standard x64 processors.

The first, more limited beta, was announced at Microsoft's Professional Developers Conference in Los Angeles in September.

"We expect our customers to start pre-production deployment evaluations of the product with this beta," Kyril Faenov, Microsoft's director for High Performance Computing, told eWEEK, adding that the team is not planning for a third beta, and a release candidate will be next.
Asked what feedback Microsoft has received from the developers who tested beta one, Faenov said it had become clear just how important it is to provide them with diagnostic tools.

"To a large degree, it is not as much the software challenges but the hardware challenges that come into play when some cluster configurations are deployed.

"So diagnostics become really important, and we have done some good thinking and made some improvements in that area. There was also a lot of focus during beta two development on optimizing the performance," he said.

It was also important for ISVs to be able to rapidly migrate their code, and Microsoft will be showcasing about 19 different applications, such as a car crash simulation guide, a weather modeling guide, as well as a demonstration of the **BioTeam's iNquiry software**—which includes 160 open source Bio applications—on a Windows CCS cluster at its booth at the Supercomputing show this week, Faenov said.

"These companies were all able to move to Windows 64-bit if they were not already there, pick up our FDI and just run with it, and that has been very encouraging," he said, adding that the product is still on track to ship in the first half of 2006.

Asked what its strategy is towards the Linux and open-source community, given its dominance in the HPC field and the fact that Microsoft is including the Message Passing Interface—a library specification for message passing proposed as a standard by a broad-

based committee of vendors, implementers and users—in the product, Faenov said Microsoft had no such strategy per se.

"Our strategy is to deliver the best value and help grow the market, and we are making very pragmatic choices based on specific requirements we hear partners or customers want. MPI was an example of this and if there are other places where this makes sense for us, our customers and/or partners, we'll look at them," he said.

All of the major original equipment manufacturers, like IBM, HP, Dell, NEC; those developing more personal clusters, like Orion Multisystems; as well as all the major interconnect vendors, have also announced support for the product, Faenov said.

Microsoft will release a single 64-bit only version of Compute Cluster Server 2003, which would also run on all the hardware platforms supported by Windows Server 2003 SP1, on which it is based.

"We are going to work with our hardware partners to tailor the systems to particular market scenarios and then to do benchmarks and evidence generation for key applications," he said.

The announcement of beta two marks a milestone for the Redmond, Wash., software maker, which will be underscored in Gates' keynote, Faenov said, where he will stress the company's long-term commitment to high-performance computing as well as to a broader technical and scientific computing area.

There are three main pillars for Microsoft: business computing, consumer computing and now technical and scientific computing, an area in which there is room for many software advances, Faenov said.

Gates will also point to the need for software to help scientists and researchers focus on their work and not on the underlying IT, he said.

Microsoft officials first announced its plans to enter the HPC software market in June 2004 and said at that time that the product would be specifically designed for customers running scalable, parallel computing workloads in vertical market segments such as engineering, life sciences and finance.

Company officials such as CEO Steve Ballmer have also admitted that Linux is ahead in the area of high-performance computing, but have said Microsoft's staff comes to work every day looking at how to offer customers an even better solution.

Gates will also talk about Microsoft's multimillion-dollar investments in, and multi-year commitment to, 10 Institutes for High Performance Computing around the globe, where it is creating funded research positions at existing facilities such as the University of Utah,

the Cornell Theory Center, the University of Stuttgart, Germany, and the Tokyo Institute of Technology.

"The centers we chose have a unique blend of expertise as they do research in HPC systems and, in many cases, are closely affiliated with a HPC center, giving them a unique insight from both a scientific and industry perspective on what will actually help solve real problems," Faenov said.

The agreements with Microsoft will see the staff at these centers deploy its HPC and other products and then work with Microsoft staff to explore and maximize its product roadmap going forward, he said.

Looking to the next release of the product, Faenov said those plans remain on track.

He previously told eWEEK that the development team will also continue to invest in making clusters easier to manage, particularly as their use moves away from centralized resources and down to the workgroup level.

Microsoft also planned to work on improving the user experience through better power management and use of the IPMI (Intelligent Platform Management Interface) or Web services management to reboot clusters without having to power down, Faenov said.

"[The second version] will look at how we can control the hardware using the standard ways of rebooting. We will also look at how we can make the integration of clusters into the applications, including workstation applications, even easier. We are also looking at performance monitoring tools for the cluster space as a whole, which is a big area for development," he said previously.

###

ARNnet
ComputerWorld
InfoWorld

Microsoft to talk up Windows for clusters

ELIZABETH MONTALBANO, IDG NEWS SERVICE

14/11/2005 08:00:22

Microsoft will tout its version of Windows for high-performance computing at the Supercomputing 2005 show this week in Seattle, but observers said the company faces a considerable challenge in selling the OS for compute-intensive environments already dominated by Unix.

Microsoft Chairman and Chief Software Architect Bill Gates will keynote at the show on Tuesday to reinforce Windows' position as a platform for large compute clusters. Microsoft is planning to ship an OS tuned for high-performance environments, Windows Compute Cluster Server 2003, as soon as January, according to sources familiar with the company's plans.

Microsoft originally hoped the OS would be available in time for the show next week, but plans were pushed back earlier this year. Bob Muglia, senior vice president of the Windows server division at Microsoft, unveiled the first beta of Windows Server 2003 Compute Cluster Edition in September at Microsoft's Professional Developers Conference, and said the full release would be available in the first half of 2006.

A spokeswoman for Microsoft's public relations firm Waggener Edstrom said in an e-mail message Friday that Microsoft aims to release the OS in the first half of 2006.

Microsoft also could unveil a new build or another beta of the product next week, depending on the early feedback it has gotten on beta 1 so far, sources said.

Windows Compute Cluster Server 2003 comprises both the Compute Cluster Edition of the OS and Microsoft Compute Cluster Pack, a combination of interfaces, utilities and management infrastructure.

With Windows Compute Cluster, Microsoft aims to compete against Unix and Linux to

run server clusters in compute-intensive environments, such as those running multiple simultaneous transactions or computations involving large amounts of data.

To make sure the company is building a competitive product, Microsoft has been working with university researchers to fine tune the OS, said Andrew Brust, chief of new technology at Citigate Hudson, a consulting company in New York.

"If they want something credibly competitive with Linux, going to academia is a good idea," he said. "Most cluster stuff has historically been on Unix and more recently Linux, and that's come out of the academic environment."

Brust said Microsoft hopes to expand Windows' presence in financial markets and also in enterprises for data mining and business intelligence with Windows Compute Cluster Server 2003.

Phil Papadopoulos, program director of grids and clusters computing at the San Diego Supercomputer Center, said that engineers in compute-intensive environments are very comfortable working with Unix and Linux because it is easier to remotely manage these environments than it is to manage Windows. It's unlikely they will be inclined to switch to a Windows environment unless there is a considerable valuable proposition, he said.

"Microsoft has somewhat of an uphill battle coming into the supercomputing market because it's so dominated by Unix and Linux," Papadopoulos said. "It doesn't mean they couldn't do it, it just means they have their work cut out for them."

###

InfoWorld

Customers get chance to test Windows for clusters

Redmond distributes second beta of Windows Server 2003 Compute Cluster Edition

By Elizabeth Montalbano, IDG News Service
November 15, 2005

The Redmond, Washington, software vendor will distribute the second beta of its Windows Server 2003 Compute Cluster Edition at the show, where Microsoft Chairman and Chief Software Architect Bill Gates will keynote Tuesday, said Kyril Faenov, director of high-performance computing at Microsoft.

Microsoft released Windows Compute Cluster Edition's first beta, aimed at developers, at the Professional Developers Conference in September.

Beta 2 of the product is aimed at allowing customers to test how the software will perform in their data centers, Faenov said. "It is for customers deploying evaluation versions of our product in their infrastructure," he said.

In Gates's keynote at the show, Microsoft also will reveal that it has set up 10 high-performance computing institutes around the world in universities to drive both research in those institutions as well as the company's own efforts to produce software that is appropriate for compute-intensive environments, said Zane Adam, director of marketing for Microsoft's Windows Server division.

"It's a two-way street," he said of the research centers. "Our platform is being used for a learning experience, and researchers are working on it so they can transfer their knowledge to us. It extends our research and development."

The centers have been set up at the following universities in the U.S.: University of Texas, University of Utah, University of Virginia, Cornell Theory Center at Cornell University and University of Tennessee. In the U.K., Microsoft has set up a high-performance computing center at Southampton University, and European centers have been established at the University of Stuttgart in Germany and Nizhni Novgorod University in Russia. In Asia, Microsoft has set up two institutes: one at the Tokyo Institute of Technology and one at Shanghai Jiao Tong University.

Microsoft has made a multiyear, multimillion-dollar commitment to providing software for and helping direct research in these centers, Faenov said. Each center, which is run by its own director -- often a professor at the respective university -- is responsible for participating actively in conferences within the high-performance community, providing industry training, giving Microsoft feedback on its high-performance computing software and providing technology prototypes for the next version of Windows Compute Cluster Edition.

Microsoft also will showcase support from applications vendors and hardware OEMs (original equipment manufacturers) at the supercomputing show, Faenov said. Microsoft plans to demonstrate about 20 high-performance computing applications for industries such as automotive, airspace, oil and gas and biotechnology running on Windows Compute Cluster Edition, he said.

Windows Compute Cluster Server 2003 comprises both the Compute Cluster Edition of the OS and Microsoft Compute Cluster Pack, a combination of interfaces, utilities and management infrastructure. Windows Compute Cluster Server 2003 is expected to be available in the first half of 2006.

Observers said Microsoft has a tough road ahead in selling a new version of Windows to a market where Unix and **Linux** (Overview, Articles, Company) already have a firm hold, but Microsoft believes that the prevalence of Windows clients could give it an edge.

Most servers in the high-performance computing space are sold for a price of $250,000 or less, Adam said, citing IDC research. Many of these are departmental deployments, a place where Windows already is widely deployed but in an application customers have cobbled together, he said. While Microsoft is mainly targeting this high-volume market segment with its high-performance version of Windows, the company also will play in the higher end of the market.

###

InternetNews.com

Supercomputing With Microsoft

By Clint Boulton
November 15, 2005

Microsoft's plan to use the Internet to provide software services may be garnering the lion's share of attention, but the software giant is pretty keen on making headway in a sector where it is a relative unknown: high-performance computing.

The company is expected to unveil Windows Compute Cluster Server 2003 today, a new piece of cluster software designed to bring powerful tools to customers who require advanced computing but don't have deep coffers. Microsoft Chairman Bill Gates is expected to present a keynote later today about the company's supercomputing vision during the Supercomputing 2005 event in Seattle.

The second beta of Windows Compute Cluster Server 2003 will include CDs to install Windows Server 2003 Compute Cluster Edition and Microsoft, both of which are designed to run on Windows Server 2003 SP1 x64 Edition, said John Borozan, group product manager for Microsoft's server division.

While most supercomputing solutions are pieced together and a mish-mash of different products and protocols, Microsoft aims to provide a complete Windows platform for HPC to help customers get up and running more quickly.

To wit, the software will support Ethernet and Infiniband protocols and is integrated with the following Microsoft technologies: cluster setup and administration; compute node management based on images; security based on Microsoft's Active Directory; and job scheduling and resource management.

Microsoft expects to target departmental and workgroup clusters in the $50,000 to $250,000 range for manufacturing, geosciences, life sciences, oil and gas and financial services sectors.

Borozan said Microsoft's entry into the market is predicated on some rising trends, namely the falling cost of HPC and its increase in enterprise deployments.

The executive said the cost of running supercomputers has dropped significantly: In 1991, $40,000,000 would get a customer 10 gigaflops of compute power. Today, with so many machines performing in the teraflop range, the cost of 10 gigaflops is roughly $4,000.

HPC is dipping into mainstream commercial deployments. Market research from IDC found that HPC deployments grew 70 percent in 2004, with most of the growth coming in

departmental or workgroup clusters that cost less than $50,000.

It also doesn't hurt that x86 server clusters, the area Microsoft hopes to serve, are growing.

"We think the opportunity is ripe for us to enter the market," Borozan said. "We believe the next big revolution in both science and industry is going to be data driven. There are mountains of data being produced by computational models or sensors, but until now the ability to do more with that data and crunch it in a fashion that is used in HPC has been out of reach."

Despite targeting a specific, lower end of the market, Microsoft's competition will be steep, with IBM, HP, Dell, Sun Microsystems waiting in the wings.

But Microsoft has already garnered support from 19 ISVs that have ported their applications to Compute Cluster 2003.

Moreover, Microsoft has pumped multi-year, multi-million dollar investments into joint projects at 10 HPC institutes worldwide. The groups, including the Cornell Theory Center and University of Virginia will do advanced research using Compute Cluster 2003.

Separately, Cisco Systems this week unveiled new HPC solutions. The new Cisco SFS 7012 and Cisco SFS 7024 switches use InfiniBand technology to provide a unified fabric for connecting servers together into grids of compute resources.

Working with the SFS advanced Ethernet and Fibre Channel gateway, the switches connect server grids with shared LAN and SAN resources connected through Cisco Catalyst switches or Cisco MDS 9000 storage networking switches.

Cisco also has new HPC software on tap in the form of the SFS Subnet Management software, which scales beyond 4000 server nodes and recalls a multi-thousand node InfiniBand cluster fabric in less than a minute.

###

IT Jungle

Linux Clusters Continue to Expand in Top 500 Supers Ranking

By Timothy Prickett Morgan

The number of parallel supercomputing clusters and the aggregate performance of these machines both continued their expansion in the latest Top 500 supercomputer rankings. The Top 500 list was announced at the SuperComputing 2005 tradeshow in Seattle, which is not exactly a coincidence. Microsoft, having seen Linux take off as the core platform on cluster supercomputers in the past five years, will debut its own Windows HPC variant at SC2005 this week. The question now is: Can Windows storm the Linux stronghold?

If the Top 500 is any indication, it would seem not. In a supercomputing environment, support for Fortran, C, and sometimes Java is critical, and so is the reliability of the core operating system platform deployed on a cluster. Excepting the very high-end, exotic machines that tend to dominate the Top 500 rankings, which are basically hand-tooled, very expensive supercomputers that only government agencies and academic research institutions can afford, most big HPC server buyers are interested in using as many commodity components as possible. That means X86 and now X64 processors and usually a Linux operating system. Windows can run on the same iron, of course, which is why it has a much better shot than Unix. But all of the compilers, workload management, job scheduling, cluster management, and resiliency software that has been either ported from Unix to Linux or created out of thin air by the open source HPC community has to be moved to Windows for Windows to have a chance in the HPC market. Bill Gates, Microsoft's founder, chief software architect, and a former compiler maker in his own right, might be a highly respected individual in the IT market, but research labs are going to expect more from a Windows HPC variant than they get on their desktops. It will be very interesting to see what penetration Microsoft gets in the Top 500 list in the next two years. There is certainly potential for Microsoft to win a lot of deals, as Linux distributors have done.

But for now, the Top 500 is about Linux and Unix, and the $7-plus billion HPC server market is one where Microsoft has almost no presence except for a few token accounts. The current Top 500 list makes this very clear. For instance, there is not a single Windows machine on the list, even though there are a few experimental Windows clusters out there in the world. The reason is simple: to get on the Top 500 list, the ante is now 1.64 teraflops, which is almost double the performance of the 500th system on the list. If Microsoft wants Windows to get a toe-hold in the Top 500 list in the coming years, it is going to have to find some very big institutions to use its software. Eating its own dog food is probably the best place to start, and it seems likely that Microsoft Research,

the software giant's research arm, will probably start moving up the Top 500 list after the company installs its own cluster; it is possible that Microsoft pumps the money for a giant Windows cluster into the Cornell Theory Center or some other academic institution as well.

Right now, Linux rules the Top 500, after entering the list in 1998. In the latest Top 500 rankings, Linux machines account for 372 out of 500 machines (74 percent), 1.19 petaflops, or thousands of teraflops, out of a total of 2.3 petaflops of aggregate sustained computing capacity (52 percent), and 333,373 processors out of a total of 731,069 (46 percent) across those 500 machines. Various Unixes (including AIX, Solaris, HP-UX, the open source BSDs, and the BSD-derived Mac OS X) account for 109 machines, and the remaining 19 machines on the list are mixed clusters that use both Linux and Unix.

By processor architecture, there are only 14 vector machines in the whole Top 500, which shows how far clusters based on scalar, commodity processors have come in the past 15 years. In June 1993, there were 334 vector machines on this list accounting for about a third of the aggregate 1.1 teraflops of computing power in the Top 500 list (the first one ever published, in fact), and the 131 other machines were exotic parallel clusters built by companies such as Kendall Square, Thinking Machines, MasPar, and a bunch of smaller vendors that no longer exist. And tellingly, IBM did not have a single machine on the list. Most of these machines had proprietary interconnects, proprietary processors, and a basically proprietary implementation of Unix.

Fast forward to November 2005's list, and most of the machines are built on commodity Intel or AMD processors: There are 206 machines built from 32-bit Xeons, 46 built from Itaniums, 81 built from 64-bit Xeons, and 55 built from Opterons; IBM's Power family of processors have a nice showing, with 73 machines.

The dominant interconnection method is off-the-shelf Gigabit Ethernet, with 249 machines in the November 2005 list, followed by 70 machines using Myricom's Myrinet interconnect, 26 using InfiniBand, 42 using one or another variant of IBM's SP switching technology from its pSeries Unix family, 19 using Silicon Graphics' NUMALink, 31 using Hewlett-Packard's HyperPlex. There's a mix of crossbars and other interconnects in the list, but clearly people are perfectly happy to build X64 and GigE clusters, and are doing so with wild enthusiasm. On the current list, 360 of the 500 machines are clusters, with 104 being classified as massively parallel processors (which means they have a high-bandwidth interconnect that more tightly couples the machines together) and another 36 being so-called constellations (which are clustered SMP architectures that offer bigger computing nodes and bigger main memories than clusters of tiny servers).

These days, IBM systems utterly dominate the Top 500 rankings, and that is so by design. Big Blue has been very aggressive in the supercomputing market since the advent of its "Deep Blue" RS/6000 PowerParallel SP machines in 1993. (By the way, the Cornell Theory Center got the first SP box in 1994, and that might make this center, which is not the most prestigious in the world in terms of aggregate computing power, a very

interesting leading indicator nonetheless. Watch the CTC, which was funded by IBM in the 1980s to build its first parallel supercomputer prototypes and which has been courted by Microsoft, very carefully.)

The biggest IBM computer on the list is also the most powerful in the world, the Blue Gene/L machine at Lawrence Livermore National Laboratory, built by the U.S. Department of Energy to simulate nuclear explosions and manage the nuclear arsenal. This machine has 131,072 32-bit PowerPC 440 processors running at 700 MHz, which runs a cut-down Linux kernel on its processing nodes. It is rated at 280.6 teraflops, which is a bit higher than IBM expected. The number two machine on the list is also a Blue Gene machine, but one at IBM's T.J. Watson Research Center with a rating of 91.3 teraflops. Blue Gene was a research project launched in 1999 that was aimed at creating a box to do protein folding simulations and reaching a 1 petaflops performance level by 2010, but it has come out as a commercialized product that could end up making IBM money. Dave Turek, IBM's vice president of deep computing, says that 19 of the 219 systems IBM has on the list are Blue Gene boxes. The reason is simple. "Blue Gene is not just about performance, but about the cost of electricity and cooling," he explains, adding that the Blue Gene project has now more than paid for itself and that interest in the HPC marketplace for the BlueGene box, which was just commercialized late last year, remains quite high. IBM also has another big RISC/Unix box on the Top 500 list, the ASCI Purple parallel pSeries server, also at LLNL. This box has 10,240 1.9 GHz Power5 processors (housed in p5 575 chasses) running AIX and using the "Federation" high performance switch. It is rated at 63.4 teraflops. The "MareNostrum" cluster at the Barcelona Supercomputer Center in Spain, which is a cluster of PowerPC blade servers running Linux and using Myrinet interconnect, is number eight on the list with 27.9 teraflops pf power; a baby Blue Gene box at the University of Groningen in the Netherlands rated at 27.45 teraflops rounds out IBM's entries in the top 10 rankings. Those 219 IBM supers in the list have a combined 1.2 petaflops of computing power (52 percent of the total installed capacity) and over 444,654 processors--many of them being Xeons and Opterons as well as Powers and comprising an impressive 61 percent of total processors installed.

###

MicrosoftWatch

Microsoft Pushes Users Down the 64-Bit Plank

By Mary Jo Foley
Tuesday, November 15, 2005

Microsoft has drawn a line in the sand. On Tuesday, at the IT Forum in Barcelona, officials went public with a list of software which the company is planning to deliver 64-bit versions only.> The list: Exchange Server 12, Windows Compute Cluster Server 2003, the Longhorn Small Business Server, Microsoft's midmarket server, code-named "Centro," and Windows Longhorn Server R2 (the Windows Server release expected around 2009) . All of these "will be exclusively 64-bit and optimized for x64 hardware," Microsoft officials said. In related news, Microsoft released Beta 2 of its Compute Cluster Server 2003 on Tuesday, as well.

New York Times
International Herald Tribune

Microsoft's high-performance goals

By John Markoff The New York Times
TUESDAY, NOVEMBER 15, 2005

SAN FRANCISCO In January a group of Microsoft researchers set out to discover how much computing power they could buy for less than $4,000 at a standard online retailer.

They found the answer at NewEgg.com, where they were able to purchase - for just $3,632 - 9.5 gigaflops of computing speed. That is the amount of computing power that a Cray Y-MP supercomputer offered in 1991, at a cost of $40 million.

The plunging cost of computing power is both an opportunity and a challenge to Microsoft, which on Tuesday was to unveil its first entry into the market for high-performance scientific and technical computing.

The company's Windows Computer Cluster Server 2003 software is scheduled to become available in the first half of next year and is intended to give scientists and engineers a simple way to gain high-performance computing from their existing desktop computers.

The high-performance market is growing faster than the rest of the server market, according to International Data Corp. Last year, the proportion of high-performance servers grew from 7 percent to 10 percent of the computer server market using X86 industry standard chips.

The new Microsoft software is intended for systems with as many as 64 processors but can be extended to much larger machines as well, if they are linked internally on high-speed data networking connections, the company said.

"Our focus is not on the very highest-end systems but on divisional and departmental computing systems," said Kyril Faenov, Microsoft's director of high-performance computing.

The entry is significant, because until now Microsoft has largely been excluded from the high-performance market, which is dominated by Linux and Unix software.

But Microsoft is planning a significant marketing push into the field with a keynote speech by Bill Gates, the company's co-founder and chairman, on Tuesday at the annual supercomputing trade show this week in Seattle.

Microsoft, based in nearby Redmond, Washington, was also to announce that it was

planning to provide international support for 10 supercomputing institutes around the world, including Cornell University, the Universities of Utah and Stuttgart and Shanghai Jiao Tong University.

"We think the big deal here is to give a lot more people access to a level of computation that was not available before," said Craig Mundie, a senior vice president at Microsoft and one of its three chief technology officers.

Microsoft is hoping to leverage its near-monopoly position in desktop computing by offering scientists and engineers a single computing environment with the thousands of applications now available for desktop and server operating systems.

The low end of the technical computing market is, however, already highly competitive with software systems, many of them available as inexpensive open-source software programs.

For example, the Linux Rocks program has been developed for more than a decade by a small group of engineers led by Philip Papadopoulos, program director of grid and cluster computing at the San Diego Supercomputer Center.

That program, which is freely available and is used by more than 500 academic and technical computing sites, can be installed on a 128-processor system in as little as eight minutes with the aid of BitTorrent file-sharing software.

To move into the scientific and technical computing world, Microsoft will have to overcome several obstacles, Papadopoulos said.

"Most users are Unix-friendly; that's the environment they work in," he said. If Microsoft wants to move scientists and engineers into its software environment, he said, it will have to demonstrate compatibility and prove that it offers an easier environment in which to do parallel programming.

Microsoft executives said they were still refining pricing for the Computer Cluster Server software, but that prices would be similar to those for the company's multiprocessor server system.

###

San Jose Mercury News
Biloxi Sun Herald
Bradenton Herald
Centre Daily Times
Columbus Ledger-Enquirer
Duluth News Tribune
Grand Forks Herald
Kansas City Star
Kentucky.com
Macon Telegraph
Monterey County Herald
Myrtle Beach Sun
Pioneer Press
San Luis Obispo Tribune
SiliconValley.com
Tallahassee.com

Leaders of supercomputing industry gathering in Seattle

By Therese Poletti
Posted on Mon. November 14

Makers of the world's biggest, fastest computers are gathering this week in Seattle for their annual conference, where they'll talk about how their systems are used in everything from conducting virtual autopsies to designing Pringles potato chips.

They'll also discuss major trends, such as the emergence of China in supercomputing, the ongoing shift to lower-cost systems for tackling the world's most sophisticated computing tasks and the decline in U.S. research spending.

Supercomputers are massive computers or bunches of smaller systems networked together to solve scientific or commercial problems that take a huge amount of computing power. These can range from weather forecasting to oil exploration to designing products, such as baby diapers or automobiles. A supercomputer, for instance, was used to study how the aerodynamics of a Pringles potato chip affects its path along a conveyor belt.

Supercomputing, sponsored by the Association of Computing Machinery and the IEEE Computer Society, started Saturday and runs through Friday.

Today at the show, the $7.2 billion industry's list of the top 500 fastest supercomputers is being released. The latest results offer further evidence of the shift away from the huge supercomputers of yesteryear toward clusters of lower-cost systems designed around chips from Intel and Advanced Micro Devices. According to the researchers who compile the list, Intel chips are at the heart of two-thirds of the top 500 systems. AMD's Opteron was in 55 systems, vs. 25 systems a year ago.

“What's happened in the industry is the shifting toward lower-margin products," said Dave Patterson, president of the ACM and a well-regarded computer science professor at the University of California-Berkeley. ``It's a great deal for the customers, but with less margins, companies are doing less long-term research on their own."

Analysts said IBM is an example of a company still able to invest in research involving high-performance computing, because it has other more profitable businesses that can help subsidize its longer-term research.

IBM took the top three spots on the list, with its mostly custom-designed systems using IBM's PowerPC and Power5 processors.

The No. 1 fastest supercomputer was once again IBM's BlueGene/L system, a joint development between IBM and the Department of Energy's Lawrence Livermore National Laboratory in Livermore. The system has doubled in the last two months, as IBM added more processors. It reached a new record of 280.6 teraflops, or 280.6 trillion operations per second.

The Columbia system, developed by SGI in Mountain View for the National Aeronautics and Space Administration, fell from the No. 3 slot six months ago to No. 4. The Columbia was developed with Intel Itanium chips.

The number of systems based on Itanium chips fell by almost 50 percent, while the number of systems using Intel Xeon chips increased.

``The list turns over every six months," said Jason Waxman, a director in the server platform group at Intel in Santa Clara. ``We continue to see Itanium going very strong in supercomputing."

Waxman said the large presence of Intel chips on the overall list “is a clear testimonial…that Intel-based supercomputers are the future for the scientific community."

In a first at Supercomputing, Bill Gates, the chief software architect and co-founder of Microsoft, will give a keynote address Tuesday morning.

Gates is expected to share his thoughts on the impact high-performance computing will have on science and engineering, and the world at large. He will also talk about how high-performance computing is being redefined as scientists wrestle with growing amounts of data and the need for better collaboration.

Analysts said they expect Microsoft to introduce a “supercomputing friendly" version of its Windows server software for clusters, in a belated move to fend off Linux. A Microsoft official declined to comment.

“It's important for Microsoft to gain a foothold in the high performance computing sector, because . . . over time, these (high-performance computing) developments start spilling over to the commercial world," said John Abbott, an analyst with 451 Group.

About 9,000 attendees are expected at the conference.

###

The Seattle Times

Microsoft supersizes Windows

By Brier Dudley
Seattle Times technology reporter

Tuesday, November 15, 2005

Microsoft made its fortune selling the equivalent of Fords, not Ferraris, but today it's wading into the rarefied world of supercomputing with a new version of Windows for managing massively powerful computing systems.

Chairman Bill Gates is introducing a test version of a new product, Windows Compute Cluster Server 2003, at the SC\05, a supercomputing conference held this year in Seattle.

Among the 8,600 people at the weeklong event are vendors and scientists who use computers to solve enormously complicated problems, including analyzing proteins, predicting the effects of nuclear explosions and modeling what might happen if an asteroid exploded near Earth.

At this conference, in the Washington State Convention and Trade Center, Gates is also expected to announce support for university research in the field and describe experiences of customers testing its new software, including Rosetta Inpharmatics in Seattle.

Microsoft sees a growing market at research-intensive companies for clustered systems, created by linking together a series of standard PC components to multiply their processing power.

Clusters now account for about 10 percent of all server sales, but Microsoft expects them to grow dramatically as prices continue to fall, the systems become simpler to manage and new applications are developed.

"What we're seeing is, as the price comes down and advanced applications trickle down from academia, they're really being picked up by enterprise [businesses]," said Kyril Faenov, Microsoft director of high-performance computing.

But Microsoft is wading into a largely academic field where perhaps 80 percent of the systems use freely shared, open-source software. The company rarely reveals the inner workings of its software to customers.

It's also an open question whether the Microsoft-based systems are true supercomputers. Traditionally, the term has referred to cost-is-no-object, room-filling systems used mostly by governments and universities.

With costs falling, Microsoft envisions high-powered clusters proliferating in the business world and appearing in every department or work group in biotechnology, energy, financial services and other industries.

Microsoft engineers built demonstration systems for under $4,000 using hardware bought from local computer stores, according to Craig Mundie, a senior vice president who started Microsoft's push into high-performance computing two years ago.

Microsoft is also entering the field to prepare for the more powerful computers expected to be widely available to consumers in 10 to 15 years.

Today's laptops are more powerful than the $1 million supercomputers Mundie built in 1991 at his previous company.

Supercomputer performance has grown apace.

A widely used ranking of the top systems, released Monday, reported that the fastest system of all — IBM's BlueGene/L — doubled in power over the past year to 280.6 teraflops. A teraflop is a trillion calculations per second.

"I don't know about the definition of a supercomputer these days; it's a constantly moving target," said David Bernholdt, a conference speaker and senior research staff member at Oak Ridge National Laboratory in Tennessee.

Bernholdt said researchers are benefiting from a diversity of suppliers entering the market, and smaller systems such as those powered by Microsoft's software are likely to find plenty of buyers. But supercomputer is a term that refers to bleeding-edge systems.

Seattle-based Cray pioneered the field of supercomputing more than 30 years ago when its namesake founder set out to build the world's fastest computers.

A shift began about 10 years ago with the arrival of clustered computer systems that were less expensive to assemble, said David Patterson, a University of California, Berkeley, professor and president of the Association for Computing Machinery, the conference's co-sponsor.

"That's kind of the new wave; for some people, they can do the problems they want to solve with basically a lot of desktop computers," he said.

Because these new systems are based on standard components, researchers can often write software for the machines on their laptops or desktop computers, said Marc Hamilton, director of technology for global education at Sun Microsystems.

Sun is announcing a supercomputer deal with Tokyo Institute of Technology today.

Separately, Cray announced the U.S. Department of Energy's Sandia National Laboratories is increasing the size of its "Red Storm" system from 10,848 processors to 14,348 processors next year.

Red Storm, developed by Cray, has been used to simulate atmospheric conditions and predict the effects of a nearby asteroid's explosion.

###

The Seattle Times

Isilon eager to make its mark

By Kim Peterson
Seattle Times technology reporter

Isilon Systems Chief Executive Steve Goldman has high hopes for his company, to put it mildly.

"I think we could be the next big company in Seattle," he said.

That's some serious optimism, especially for a company that isn't profitable. But sales this year tripled from 2004, and next year Goldman expects them to triple again. Profit, he said, will follow.

Isilon is trying to make a splash this week with the supercomputing industry, which lies far outside the core customer base it has courted in the past. The industry is converging in Seattle this week for SC|05, an annual supercomputing conference and trade show that began Saturday and extends through Friday.

About 9,000 attendees are expected for the 18-year-old conference, which grows by about 1,000 attendees every year.

Supercomputing, a historically volatile business, is seeing new interest as software applications advance and the need for more powerful data crunching grows. The field involves the development of computers with massive processing power, which are used in areas such as land exploration and drug development.

One leading supercomputer maker, Cray, is based in Seattle.

SC|05 is sometimes considered to be the next Comdex, the now-defunct trade show that once was a must-attend for computer companies worldwide. Organizers aren't angling to fill that void, though.

"We're not as commercial as Comdex," said Kathryn Kelley, who oversees communications for the conference. "It still fills a very particular niche."

Microsoft also will be trying to make an impression with attendees, many who are hardcore loyalists to open-source software. Chairman Bill Gates is keynoting the conference for the first time.

Industry pillar Sun Microsystems is expected to announce the largest supercomputing deal in its history.

Isilon will work the crowd on a much smaller scale, introducing a new product and trying to convince attendees that high-powered computing also needs high-powered storage.

Media focus first

Isilon started in 2001 as a storage provider for the media and entertainment industries — partly because its founders had previously worked at digital-media company RealNetworks and it needed to carve out a niche to get started.

The storage business is dominated by large incumbents, including Sun and IBM.

Sports Illustrated and NBC took Isilon's storage system to Athens to store the video and digital images captured at the 2004 Summer Olympics. Sony Pictures Imageworks, a visual-effects and animation company, uses the system to store motion-picture files.

Now, Isilon is broadening its user base to other industries that need to store, and quickly access, large data files. Some of its newest customers are oil and gas companies that are developing complicated three-dimensional models with data from seismic surveys.

Such companies are continually exploring for undiscovered oil and natural-gas deposits. The search is made more urgent as the need for energy sources becomes more pressing, making the field a good niche for Isilon to pursue.

Isilon, which launched its first product in 2003, has doubled its customer base in the past six months to more than 100. Growth came as Isilon turned its focus to sales, recruiting more than 50 reseller partners in the last year.

As a private company, Isilon doesn't disclose its finances. It says it will reach profitability on the strength of the $20 million funding raised in May.

Isilon has raised $59.9 million to date. Primary investors have included Lehman Brothers Venture Partners, Focus Ventures and Sequoia Capital.

New product

Today, in conjunction with the show, the company is introducing a new product called the Isilon IQ Accelerator, designed to increase the performance of its data-storage products by ramping up the speed at which data can be stored and accessed.

It is also partnering with Microsoft to show how its clustered storage works with Microsoft's server systems.

Clustered storage offers connected "nodes" of storage, with the idea that smaller segments acting in concert would be more efficient than one large depository.

Storage hardware is a $13 billion market, with sales growing about 10 percent a year, said Stephanie Balaouras, an analyst who covers the storage industry for Forrester Research. When storage networking and software is added to the mix, the market rises to at least $20 billion, she said.

Isilon's biggest obstacle right now is convincing new customers it's a trusted company that won't fold in a few years, Balaouras said.

"It is a risk to go with an emerging company versus an incumbent," she said. "If that company's not around in two years, who's going to provide the storage support for the system you've invested in?"

Isilon executives acknowledge the challenge.

"The biggest barrier to growth is getting our name out," Goldman said.

He added that the company is best known in media and entertainment, a relatively small industry. With the help of new resellers, Goldman is hoping that will change.

"We're well past being a startup," he said. "A year from now, we'll be trying to go as fast as we can to keep up."

###

Seattle Post -Intelligencer

Windows gets super powers

Microsoft taking plunge into high-performance computing

By TODD BISHOP SEATTLE POST-INTELLIGENCER REPORTER

Tuesday, November 15, 2005

Microsoft Corp. is moving into the world of high-performance computing with a new version of Windows that aims to make the ultrapowerful systems far more common in businesses and research institutions.

This morning at SC/05, an international supercomputing conference in downtown Seattle, Bill Gates will lay out Microsoft's strategy for entering the high-performance computing market with a program called Windows Compute Cluster Server 2003.

The program increases the competition between Windows and open-source rival Linux, which currently leads the high-performance computing market. But rather than simply trying to take market share from Linux, Microsoft says it will try to expand its potential market by increasing the use of high-performance computing overall.

"It should just be available to a lot more people," said Craig Mundie, Microsoft senior vice president and chief technical officer, describing the company's strategy in an interview.

High-performance computers are fast and powerful machines used by researchers and companies to work with huge amounts of data. The fastest are known as supercomputers.

Microsoft's Windows Compute Cluster Server is designed to run integrated collections of computers operating together as high-performance systems. The company will release a second test version of the program today. A final version is due out in the first half of next year.

Cluster computing makes up about half the high-performance computing market and is "clearly one of the higher-growth areas," said Bruce Toal, marketing and solutions director for Hewlett-Packard's high-performance and technical computing group.

One of Microsoft's goals, Mundie said, is to make such systems easier to operate by offering an integrated package that includes and connects to many of the same administrative features that many companies already use to run Windows PCs and computer servers. Examples include security, automatic updates and user authentication.

Microsoft says it envisions companies using its high-performance computing software in such areas as manufacturing, geosciences, life sciences, oil and gas and financial services.

At the conference today, Hewlett-Packard plans to announce that it will offer high-performance computing systems running Windows as an option for customers alongside Linux-based systems. Microsoft says it is also in talks with other major manufacturers of high-performance computing systems.

Microsoft also is showing a series of major high-performance computing applications that already have been adapted for the new Windows version.
But a separate announcement by HP underscores the position of Linux as a major force in the market. The company will base a new high-performance computing product, the HP Scalable Visualization Array, on the open-source operating system.

Linux is "one of the things that Microsoft has to contend with more in high-performance computing, because it's so well-entrenched," said analyst Carl Claunch, a Gartner Inc. research vice president. However, he said, Microsoft will benefit from the fact that many of its target customers already are running other Windows systems.

Microsoft hasn't said how much it expects to charge for Windows Compute Cluster Server. The Linux operating system is available for free. As with its main Windows Server software, however, Microsoft says it is aiming for a lower "total cost of ownership," taking into account the cost of such things as operation and maintenance. About 100 people at the company are working on Windows Compute Cluster Server 2003. But that number could grow, said Zane Adam, a director in the Windows Server division. "We'll invest what's needed to make this happen," Adam said.

Microsoft's Mundie led a supercomputing company called Alliant Computer Systems before joining Microsoft in the early 1990s.

About two years ago, Mundie said, he began encouraging Gates to move the company into high-performance computing, based on what he was seeing in the market. High-performance computing systems, once confined to high-end research labs and large businesses, have become less expensive and more common in recent years, as computer hardware costs have dropped.

And although proprietary systems still exist, more are being built on the standard AMD and Intel microprocessors that Microsoft is accustomed to working with.
Microsoft cites data from the IDC research firm showing a 70 percent increase in high-performance systems last year.

"The volumes have gone up in absolute terms to the point where it's an interesting business," Mundie said. "Many of the attributes that have been a part of Microsoft's

historical strength, in distribution and support, now have relevancy to this market."

Kyril Faenov, Microsoft's director of high-performance computing, compared the current growth in the field to past trends in markets where Microsoft has experienced success.

"We're going through an inflection point in this market very similar to what business and personal computing have gone through over the last couple of decades," Faenov said. Microsoft's challenge, he said, is now to make high-performance computing "digestible and usable" by the much larger number of users that can now afford it.

SC/05 IN SEATTLE

Supercomputers are the fastest and most powerful in the world, used by large companies and research institutions to perform complex modeling and analysis with huge amounts of data. The broader field is known as high-performance computing. More than 8,000 people from the industry are in Seattle this week for the SC/O5 conference.

Microsoft is seeking a bigger role in the industry with a Windows version designed for high-performance computing clusters. Although most of those systems won't technically qualify as supercomputers, Microsoft says it wants to ultimately demonstrate the potential of Windows Compute Cluster Server by running a system on the vaunted list of the world's Top 500 supercomputers.

A new version of that list is being released in conjunction with SC/05 this week: **www.top500.org**.

P-I reporter Todd Bishop can be reached at 206-448-8221 or toddbishop@seattlepi.com.

###

Seattle Post-Intelligencer

Software Notebook: Computing power is star of event

By TODD BISHOP SEATTLE POST-INTELLIGENCER REPORTER

Monday, November 14, 2005

Microsoft's Bill Gates will headline a major computer conference in Seattle this week -- but for some in attendance, the real technological star of the event will be the infrastructure.

The international supercomputing confab, SC/05, will bring together huge amounts of computer processing power and storage, running on what the conference leader calls "the world's most intense network" -- with a capacity of nearly half a terabit a second, or 100,000 times more than the fastest home connections.

Researchers and experts in high-performance computing will use all that power and capacity to conduct demonstrations and scientific experiments on a level not possible on their own. Examples include a massive analysis of data from telescopes, seeking to better understand how the universe will evolve. Researchers will also take on earthly challenges -- in such areas as automobile traffic and music information retrieval.

The experiments will take advantage of the ability to have "so much networking, so much storage and so much computing in one place, along with the expertise to do the analysis," said Bill Kramer, the SC/05 conference chairman and general manager of the National Energy Research Scientific Computing Center.

The idea is to anticipate the level of high-performance computing and networking to be available at individual research facilities several years from now, as technology advances, he said.

"Once a year, the community comes together and we build unique things, and we can do stuff that we can't do in any one normal place," Kramer said. "It's an interesting conference that doesn't just exchange papers. It tries to push the technology."

Crews have been working for weeks to build the infrastructure for the conference, which began Saturday and runs through Friday at the Washington State Convention and Trade Center. More than 8,000 people are expected to attend the conference.

A supercomputer is a machine that leads the world in processing power when it's introduced. The broader discipline is known as high-performance computing, and Microsoft is moving into the lower end of the spectrum with a version of its Windows operating system designed for collections of computers that work together. The program, known as Windows Compute Cluster Server 2003, has been released in test form.

Earlier this year, Microsoft Chief Executive Steve Ballmer acknowledged that high-performance computing is one of the areas where Microsoft has fallen short of open-source rival Linux.

Gates is scheduled to speak at the event Tuesday morning on "the impact high-performance computing will have on science and engineering, and therefore the world at large," according to the conference agenda.

In addition to the technical portions of the conference, it's also a place where companies involved in high-performance computing meet with customers and develop their businesses. Cray Inc., the Seattle-based supercomputer company has reserved multiple meeting rooms at the downtown Seattle Sheraton that are "booked literally solid for three days straight" for meetings with customers and partners, said Steve Scott, Cray's chief technology officer.

The annual Supercomputing conference, which began in the late 1980s, "has really grown into this mega-conference," with a big exhibit floor, technical sessions, business briefings, workshops, panel sessions and other events, Scott said.

At the conference today, Scott is scheduled to give an update on Cray's efforts as one of three companies working under a grant from the U.S. Defense Advanced Research Projects Agency, or DARPA, to develop a new generation of high-performance computing systems.

The company's work on the DARPA project is on the upper end of the high-performance computing spectrum, focusing on specially designed and integrated supercomputer systems.

"Commodity processors clustered together, while they can perform very well on certain types of problems, don't perform well on the most challenging sorts of problems," Scott said. "The most challenging problems really need systems that are custom-built, purpose-built for high-performance computing."

Cray has headquarters in the Pioneer Square neighborhood, with facilities in Minnesota, Wisconsin and British Columbia. The company this year cut its work force by 90 people, or 10 percent, as part of a bid to reach profitability. Its customers include companies such as golf club manufacturer PING Inc. and big science facilities such as the Oak Ridge National Laboratory.

In its discussions with customers this week, Cray will focus on how to make high-performance computing systems more productive.

At the conference, attendees using their laptops will also benefit from the robust infrastructure -- which will be capable of delivering wireless access to 8,000 people simultaneously, conference organizer Kramer said. But future users of the convention center need not waste time dreaming about possibilities. The network will be disassembled after SC/05 ends.

###

Silicon.com
ZDNet UK

Microsoft backs UK supercomputing move

Teams up with Southampton University...

By Dan Ilett

Published: Tuesday 15 November 2005

Microsoft is backing a UK universities project to research new ways of cutting the cost of supercomputing.

High-performance computing is currently expensive because of the high costs associated with processing power and custom-built applications.

The University of Southampton is today launching the Microsoft Institute for High Performance Computing, which will look at cheaper ways of using supercomputing in areas including the aerospace, automotive, marine and telecommunications sectors.

Jeff Price, a senior director for Windows Server at Microsoft, told silicon.com: "It's a market we don't do much in today. We think it's our job to make it less expensive and easier to do. We're spending a lot of money in trying to enable more research here."

The move is part of a larger Microsoft research and development plan to create 10 supercomputing centres around the world, Bill Gates will announce later today.

Speaking at Supercomputing 2005 in Seattle, Gates, the chairman and chief software architect at Microsoft, will outline his plans to make high-performance computing more widely available.

Microsoft is also promoting faster computing speeds using 64-bit processing and virtualisation technologies at the company's IT Forum conference in Barcelona today.

Professor Simon Cox of Southampton's School of Engineering Sciences is leading the project. He was unavailable for comment - he is attending the conference in Seattle - but a source close to the partnership said the university will also be working on computational algorithms for researching Antarctica, coastal behaviour and evolutionary biology.

The University of Stuttgart in Germany is also taking part in the scheme, creating a virtual power plant, including a customised water turbine. The goal of the project is to learn how to develop virtually without having to create prototypes.

Seattle Post-Intelligencer

Mundie on supercomputing

Todd Bishop's Microsoft Blog @ SeattlePI.com
http://blog.seattlepi.nwsource.com/microsoft/archives/100530.asp

Microsoft's move into high-performance computing is a homecoming of sorts for Craig Mundie, the chief technical officer who ran a supercomputing company before joining Microsoft in the early 1990s.

A couple years ago, Mundie saw trends in the high-performance computing industry that led him to encourage Bill Gates to move Microsoft in that direction. Gates will outline the company's plans this morning, at the SC/05 supercomputing convention. Among other things, the company will demonstrate Windows Compute Cluster Server 2003 to the attendees. See this morning's story for more details.

About 100 people are working on the Compute Cluster Server at Microsoft. Speaking with Mundie yesterday, I asked him to respond to the question of why Microsoft is expanding into another new area, rather than putting those resources toward making Windows for PCs better and more secure. Here's what he said:

"I think we feel like we're investing appropriately in that tack. But it does come back to one of the three points I made about why we think this market has interest that goes beyond just the absolute size of the server software business. And that's the fact that we do expect a fairly significant evolution in the structure of computers that will affect both the traditional desktop and server markets over a period of the next five to 10 years. And many of the attributes of these high performance computing environments are going to continue to shrink down and they'll ultimately show up as what it looks like to be one machine, instead of a whole data center or a whole cluster. And so there is this additional benefit of being able to both work with the community and, for us, to experiment with how we should adjust the software such that when the data center shrinks to the desktop [the company will know] the right way to maintain it and operate it and get benefit from it. You have to have a long-term perspective, but we think that is another benefit of the company playing in this space now."

The Wall Street Journal

IBM Technology Gains in Fastest Supercomputers

By Don Clark

14 November 2005

Supercomputers, a name applied to the biggest machines, have expanded from military and scientific applications to play major roles in commercial fields such as product design and film production. In one sign of their popularity, Microsoft Corp. Chairman Bill Gates is scheduled to speak at a supercomputer conference in Seattle tomorrow.

Bloomberg

Microsoft Plans High-Performance Computer Software, NYT Says

By Dan Hart

2005-11-15 05:54 (New York)

Microsoft Corp. plans to offer software that will give scientists and engineers a way to gain high-performance computing from their existing desktop computers, the New York Times reported.

Microsoft is set to release its Windows Computer Cluster Server 2003 software in the first half of next year, which is intended for systems that have as many as 64 processors and can be extended to larger machines, the newspaper said.

The software developer will be competing with systems that run Linux or Unix software, the newspaper said.

Microsoft Chairman Bill Gates will begin the marketing effort with a keynote speech today at the annual supercomputing trade show in Seattle, and the company plans to provide support for 10 supercomputing institutes worldwide, including Cornell University and the University of Utah, the newspaper said.

The Wall Street Journal

IBM Technology Gains in Fastest Supercomputers

By Don Clark

14 November 2005

A new list of the world's fastest supercomputers shows the growing popularity of technology from International Business Machines Corp. to solve the toughest scientific problems.

The ranking also shows wider use of Intel Corp.'s Xeon microprocessor, particularly for systems called clusters that use components that evolved from personal computers. But the number of systems on the list using Intel's high-end Itanium chip declined to 46 from 84 in the past year. Advanced Micro Devices Inc.'s Opteron chip was used in 55 machines, up from 29 a year ago.

Those are among the trends in the latest ranking of the 500 biggest systems in the world, issued twice a year by researchers at the University of Tennessee, Lawrence Berkeley National Laboratory and the University of Mannheim, Germany. Systems are ranked by how fast they complete a set of calculations.

Supercomputers, a name applied to the biggest machines, have expanded from military and scientific applications to play major roles in commercial fields such as product design and film production. In one sign of their popularity, Microsoft Corp. Chairman Bill Gates is scheduled to speak at a supercomputer conference in Seattle tomorrow.

Such systems also are a matter of national pride. After Japan's NEC Corp. took the No. 1 spot several years ago, companies and U.S. government agencies mounted an effort to regain the lead.

A collaboration between IBM and Lawrence Livermore National Laboratory helped to produce BlueGene/L, a supercomputer used to simulate nuclear reactions. That system, recently expanded, reinforced its No. 1 position in the latest list by executing 280.6 trillion calculations per second, doubling its own speed record.

IBM, besides selling more supercomputers than any other vendor, more than doubled the number of Power chips used in systems on the list. Some 327,622 of those chips were used in 73 supercomputers, including five of the 10 fastest machines. By comparison, Intel boosted the number of Xeon chips used by 57% to 243,112, in 287 systems, with one Xeon-based system and one Itanium machine among the top 10.

Government Computer News (GCN)
Washington Technology

Microsoft joins supercomputing fray with new server software

11/15/05
By Joab Jackson,
GCN Staff

With today's beta release of Windows Compute Cluster Server, Microsoft Corp. is hoping to enter the high-performance computer market—or at least the lower reaches of that market.

The Redmond, Wash.-based company is pitching the operating system to engineers and researchers "who would like some power, but don't want to have to take their computer jobs to a supercomputer center," said John Borozan, group product manager for the Windows server division. By running this operating system on groups of servers, offices and departments could run small supercomputing clusters.

A departmental implementation of a high-performance computer (HPC) cluster may run about $50,000 to $250,000, Microsoft estimates. The cluster sizes could range from four to 64 nodes. According to Borozan, a typical four-node system may yield about 10 billion floating-point operations per second (GFLOPS) of computational power, depending on the hardware used.

The Windows Cluster Server is based on the Windows Server 2003 server software, but has additional features that allow computers to be yoked together to work in parallel on computationally intensive tasks.

The package includes an open-source message passing interface (MPI) library for passing data among multiple machines. The MPI library, called MPI Chameleon or MPICH2, was developed by Argonne National Laboratory. Microsoft added performance and security enhancements.

The server software will also come with cluster management software and a job scheduler to coordinate workloads.

Traditionally, Unix and Linux-based machines have dominated the supercomputer cluster space. As such, Microsoft is working with independent software vendors and universities to generate more HPC tools for its own platform. A number of HPC software vendors have ported their applications over to this HPC version of Windows, including MatLab from The MathWorks Inc., and Mathematica from Wolfram Research Inc.

Microsoft will offer a version of its Excel spreadsheet that can be run across a multi-node cluster, according to Borozan. Users can also build their own applications using Microsoft Visual Studio and the .NET framework.

Beyond the .NET languages, Microsoft is also working with third parties to build compilers for other languages favored in the HPC space, such as Fortran.

To pull data from Unix file systems, the Compute Cluster Server comes with a copy of

Microsoft Services for Network File System, which allows Microsoft Windows machines to access NFS-based drives.

On the hardware side of things, Windows Cluster Server can run on x86 64-bit microprocessors—specifically Intel's EM64T Xeon and Pentium chips and AMD's Athlon and Opteron. It will not run on Intel Itanium processors, Borozan said. Out of the box, the software will work Gigabit Ethernet and Infiniband interconnects.

The full production release of the server software is scheduled for the first half of 2006.

###

SEATTLE POST-INTELLIGENCER

Gates addresses supercomputer experts

Wednesday, November 16, 2005

By TODD BISHOP SEATTLE POST-INTELLIGENCER REPORTER

In the realm of unlikely places for Bill Gates to speak, a supercomputing conference might not top a convention of open-source software fanatics, but it's definitely on the list.

The Microsoft chairman and personal-computing icon stirred up a mix of curiosity and skepticism Tuesday among the crowd of high-performance computing experts who listened to him outline the company's plans to expand into the field.

"It may seem strange to have somebody who works at the very lowest end of computing here addressing the supercomputing audience," Gates acknowledged at the outset of his speech at the SC/05 convention in Seattle.

In fact, the reaction from some attendees suggests that Microsoft has some perceptions to overcome before it is fully accepted into the supercomputing world.

The field, in which research labs and companies use ultrapowerful computers to work with huge amounts of data, has long been the domain of the Unix and Linux operating systems.

Jill Gemmill, assistant director in the University of Alabama-Birmingham's Academic Computing department, said she had been surprised the night before to see banners touting Microsoft's first high-performance computing product, Windows Compute Cluster Server 2003.

Pitching her on the product, the people at Microsoft's booth pointed out the software's benefits -- such as integrated tools for rapidly deploying security patches, Gemmill said.

"I just smiled and said, 'But if I wasn't using your operating system, I wouldn't have that problem,' " she said. Even so, she said the vision outlined by Gates was "kind of interesting."

Microsoft's plan is to expand the market for high-performance computing by making the machines easier for a broader range of businesses and research institutions to operate.

During his speech at the conference, Gates predicted that there will be "supercomputers of all sizes." An accompanying slide referred to it as the "rise of the personal supercomputer."

Kyril Faenov, Microsoft's director of high-performance computing, demonstrated the concept for the crowd by using a "Personal Compute Cluster," running on Windows Compute Cluster Server and housed in a casing small enough to sit next to a desk.

The deskside machine let Faenov conduct an initial analysis of a large set of blood samples used in cancer research.

But for a more intense analysis, he tapped into a larger system that included a computing cluster backstage, running on Linux, and an even larger cluster, running on Windows and accessed over a network from Intel Corp.'s facility in DuPont.

"The arrival of lots of low-cost microprocessors (and) the ability to connect those together with a fairly high-speed data network has led to an exciting new way of thinking about running these high-end applications," Gates said.

Gates said Microsoft expects it to be the norm for Windows and other operating systems to work together in high-performance computing systems, as shown in the company's demonstration.

"Microsoft wants to play a role here -- to be a participant and work with partners to see how our software fits in these solutions," he said.

"These solutions will often be extremely heterogeneous. Unix and Linux have been very important in this environment. ... All these things are going to work together."

More than 9,000 people are attending the SC/05 conference this week, and many of them crowded into a large assembly hall at the Washington State Convention and Trade Center to hear Gates speak.

During a question-and-answer session afterward, attendees asked him about topics ranging computer security to national energy policy.

Even though it is new to the field, Microsoft may be able to apply some of the lessons it learned in personal computing.

The Washington, D.C.-based Council on Competitiveness commissioned a study that examined high-performance computing software programs and found, among other things, that the on-screen interface made them difficult to use, said Melyssa Fratkin, policy director for the council's High Productivity Computing Project.

That's an area where Microsoft may have an edge.

"The user interface may be the one thing that they have that is going to get them into a space that they wouldn't ordinarily be in," Fratkin said.

"People who have never used it before can look at the interface and say, 'Oh, I know how this works. This is Microsoft.' "

###

IT Jungle

Gates Lays Out Vision of Future of Supercomputing

by Alex Woodie
November 16, 2005

Microsoft chairman and chief software architect Bill Gates shared his vision of the future of supercomputing yesterday during a keynote event at the SC 05 supercomputing conference in Microsoft's hometown of Seattle. While the era of the sub-$10,000 desktop supercomputer cluster is upon us, users still need better software to harness the computing power, and to consume its results. Microsoft, which also unleashed a public beta of Windows Compute Cluster Server 2003 yesterday, says Visual Studio and Office may fit the bill.

Gates started his keynote address by talking about the challenge of parallelism, which he says is "becoming a very important challenge, particularly as the clock speed of the microprocessor won't be increasing at the rates that it did in the past." When clock speeds top out at around 6 GHz to 8 GHz, "we'll have to bring in techniques around parallelization and those are techniques, of course, that the high-end has been working on and understanding for a very long period of time."

Gates commented on a tantalizing feat of having the world's entire body of knowledge at ones' fingertips. "As [computers] become cheaper, smaller, as the screens become better, we can think of even something like reading moving over to the digital realm and so that instead of having a library of papers, simply by being connected to the Internet through wired Internet or wireless Internet you can browse the full library of knowledge and annotate that and communicate relative to that in a far better way than paper systems would have allowed in the past."

In the future, Tablet PCs will enable users to interact with computers not just through the keyboard, but also through speech, Gates says. "And we'll have lots of cameras so that we can see what's going on, understanding whether you're paying attention to the computer or talking to someone else," he says.

In fact, one could say we're moving toward the "digitization of everything," Microsoft's founder says. "A lot of compute clusters now are doing things like simulating new product designs. A new plane, a new car will be done on a digital basis, and there should be a deep understanding of what happens in a crash or how does it wear over time by using these digital models instead of having to figure those things out in the physical world."

The ability of scientists to access the right data and apply it the right way will be critical for many scientific disciplines, including astronomy and even social sciences. "The life sciences with genomic and proteonomic data, it's an explosion of information that people around the world should be able to benefit from the results that other people are gathering

together," he says. "And so in the same way that mathematics has been a tool of all the sciences, in the years ahead we can think of computer software that manages lots of information, that lets you mine that information, that will likewise be a very critical tool."

Of course, it's the role of supercomputers to break up those giant accumulations of data, and to make some sense out of it. Microsoft is getting into the supercomputer business through a new release of its operating system called Windows Compute Cluster Server 2003, which was released for a public beta test yesterday, as well as through working with ISVs, OEM providers of supercomputers, and HPC centers at universities around the world (see "HPC Version of Windows Server Goes to Public Beta").

"We'll have supercomputers of all sizes," Gates says, including some that will cost less than $10,000 and sit at your desk. These "desktop clusters" will be good for getting preliminary results for simple problems, but will also be able to communicate with bigger supercomputers. "That means that when you want to take that same computation and do it with a finer level of detail, submitting it off to a cluster that's dramatically larger and can do more will be extremely straightforward." What's more, since all this data will be in XML formats and carry metadata tags, it will be a relatively simple matter to see where that data came from and its context, Gates says.

While supercomputer hardware that can do tremendous amounts of processing has become very affordable, and the data format (XML) has basically been standardized, Gates says the actual software used in HPC computations could use some work. "We need to take the techniques that have been developed for things like business intelligence and data mining . . . [and] apply those in these realms as well." This software should be very visual, very easy to use, and should "only require as much software understanding as is absolutely necessary."

Gates mentioned the upcoming new version of Office and an upcoming server version of Excel that's multi-threaded and much more scalable, as examples of the kind of software that could be used to make HPC applications more visual.

At the same time users get new ways to visualize their supercomputer programs and data, there should be pains taken to ensure interoperability with other platforms. "These solutions will often be extremely heterogeneous and of course Unix and Linux have been very important in this environment," Gates says. "Making sure that all these things are going to work together . . . is just one element of how software can do a better job for scientists and designers."

Microsoft, which no doubt employs several desktop clusters at its Redmond, Washinton, headquarters, is very good about providing timely transcripts of its executives' speeches. A transcript of Gates speech at the SC05 event yesterday is available on his Web site.

###

IT Jungle

HPC Version of Windows Server Goes to Public Beta

by Alex Woodie

November 16, 2005

While Microsoft is a relative newbie on the competitive high performance computing (HPC) circuit, the software giant is determined to gain some respect and to show that it's serious about HPC. It took a big step in this direction at a supercomputer conference being held in Seattle this week, where it issued a public beta of Windows Compute Cluster Server 2003, and outlined its ambitions for the HPC market.

The list of the world's 500 biggest supercomputers used to be dominated by proprietary RISC and Unix-based systems from the likes of Cray, SGI, Hewlett-Packard, and IBM. These supercomputers were tremendously expensive, which meant they could only be used by the biggest governmental organizations tackling the most important problems, such as simulating nuclear explosions.

In recent years, computer technology has evolved considerably, and the list of the 500 biggest HPC systems, which is updated every June and November at www.top500.org, has become increasingly composed of X86 and X64 clusters (Intel- and AMD-based systems now account for more than half), and the vast majority are now running the open-source Linux operating system. This shift toward inexpensive computers and free software has enabled less wealthy organizations to tackle a range of problems such as drug discovery and crash-test simulation.

Microsoft sees this shift, and wants its share of the glory, if not the business, which amounts to about $7 billion annually for HPC, just a fraction of the overall market. Microsoft's HPC ambitions were first revealed in May 2004, when it revealed it was working on a version of Windows Server for supercomputer-type workloads. A select group of beta testers got their first glimpse of Microsoft's HPC product, called Windows Compute Cluster Server 2003, at the Professional Developer's Conference in Los Angeles in September.

Yesterday marked the beginning of the public beta test for this product, which is expected to ship during the first half of 2006. The company also announced it would spend millions of dollars to set up 10 HPC centers at universities around the world, and highlighted some of the ISVs it is working with to provide the HPC applications that will run on its Windows clusters.

'Flops are Free'

Microsoft officials say the cost of supercomputers will continue to decline to the point where HPC systems will become much more prevalent. "HPC as a computing tool is growing and is no longer restricted to a niche high-end for academic research," says John Borozan, the group product manager in Microsoft's Windows server group in charge of HPC marketing. "The idea of personal supercomputing, or having a desktop cluster, is becoming much more of a reality."

We're on the cusp of a new era of HPC, according to Microsoft. In 1991, a 10 Gigaflop supercomputer from Cray would have cost $40 million, which put the system out of reach of everybody except the federal government. By 1998, the cost of a 10 Gigaflop system, such as Sun's HPC10000 computer, had dropped by a factor of 40, to about $1 million. Today, that 10 Gigaflop Shuttle system is available for about $4,000 online from Newegg, according to Microsoft.

The massive amounts of computing capacity available today is reflected in HPC system sales. According to IDC, the number of HPC installations grew by 70 percent in 2004, to about 150,000, and the growth was driven primarily by clusters that cost less than $50,000. This sub-$50,000 portion of the market will continue to grow at a healthy clip, and will eventually come to dominate the entire HPC market in terms of the number of installations, according to IDC. Larger HPC systems, including those costing between $50,000 and $250,000, and those costing between $250,000 and $1 million, will also experience growth, according to IDC charts.

Microsoft's plan is to target workgroup (sub-$50,000) and departmental ($50,000-to-$250,000) segments of the HPC market, with focuses on manufacturing, geosciences, life sciences, oil and gas exploration, financial services, and public market sectors.

Microsoft Targeting Linux for HPC

While the majority of HPC systems today run the Linux operating system on relatively affordable X86- or X64-based clusters, that's not to say the systems are easy to use, Borozan says. "Despite how economical these systems have become, they're still pretty tough to deploy. You have to be pretty savvy" in stitching everything together, he says.

Instead of requiring chemists or biologists to become computer scientists as well as experts in their own fields, Microsoft's goal with HPC is to make the deployment as easy as possible. "That's where we think we can remove a lot of the complexity for end users and developers," Borozan says. "That's the core part of our value proposition."

Microsoft says that by working with its ISV partners and OEM partners like Dell, HP, and IBM, it can do much of the work of building HPC systems beforehand, so users can get as close to a finished "turnkey" product as possible. To this end, Microsoft has

pledged to closely follow the existing standards in the marketplace, including Infiniband and Gigabit Ethernet inter-connects, use of Winsock Direct and OpenIB protocols, and the open-source MPICH2 messaging interface from Argonne National Lab.

Microsoft has done some work to make sure that MPICH2, which is used for dividing HPC workloads among the various nodes, works well with its Windows HPC variant. These security and performance enhancements have been contributed to the open-source community, Borozan says. The company is also planning to include a job scheduler (often one of the most expensive components of a cluster) with Windows Compute Cluster Server 2003. It has worked with some of the premier developers of job schedulers for clusters to get those to support Windows.

In terms of ISVs, Microsoft has identified 19 ISVs that account for 80 percent of the HPC application sales in its targeted markets, and partnered with them to bring their applications to Windows. These ISVs each have applications running on Windows now, and some will participate in benchmarks prior to the release of Windows Compute Cluster Server 2003 next year.

The ISVs Microsoft is working with include: Accelrys; Ansys; BioTeam; CD-Adapco; Dassault; ESI Group; Fluent; Landmark; Livermore Software Technology Group; MathWorks; MSC Software; Parallel Geoscience; Platform Computing; Schlumberger; the University of Illinois Urbana's NAMD group; and Wolfram Research.

Borozan says that while the initial upfront acquisition cost of a Windows HPC system might not be much different than competing Linux-based systems, over time, users will save time and money because it is easier to manage. This will be helped by the fact that Windows Compute Cluster Server Edition 2003 will work with Microsoft Operations Manager and Systems Management Server. Pricing for Windows Compute Cluster Server Edition 2003 has not been set, but don't expect it to be too high, Borozan says. Think more along the lines of the Standard Edition of Windows Server 2003 than the high-end Datacenter Edition.

Up to this point, Microsoft has worked very closely with Cornell University in Ithaca, New York, which has deployed Windows clusters and been on the leading edge of using Windows for HPC. Microsoft is now working with the HPC programs at nine other universities. These universities include: the University of Texas; the University of Utah; the University of Tennessee; the University of Virginia; Southampton University in the U.K.; the University of Stuttgart in Germany; the Nizhni University in Russia; the Tokyo Institute of Technology in Japan; and Shanghai Jiao Tong University in China.

###

ENT News

Gates: Microsoft Ships Beta 2 of HPC Cluster Version of Windows

by Stuart Johnston

11/16/2005 — SEATTLE – Not satisfied with a place in the data center, Microsoft is making a play for the laboratory, too. Chairman and chief software architect Bill Gates told a standing room only audience at the Supercomputing 2005 conference held here this week that it has released the second beta test version of its upcoming high-performance computing (HPC) edition of Windows.

Microsoft had originally hoped to ship Windows Computer Cluster Server 2003 by the end of the year, but now it's targeted for release in the first half of 2006. When it does ship, the software will only run on systems that support Intel's and AMD's 64-bit memory addressing technologies.

The system is intended for use in clusters of inexpensive machines that work simultaneously on the same problem. The server's base code is Windows Server 2003 Service Pack 1. Windows Computer Cluster Server will consist of two CDs -- a Compute Cluster Edition based on the x64 release of Windows Server 2003, and a Compute Cluster Pack consisting of the MPI layer, job scheduler, administrative console and management tools. The components can be used together or separately.

Microsoft has offered failover clustering to support high availability computing for years. However, this is the company's first foray into the world of high-performance computing. Gates also announced that the company is investing in ten HPC institutes worldwide. The multiyear, multimillion-dollar plan will fund work at U.S. universities Cornell University, the University of Tennessee, the University of Texas at Austin, the University of Utah and the University of Virginia. Additionally, it will fund overseas institutes at the Nizhni Novgorod State University in Russia; Shanghai Jiao Tong University in the Peoples Republic of China, the Tokyo Institute of Technology in Japan, the University of Southampton in England and the University of Stuttgart in Germany.

At the high end, the supercomputing world is increasingly dominated by massive clusters of inexpensive Linux servers. Microsoft's new investments may broaden its presence among those high-end systems, but the company's plan to make money from the technology involves much less grandiose applications.

"We see as a key trend here … that we'll have supercomputers of all sizes, including one that will cost less than $10,000 and be able to sit at your desk or in your department," Gates said Tuesday. Users will be able to employ such "personal supercomputers" for preliminary results or relatively simple problems. Architectural continuity between the personal or workgroup supercomputers and a much larger supercomputing cluster at, say, company headquarters could allow the same computation to be run with a finer level of

detail. "We need an approach here that scales from the smallest supercomputer that will be inexpensive up to the very largest," Gates said.

John Borozan, a group product manager in the Windows Server Division, says Microsoft sees a huge opportunity at the low-end of supercomputing. He points to IDC research showing 8 percent of all x86 servers are purchased for high-performance computing clusters today. "In spite of amazing growth in this market, it's still very difficult to build an HPC cluster. You need to get the hardware, you need to find interconnects, you've got to put an OS on there, you've got to put an MPI on there, a job scheduler," Borozan said.

"We're going to focus on the departmental and workgroup level. One of our acknowledged strengths as a technology vendor is bringing things mainstream and making them easy to use," Borozan said.

Just because Microsoft aims to bring clusters down from government labs into corporate departments doesn't mean Microsoft is settling for a small piece of the existing revenue pie. The company is placing a bet that high-performance clusters will become a major way that data is processed in the near future.

"We think the next big revolution in science and industry is going to be driven by data -- flooding in from sensors as well as more traditional data from computational models," Borozan said.

During his keynote, Gates provided an example of sensors that might one day feed a supercomputing cluster. The University of Washington's NEPTUNE project involves dropping thousands of sensors along the undersea fault off the West Coast of the United States to study "black smokers," which scientists think may hold clues to the origin of life on Earth. "We can see this in many of the sciences, that low-cost sensors that give us overwhelming amounts of data and yet that we want to control in real time, will be feasible," Gates said.

Scott Bekker contributed to this report.

Stuart J. Johnston has covered technology, especially Microsoft, since February 1988 for InfoWorld, Computerworld, Information Week, *and* PC World, *as well as for* Enterprise Developer, XML & Web Services, *and* .NET *magazines. Contact him at stuartj1@mindspring.com. You can contact Stuart about "Gates: Microsoft Ships Beta 2 of HPC Cluster Version of Windows" at stuartj1@mindspring.com.*

###

The Seattle Times

Colossal computers turn on technologists at SC|05

By Kim Peterson *Seattle Times technology reporter*
Wednesday, November 16, 2005

A crowd milled around Sun Microsystems' booth Tuesday hoping to score one of the most popular freebies at the SC|05 conference: a T-shirt with the phrase, "No, I will not fix your computer."

The shirt's popularity could mean two things: Either the swag was sadly lacking or this conference was one big mob of tech-heads.

A tour around the conference hall proved both to be true.

In its 18th year, SC|05 has become the most prominent annual event of the supercomputing industry. About 8,700 people are at the show, running through Friday at the Washington State Convention and Trade Center.

The conference was once the domain of academics and government researchers who worked with enormous computer systems known as supercomputers.

It has expanded and become more mainstream in recent years, as has the definition of supercomputing.

But that doesn't mean the conversations are dumbing down.

"We're hitting over a teraflop of sustained global bandwidth!" exclaimed John Levesque, a senior technologist at Seattle-based Cray, to a crowd of about 45 people.

As Levesque spoke at the supercomputer company's booth, Eric Rudder wandered by, alone and unnoticed. That Microsoft Chairman Bill Gates' top company adviser could attract so little attention at the show says something about the challenge Microsoft faces in an industry that embraces open-source computing.

Gates did appear at the show Tuesday, outlining in a keynote address Microsoft's vision of "personal supercomputing."

That sounds like an oxymoron, but Microsoft believes businesses will increasingly use ultra-high-performance clustered systems, as prices fall and new software makes it easier to manage and run the systems.

Gates said technologists on the high and low end of the computing spectrum can work together on common challenges, such as parallel computing.
He also suggested that products such as Microsoft's Visual Studio 2005 programming toolkit can help the supercomputing crowd simplify software development.

Showing his chops, Gates dove into the technical details of a University of Washington research project monitoring undersea activity along fault lines. The researchers are taking advantage of XML software language to collect and analyze data.

He also described how the UW, Fred Hutchinson Cancer Research Center and Microsoft Research are working together on methods for developing an AIDS vaccine. Outside the speech, the conference exhibits lacked the glitz of better-known tech gatherings, such as the annual Consumer Electronics Show in Las Vegas. Exhibitors mostly used poster boards with diagrams on them or PowerPoint presentations of grids and helixes.

At the Sun booth, a magician performed tricks for a large crowd. The company spares no expense in reaching these attendees, said Rich Brueckner, Sun's marketing manager for high-performance computing.

"A lot of them, this is the only show they go to and they are shopping for multimillion-dollar computer systems," he said. "This is the only [conference] that's targeted at this market."

SC|05 is unusual in that the exhibitors are also the customers. Some of the most sought-out buyers at the show were the research facilities and supercomputing centers that exhibited their work for the crowd.

Companies also checked out what rivals were up to.

"It's a lot of competitors talking to competitors about the state of the industry," said Alex Lesser, a vice president at PSSC Labs, a Lake Forest, Calif., company focusing on high-performance computing. Lesser displayed new water-cooling technology designed to absorb the internal heat created by powerful computing machines.

The San Diego Supercomputer Center offered drinks from an espresso cart and showed off its research. The center conducted an earthquake simulation last year that depicted the potential impact from a major quake originating near Palm Springs, Calif.

The simulation took four days to compute and created 43 terabytes of data, said Lynn Ten Eyck, associate director for scientific research and development at the center. A terabyte is 1,000 gigabytes.

The center is an exhibitor but also does careful comparison shopping at the conference, he said.

"We are high-profile customers," he said. "Our user base is a very large scientific community, and these people talk not only to one another but also to policymakers."

Numerous research institutions discussed their supercomputing projects. Indiana University showed how it is simulating physical systems to understand how geologic oil deposits are formed.

Researchers are also studying the way air flows over airplane wings.

At the Maui High Performance Computing Center, part of the Air Force Research Laboratory, scientists are studying combat models in battlefield simulations.

The University of Iowa is creating a digital human that performs virtual missions, such as maneuvering around obstacles.

The Ohio Supercomputer Center was touting an initiative called "blue-collar computing," which aims to make supercomputing more accessible to companies and researchers. The project is pushing the industry to review how supercomputing software is written and how hardware is designed.

Down the road, the center said, the large jobs of today could become routine tasks.

Following the blue-collar theme, the center's researchers wore blue shirts with their names sewn on patches.

"A lot of people come around here looking for an oil change," joked Paul Buerger, a systems and operations leader with the center.

Seattle Times technology reporter Brier Dudley contributed to this story. Kim Peterson: 206-464-2360 or kpeterson@seattletimes.com

###

SYS-CON Britain-UK

Gates @ Supercomputing 2005 – Announces Public Beta of Microsoft Windows Compute Cluster Server 2003

By: Staff Writer, News Desk
Nov. 16, 2005 12:30 AM

Microsoft Chairman and Chief Software Architect Bill Gates delivered the opening keynote address at Supercomputing 2005 Tuesday morning in Seattle to more than 7,000 attendees, sharing his vision of how the software industry can contribute to accelerating scientific research and engineering innovation, and calling for broad collaboration between the computing industry, academia and government to make technical computing easier and more productive.

In Gates's SC|05 Conference address, entitled "The Role of Computing in the Sciences," he also announced that, as part of the company's steps toward realizing this vision, Microsoft is funding joint research projects at 10 academic centers worldwide and has released the beta 2 version of Windows Compute Cluster Server 2003. Windows Compute Cluster Server 2003 is Microsoft's first product offering for high-performance computing.

"Technical computing is crucial to the many discoveries that impact our quality of life -- from making safer, more efficient cars and airplanes to addressing global health issues and environmental changes," Gates said. "Moreover, most sciences are becoming computational sciences, which is why advanced computing capabilities need to be seamlessly integrated into the end- to-end scientific process. We see many opportunities to collaborate with the scientific community on innovative solutions that will accelerate the pace of insight and discovery."

With the delivery of Microsoft Windows Compute Cluster Server 2003 Beta 2 now and general product availability scheduled for the first half of 2006, Microsoft will offer an HPC platform that accelerates the time to insight, Microsoft says. To help ensure simple integration of Windows Compute Cluster Server 2003 with a broad set of applications and tools, the company addd in a statement, Microsoft is working closely with software, application and hardware partners across the industry: "This collaboration is designed to meet customers' unique needs by enabling them to choose among and run a variety of compatible HPC applications. Microsoft is being joined in its booth at Supercomputing 2005 by 17 software partners and a number of hardware partners, including Dell Inc., HP, IBM Corp., NEC Corp., Orion Multisystems Inc., Tyan Computer Corp. and Verari Systems Inc., to demonstrate 20 industry-specific HPC applications."

Gates also announced Microsoft's multiyear, multimillion-dollar investment in joint research projects at 10 institutes, aimed at helping guide ongoing software research and product innovation at Microsoft to address the most challenging technical computing

problems. These institutes are:

- Cornell University (U.S.)
- Nizhni Novgorod State University (Russia)
- Shanghai Jiao Tong University (China)
- Tokyo Institute of Technology (Japan)
- University of Southampton (England)
- University of Stuttgart (Germany)
- University of Tennessee (U.S.)
- University of Texas at Austin (U.S.)
- University of Utah (U.S.)
- University of Virginia (U.S.)

"Easier access to workgroup and departmental HPC resources, such as that provided by Windows Compute Cluster Server 2003, will play an ever-increasingly important role in furthering scientific and engineering innovation in academia and industry," said Jack Dongarra, Ph.D., distinguished professor and director of the University of Tennessee Innovative Computing Lab. "Our research project will focus on self-adapting high-performance applications and software on the Windows platform, which could help scientists and engineers reduce the time and effort to discovery. I've found that Microsoft's passion for technology innovation around an integrated development and deployment software platform will help expand the potential for many who haven't had access to high-performance computing resources in the past," Dongarra added.

In the closing statements of his keynote address, Gates reiterated Microsoft's long-term commitment to working with the HPC community and emphasized his enthusiasm for the future role of computing in scientific discovery.

###

People’s Daily Online

Microsoft aims for super computing: Gates

November 16, 2005

Microsoft chairman Bill Gates forecasted on Tuesday the merging of the world's most powerful super computer with the Windows operating system.

Addressing the Super-Computing 2005 Conference in Seattle, Gates said he believes that as chips reach gigahertz speed limits, the need for "parallelism" becomes more important.

"Its exciting to think that we can get the best brains from super-computing and from mass computing and bring those together and make great progress in the decade ahead," he said.

He also sketched out an environment in which desktop supercomputers are linked to more powerful clusters, and a super-computing system will cost less than 10,000 US dollars.

"Microsoft wants to play a role here, to be a participant and work with partners to see how our software fits in these solutions, " he said.

"These solutions will often be extremely heterogeneous," he said, adding that making certain all these systems work together "is just one element on how software can do a better job."

The Microsoft chairman also said his company is reaching out to more super-computing centers to understand "what should we be doing with our software, how can it connect up to the other software that they have in a better way?"

By now, Microsoft's Windows has barely won a stand in high-end computing. It is not mentioned at all on the latest world's Top500 super-computing list released this week.

But experts predicted that clusters of cheap computers with the Windows operating system are almost inevitable in many respects.

The notion of "mass computing" is pretty evident especially as researchers turn to low-cost commodity clusters to solve problems, according to William Kramer, general chairman of the super-computing conference.

Kramer, head of high-performance computing at the US Energy Department's Lawrence Berkeley Laboratory, said Gates' appearance is an indication of the growing awareness of super-computing's importance.

"The output of (high-performance computing) activities are no longer hidden behind a curtain, if you will," he said. " Super-computing is being scaled down so more people can make use of these very complicated tools, and I think that's one of the indications of Microsoft's interest here."

Source: Xinhua

All Headline News

Microsoft Developing Supercomputer Software

November 16, 2005 11:00 a.m. EST

Hector Duarte Jr. - All Headline News Staff Reporter

Seattle, WA (AHN) – Software giant Microsoft Corp. is in the midst of developing software for high-performance computers commonly used in engineering and science research; a move that puts the company in head-to-head competition with open-source developers.

At a supercomputing industry conference on Tuesday, Microsoft Chairman Bill Gates announced the release of another test version of the company's software.
Craig Mundie, one of Microsoft's chief technical officers, says the company is entering the market at a time when high-powered computers are becoming more affordable, allowing for significant market expansion.

Such computers are used for computing, storing and sharing data for a range of uses. While these machines used to easily cost as much as $1 million, Mundie says some are now available for as little as $4,000 or $5,000.

Windows Server 2003 Compute Cluster Edition, is slated for release during the first few months of 2006. Microsoft has yet to set a price.

eWEEK

Microsoft to Add 64-Bit Computing to Key Products

November 15, 2005
By Peter Galli

Microsoft Corp. is pushing forward with its plan to embrace 64-bit computing across all its products and has decided that the next version of Microsoft Exchange as well as the upcoming Windows Server "Longhorn" Small Business Server and its Centro infrastructure solution for midsize businesses will only be released as 64-bit and optimized for x64 hardware.

Bob Muglia, senior vice president of Microsoft Server and Tools, announced this at the IT Forum being held in Barcelona, Spain.

He also said that while the first release of the upcoming next-version Windows Server family of products, code-named Longhorn, will be both 32-bit and 64-bit, the future update release to that, known as Windows Server "Longhorn" R2, will bring the complete transition to 64-bit-only hardware, while still benefiting from 32-bit and 64-bit application compatibility.

"We're betting big on 64-bit, but we believe in the right 64-bit for the job. In that context, Exchange 12 will be 64-bit only; Longhorn R2, Centro and our small-business Longhorn Edition will be 64-bit only," Bob Kelly, general manager of infrastructure server marketing at Microsoft, told eWEEK Tuesday.

Click here to read more about Centro.

"These are important transitions for customers, and we wanted to signal early so they can plan and that the right ecosystem changes occur to enable a full set of solutions around that," Kelly said.

There were also "tremendous" benefits to customers for leveraging the 64-bit hardware with Microsoft software, Kelly said, most notably that it enabled them to run 32-bit as well, giving an easy transition over time.

"So we have the best of both worlds and what we are essentially telling customers is that we are ready when you are," he said.

Asked about the hardware requirement that would be necessary for customers making the move to 64-bit computing, Kelly said that most new hardware available today was already x64 and that customers with legacy hardware would be able to run a mixed mode of 32-bit and 64-bit.

"But they will not be able to run Exchange 12 on 32-bit gear. This is an important leap,

particularly in the case of Exchange where the mail store requires massive scalability and the limits on memory have customers bumping into that," he said.

Click here to read more about Microsoft pushing hardware partners toward 64-bit.

Moving to 64-bit Exchange brought consolidation benefits and true scalability of the mail store, "so there is real benefit to the customer. Also, if the customer wants a transition plan we'll enable them with that," Kelly said.

The hardware industry agrees that, over time, 64-bit x86 servers will become ubiquitous as most machines on the market now contain 64-bit capable processors, which can run both 32-bit and 64-bit operating systems and applications.

Advanced Micro Devices Inc. lead the way and have been offering 64-bit capable Opteron server chips since April, 2003.

Intel followed in June, 2004 with its first 64-bit capable Xeon chip, code-named Nocona. Since then, server makers have added one or both of the 64-bit capable chips to nearly all of their model lines.

The jump to 64-bits is particularly useful for server applications such as databases, as it allows servers to make use of much larger amounts of memory without resorting to techniques such as memory windowing.

But one Microsoft beta tester questioned why Microsoft's SQL Server product was not on the 64-bit-only release. While the just-released SQL Server 2005 "Yukon" supports both 32-bit and 64-bit hardware, the tester wondered why the follow-on release, code-named Katmai, isn't being designated 64-bit only.

The tester, who requested anonymity, said that of any application, a database would be the type most apt to take full advantage of 64-bit processing. Microsoft IT Pro Evangelist Eileen Brown explained Microsoft's 64-bit Exchange strategy on her Web log.

"Exchange 12 (E12) will only be made available in 64-bit. Yes, I said only … The product team has been testing E12 on 32 bit and 64 bit, and have found some significant gains and a reduction of IOs per second, which results in really good performance gains. They tested Exchange on 64 bit and found almost a 75% reduction in IOs per second compared with Exchange 2003.

"This could result in almost a 4X increase in the number of users on the same disks or require 1/4 the disks to support the same users from a throughput perspective.

"If you think about this, it's quite significant, since it will proportionately decrease the investment in storage which accounts for 80 percent of the capital cost of Mailbox servers. And that's worth noting," Brown said.

Brown added that Microsoft is expecting to provide both 32-bit and 64-bit versions of Exchange 12 during the beta process, "so you can compare the performance gains."

"And as E12 will support mixed 32-bit and 64-bit environments including legacy Exchange 2003 and Exchange 2000 servers. You can also connect to either 32-bit or 64-bit Outlook," she clarified.

Muglia also used his keynote address in Barcelona to highlight the release of the second beta for Windows Compute Cluster edition 2003, stressing how this can be integrated with a customer's existing Active Directory environment.

Next Page: MS hopes to bring virtualization to the masses.

Muglia also announced that Virtual Server 2003 R2 has been released to manufacturing as well as pricing, which comes in at $99 retail for the standard edition and $199 retail for the enterprise version.

"This is a very important release for us and is key in our vision to bring virtualization to the masses," said Muglia.

Microsoft recently decided to christen Virtual Server 2005 Service Pack 1 as Virtual Server 2005 R2. The renamed product is still due to ship later this year, with the follow-on release due out in the latter half of 2006.

The company also recently decided to simplify Windows server licensing to allow for virtualization to become more pervasive.

Licenses for the upcoming Windows Server 2003 R2 Enterprise Edition, expected to ship later this year, will allow customers to run as many as four virtual instances on one physical server at no additional cost, extending the savings customers can realize through server consolidation on the Windows Server platform.

Click here for more on Microsoft's Windows Server licensing moves.

Muglia also discussed a number of technologies and value propositions that Microsoft saw as critical to driving forward its DSI (Dynamic Systems Initiative), its vision and technology road map for reducing the cost of managing and securing enterprise systems.

Muglia announced the December RTM of System Center Capacity Planner 2006, a tool for IT professionals to do performance analysis and planning of Exchange Server 2003

and MOM (Microsoft Operations Manager) 2005 system deployments.

He also covered a plan to build a midmarket offering called System Center Essentials, a new product that will bring the commitment of the Dynamic Systems Initiative to midsize companies, an underserved IT segment that was "suffering from a lot of pain of management and high complexity."

The new product will allow users to secure, update, monitor and track their IT environment and better support end users, said Muglia.

Microsoft also released at IT Forum a couple of solution accelerators for the Systems Center as free web downloads, specifically Desired Configuration Monitoring and the SLA Scorecard for Exchange. Kelly added that these were templates for customers to use to drive further cost out of their infrastructure.

Muglia also announced a preview of MOM version three, where technologies will allow proactive monitoring, analysis and managementof service levels "to drive an infrastructure view of the health rather than a boxed version of the health," Kelly told eWeek, adding that the product is slated to ship late in 2006.

Microsoft is releasing all of this information about future products and plans as part of its commitment to be more transparent about where it was going.

The company said it wants to give customers enough lead time to make the transition to 64-bit computing, as well as plan for the new features and functionality the company plans to deliver to the market, he said.

"This is a big set of news and enables us to help customers map out their next year. We also tend to have two big management conferences a year: IT Forum and the Microsoft Management Summit, the next one of which will be held in about six months time.

"You should expect more revelations there as to what the roadmap looks like so customers can clearly plan for these transitions," Kelly said.

Editor's Note: *This story was updated to include information and comments from Microsoft executives.*

Additional reporting by Mary Jo Foley and John Spooner

###

Internetnews.com

Microsoft's 64-Bit Support Rains In Spain

By Clint Boulton
November 15, 2005

The year may be starting to wind down, but Microsoft's deluge of last-quarter software releases is still going strong.

The company announced more server software updates and timetables for upcoming business applications during IT Forum 2005 in Barcelona Spain today.

With an emphasis on software support for 64-bit computing (define), company officials said Microsoft Exchange Server 12, the new Windows Compute Cluster Server 2003, Windows Server Longhorn Small Business Server, and Microsoft's Centro infrastructure will be exclusively 64-bit and tailored to run on x64 hardware.

Bob Muglia, senior vice president of Microsofts Server and Tools division, said customers will get a taste of this 64-bit hardware transition in a future release of Longhorn, R2.

The news follows its recent launch of SQL Server 2005, Visual Studio 2005 and Virtual Server 2005 R2, all of which are primed for 64-bit architecture.

Microsoft's increasing support of 64-bit computing is a sign of its faith in the future of the architecture, which boasts greater physical and virtual memory, among other improvements over its 32-bit ancestor. IBM, Dell, Sun Microsystems and HP are all pouring the 64-bit solutions on, making the transition an obvious choice for a major applications provider such as Microsoft.

Muglia also said Microsoft has released to manufacturing Virtual Server 2005 R2, which will be available the first week of December. Virtual Server 2005 R2, equipped with greater performance, will cost $99 as a standard edition and $199 for an enterprise version.

The company also announced the December release to manufacturing of System Center Capacity Planner 2006, which lets customers do performance analysis of Exchange Server 2003 and Microsoft Operations Manager (MOM) 2005 software.

Microsoft is also crafting a midmarket offering, called System Center Essentials, that helps medium-sized companies secure, update, monitor and track their IT environment.

Muglia also said Microsoft's new supercomputing software, Windows Compute Cluster Server 2003 Beta 2, is now public.

During the dual-keynote, Jeff Raikes, president of the Microsoft Business Division, detailed the launch of Microsoft Dynamics GP, business management applications that combine the capabilities of Microsoft's Great Plains assets with an interface that looks like the company's popular Office suite.

The fusing of business applications with Office tools is part of Microsoft's plan to offer customers a user-friendly experience and improve collaboration across the workforce.

Microsoft Dynamics GP features role-based user templates to provide information tailored to an individual's daily responsibilities; better business intelligence tools; new extranet capabilities; and tighter integration with the forthcoming Microsoft Dynamics CRM 3.0 through a new processing engine based on XML.

A standard edition of Dynamics GP 9.0. starts at $3,500 for up to four users, while the professional version starts at $7,500, also up to four users.

Lastly, Raikes unveiled a new version of its Windows Desktop Search tool for enterprise users.

Windows Desktop Search provides a single search starting point from which users can find information on their PCs, in e-mail, on networks, or across intranets and the Web.

###

InfoWorld

Customers get chance to test Windows for clusters

Redmond distributes second beta of Windows Server 2003 Compute Cluster Edition

By Elizabeth Montalbano, IDG News Service
November 15, 2005

Microsoft (Profile, Products, Articles) Tuesday at the Supercomputing 2005 show in Seattle will unveil several milestones in its strategy to becoming a serious competitor in the high-performance computing market.

The Redmond, Washington, software vendor will distribute the second beta of its Windows Server 2003 Compute Cluster Edition at the show, where Microsoft Chairman and Chief Software Architect Bill Gates will keynote Tuesday, said Kyril Faenov, director of high-performance computing at Microsoft.

Microsoft released Windows Compute Cluster Edition's first beta, aimed at developers, at the Professional Developers Conference in September.

Beta 2 of the product is aimed at allowing customers to test how the software will perform in their data centers, Faenov said. "It is for customers deploying evaluation versions of our product in their infrastructure," he said.

In Gates's keynote at the show, Microsoft also will reveal that it has set up 10 high-performance computing institutes around the world in universities to drive both research in those institutions as well as the company's own efforts to produce software that is appropriate for compute-intensive environments, said Zane Adam, director of marketing for Microsoft's Windows Server division.

"It's a two-way street," he said of the research centers. "Our platform is being used for a learning experience, and researchers are working on it so they can transfer their knowledge to us. It extends our research and development."

The centers have been set up at the following universities in the U.S.: University of Texas, University of Utah, University of Virginia, Cornell Theory Center at Cornell University and University of Tennessee. In the U.K., Microsoft has set up a high-performance computing center at Southampton University, and European centers have been established at the University of Stuttgart in Germany and Nizhni Novgorod University in Russia. In Asia, Microsoft has set up two institutes: one at the Tokyo Institute of Technology and one at Shanghai Jiao Tong University.

Microsoft has made a multiyear, multimillion-dollar commitment to providing software

for and helping direct research in these centers, Faenov said. Each center, which is run by its own director -- often a professor at the respective university -- is responsible for participating actively in conferences within the high-performance community, providing industry training, giving Microsoft feedback on its high-performance computing software and providing technology prototypes for the next version of Windows Compute Cluster Edition.

Microsoft also will showcase support from applications vendors and hardware OEMs (original equipment manufacturers) at the supercomputing show, Faenov said. Microsoft plans to demonstrate about 20 high-performance computing applications for industries such as automotive, airspace, oil and gas and biotechnology running on Windows Compute Cluster Edition, he said.

Windows Compute Cluster Server 2003 comprises both the Compute Cluster Edition of the OS and Microsoft Compute Cluster Pack, a combination of interfaces, utilities and management infrastructure. Windows Compute Cluster Server 2003 is expected to be available in the first half of 2006.

Observers said Microsoft has a tough road ahead in selling a new version of Windows to a market where Unix and **Linux** (Overview, Articles, Company) already have a firm hold, but Microsoft believes that the prevalence of Windows clients could give it an edge.

Most servers in the high-performance computing space are sold for a price of $250,000 or less, Adam said, citing IDC research. Many of these are departmental deployments, a place where Windows already is widely deployed but in an application customers have cobbled together, he said. While Microsoft is mainly targeting this high-volume market segment with its high-performance version of Windows, the company also will play in the higher end of the market.

###

Chronicle.com

Microsoft, Venturing Into Supercomputing, Will Underwrite 10 New Programs Worldwide

By VINCENT KIERNAN

Wednesday, November 16, 2005

In its first major foray into academic supercomputing, the Microsoft Corporation announced Tuesday that it would pay for new Institutes for High-Performance Computing at 10 universities around the world.

Bill Gates, the company's chairman and chief software architect, told a supercomputing conference in Seattle that joint research projects at the institutes would help guide software research and development at Microsoft. The institutes will help Microsoft understand researchers' needs, Mr. Gates said, according to a transcript of his talk provided by Microsoft.

The universities involved are Cornell University; Nizhni Novgorod State University, in Russia; Shanghai Jiao Tong University, in China; Tokyo Institute of Technology, in Japan; the University of Southampton, in England; the University of Stuttgart, in Germany; and the Universities of Tennessee, Texas at Austin, Utah, and Virginia.

Microsoft refused to say specifically how much money it is spending on the effort. Mr. Gates called it a "multiyear, multimillion-dollar investment."

But the dollar amounts involved appear to be modest. Jack Dongarra, director of Tennessee's Innovative Computing Laboratory, said that his lab was receiving $300,000 in cash from Microsoft this year, as well as some Microsoft software. Mike Kirby, an assistant professor of computer science at Utah, said that the university expected to receive $150,000 to $200,000 annually for at least three years.

Both researchers said they expected to use most of the money to support graduate students conducting research in high-performance computing. Mr. Dongarra said that his laboratory might provide some support to undergraduates as well.

Mr. Dongarra said the funds were a gift, with no requirement for Tennessee to produce any reports or complete any projects for Microsoft. As a gift, the money is not subject to an internal tax levied by the university's central administration on research grants, he said. "The money goes a lot further because of that," he said.

Kyril Faenov, director of high-performance computing for Microsoft, said the company selected the institutions on the basis of their research projects, their past interactions with Microsoft, and their locations.

"We wanted to make sure this is a global program," he said.

Mr. Kirby said Microsoft had reassured researchers at Utah that the company would not interfere with their work. "We have been told over and over that they highly encourage us to disseminate our results," he said. "They want us to act as scientists."

"We're investing in people," Microsoft's Mr. Faenov said. "We believe ultimately that is the most precious resource in innovation."

Mr. Gates told people attending the conference that Microsoft and other software makers must take steps to reduce the "time to insight" -- the amount of time it takes a scientist to get useful results from a supercomputer. "If you're a scientist who wants to gather a lot of information, compute around that information, look into that information, how can we make all the steps involved there far more efficient than we are today?" the transcript quoted Mr. Gates as saying.

###

HPC Wire
11/11/05

Features:

Kramer Spearheads Impressive SC05 Showing

William T.C. Kramer must be using his HPC prowess to somehow get 28 hours of each day. It's the only way to explain how he could maintain a schedule that includes chairing this year's Supercomputing extravaganza, while serving as deputy division director at the National Energy Research Scientific Computing Center and pursuing his doctorate while also remaining a dutiful and doting husband and father.

It was not surprising that SC05 was first on his mind during a recent conversation from his NERSC office at Lawrence Berkeley Laboratory. After all, this year's conference and trade show is poised to be the biggest in recent years, measuring from many different aspects.

Perhaps most foremost, SC05 has sold out every available inch of allocated exhibit area. According to Kramer, some 220 exhibitors will be occupying "every space the fire marshal would allow us to sell." The booths range in square footage from 10-by-10s to 50-by-50s.

"The industry as a whole is doing well," said Kramer, reflecting on the reasons for the encouraging vendor support. "We're seeing more people and different people." And not just on the vendor side. In particular, he added there has been lots of interest from government initiatives.

This interest has helped Kramer expand the focus of the conference, which should be particularly noticeable to habitual attendees. There is the biggest educational program ever, highlighted by the expected attendance of dozens and dozens of teachers from high schools and colleges. Much of this increase is thanks to additional funding. "There is a renewed energy regarding HPC," said Kramer. "This relates also to how HPC is being applied in a multiple of new and existing areas of study."

Attendance is running ahead of last year's pace. This is particularly interesting because it was only the late 1990s when support of SC was rumored to have peaked and whispers began questioning its relevancy. Hopefully, 2005 will let SC wipe off its Rodney Dangerfield tag altogether. After all, when the gala bounced back with 2002's showing in Baltimore, critics attributed its proximity to the nation's capital. The word went out the next year that there was little interest in Phoenix - and SC03 drew the biggest crowd ever. But the result only led to naysayers predicting a deep drop for Pittsburgh last year. The result: another attendance record.

This kind of pessimism doesn't seem to exist for Seattle's event. The No. 1 reason could be the opening keynote speaker that Kramer secured: a guy named Sir William H. Gates, Microsoft Corp's chairman and chief software architect.

Kramer admitted his masterstroke of scoring Gates as SC05's opening speaker represents two years of working various channels at Microsoft to get an audience with the man at the top. Once he did, then it still didn't end. Then there was crafting the rational for the HPC faithful as to why the software giant—not even considered a legitimate HPC player by many in the field -- should score the cherished kickoff spot that has traditionally been the roost of non-vendor dignitaries.

After all, holding the event in proximity to Microsoft's Redmond, Wash. headquarters didn't hurt SC's chances. Yet Kramer acknowledged he did not reel in Gates without a Herculean effort from himself and others counted among the SC05 inner circle.

"Gates actually canceled a trip to China to be at SC," Kramer said. "He's certainly interested in HPC. He's been getting briefed twice a week to get up to speed on all aspects. Obviously, this is a significant sign."

One thing's for sure: the HPC faithful have been giving Kramer their opinions on Gates' high-profile selection for SC05. "People love him, people hate him," Kramer chuckled. "But even if they're somewhere in-between, it looks like they will come see him speak."

According to Kramer, those attendees that do should know Gates would not be blowing

his marketing bugle. "He's going to do a technical talk," Kramer explained. "Sure, he'll bring in Microsoft's interest in HPC, but he [Microsoft] is doing an HPC product announcement in this timeframe." A lot of the keynote also will be based around Microsoft's research efforts in the field of HPC, he added.

SC05 is bolstering its bastions for Gates' speech, having more than 3,000 seats available in addition to two overflow areas. If getting Gates to speak is considered a coup, imagine the extra credit for getting him to answer some unscripted questions in an open question-and-answer forum after the speech. This looks like a reality. "We're very happy with the set-up," Kramer said. "It should make for an interesting discussion."

Kramer also is excited about the debut of what is called SC Desktop, an experiment that will extend some of the conference proceedings to people who could not typically attend the event. "This concept may really change the way the event runs," he said. For example, graduate students who cannot make the trek to Seattle due to budgetary, time or visa restrictions will be able to view some of the happenings broadcast to their computers. SC Desktop takes in Gates' keynote, SC's technical content and one exhibitor forum. Gates' speech will also be available to all interested parties through Microsoft's pressroom on its site.

SC Desktop uses Access Grid technology to connect off-site attendees anywhere in the world to specific parts of the conference programs via an appropriate network connection. Participants then will have access to speaker plenary sessions, the Cray and Fernbach Award talks and four of the half-dozen parallel technology program tracks. The tracks include two sessions of SC Papers, one Masterworks presentation and one Exhibitor Forum.

Participants can choose to view the SC Global Showcase, which includes a keynote address by Rick Stevens, father of Access Grid, along with presentations on collaborative art, Grid technology, high-resolution visualization tools and collaborative group development.

Aside from watching the goings-on from their desktops, participants can ask questions

via email. Kramer again explained SC Desktop is "an experiment for the conference," meaning he "does expect a tremendous amount," but that said, he added, "It's exciting to get more HPC content and information to a broader audience, especially students."

Because, at the end of the day -- even long ones like Kramer's -- HPC's future is hinged on getting students and technology leaders interested in the field. "We need to continue growing the awareness," he said. "And for a modest fee, SC Desktop gets the word out." Look for the effort to expand in 2006, he said, with perhaps a virtual exhibit area.

As another way to spread the HPC gospel to the masses, Kramer said SC05 is undertaking the filming of a video. Coming on the heels of the Council on Competitiveness' outstanding DVD, expect SC's version of the HPC story to be more Ken Burns than the penguins from "Madagascar," who starred in the council's movie.

The effort, which will detail why HPC machines are a necessity in our lives, will be distributed at SC. Unlike the council's effort, Kramer, added, this one is available at no cost.

Look for Microsoft to have a huge presence on the show floor, Kramer said. The company specifically targeted the exhibit space on the bridge between the conference's two exhibit halls. According to Kramer, the software giant was treated just like any other exhibitor. In fact, senior HPC stalwarts Cray Inc., IBM and Intel Corp. all passed on the exhibit territory that Microsoft cherished from the start.

Kramer points to SC05's HPC Analytics program as a path for bringing in new users to the event. The increased interest in storage issues, which started last year, also continues to go forward, he added. According to Kramer, computing, networking and storage are the three "stool legs" upon which the conference program is built.

He also credits this year's HPC Analytics Challenge for expanding the legions of supercomputing users who will be on hand. Most notably, Kramer added, HPC heavyweight Boeing will have a bigger presence at SC this year, in part due to its proximity to the event, but other big-name HPC users also have promised stepped-up

participation.

The challenge represents an opportunity for industry, academia, government and other organizations to develop and demonstrate applications highlighting analytics techniques for solving complex, real-world problems. Submissions from around the world are under consideration, including ones focused on earthquake research, predictive modeling and traffic analytics.

According to Kramer, the rise in performance and overall computational power in HPC coupled with the rising amount of data collected have resulted in a rising interest in HPC Analytics. At the same time, the field, once the exclusive domain of government and research supercomputing communities, now has generated noteworthy interest from a range of business sectors.

All signs point to surpassing 2004's attendance. One key indicator, hotel reservations, are ahead of last year's tally to the point where additional hotels had to be secured for attendees, he said.

Meanwhile, at NERSC's Berkeley Lab, Kramer's "other" job, the facility's new Linux Networx Inc. supercomputer has been in production since August. According to Kramer, users "like it very much" and it has been kept busy with science production tasks.

There is also another new machine currently under scrutiny at NERSC, Kramer said. While it represents the lab's biggest and best-performing machine ever, it has not gained full acceptance yet, he explained.

While he categorized his demanding SC05 workload as "manageable" while the responsibilities at Berkeley Lab seem numerous, he admitted there have not always been enough hours in the day for his Ph.D. studies. He is pursuing his doctorate in computer science at the University of California at Berkeley. "I'm a semester behind now," he said. "But my advisors have been supportive and understanding." Nevertheless, his studies have been relegated to late evening and early morning hours after his young daughter has gone to bed.

While he expects her to be in attendance this month in Seattle, he remarked she isn't expected to carry on the Kramer SC legacy until 2025.

###

LiveWire Weblog

Some Initial Impressions

Dan Reed, 11/15/05

The gala opening of the exhibits last night, the release of the latest Top 500 list, and Bill Gates' keynote this morning all showed the continuum of high-performance computing -- the triumph of both the large and the small. Bill's keynote illustrated the commoditization of HPC via software extensions that access deskside multiprocessors and laboratory clusters using web and grid services. This approach provides access to ISV codes and computation accelerators for day-to-day research by large numbers of users.

At the high end, IBM's custom Blue Gene/L system at LLNL raised the bar on the complexity and scale of possible simulations. It also raised bar on the resources needed to occupy the upper echelon of the Top 500 list. Just a few years ago, a 1-5 teraflop (peak) system was a flagship capability at many national research centers and occupied a top five position on the Top 500 list. Today, to aspire to the top five, one must measure machine room power budgets in megawatts and storage system sizes in petabytes.

It seems to me that one of our great challenges is to more deeply integrate the elements of this continuum via easy-to-use software and tools. One should be able to access the most powerful systems, using desktop point-and-click metaphors or even more advanced software models, just as easily as one accesses deskside computational accelerators. As others have noted, the success of any technology is measured by its invisibility -- the ability to enable scientific discovery, create economic advantage and add public value without detailed knowledge of the enabling capability. We need to make supercomputing invisible but ubiquitous.

posted Wednesday, 16 November 2005

LiveWire Weblog

SC'05 BEGINS WITH A BANG

J.S. Hurley, 11/15/05

Supercomputing 2005 began with a bang, as this year's event extends to over six continents, 10 countries and about 9,250 participants. For the first time, there are two exhibitions with 265 exhibitors, including 58 first-time exhibitors. The conference includes 500 volunteers with key program components. The keynote speaker, Bill Gates, focused on the future of computing in the sciences, with key considerations on how we connect machines and develop parallel algorithms. The problems were perceived as being very similar between high- and low-end computing. Gates stressed the importance of developing computation tools for the sciences. Three classes of computing were denoted, each with its own special challenges, including: business, consumer and technical.

Progress on 32- to 64-bit technologies was deemed as important for large datasets, and requires important features, including:

- Parallelism -- Advancements in clock speeds are not expected to continue at the same pace.
- Moving XML into the mainstream -- Connecting the business and scientific worlds.

Digital modes are also viewed as being able to play an important role.

posted Tuesday, 15 November 2005

LiveWire Weblog

A LOOK INTO THE FUTURE OF GRID RESEARCH

Wolfgang Gentzch, 11/15/05

It is now 10 years since the research community gave birth to Grid computing; 10 years since early projects like I-Way, Globus, Legion and Unicore gave us a hint to what we could do with this new technology; after the Grid "bible" from Ian (Foster) and Carl (Kesselman) started changing our minds and imagination; but the community is still struggling with its adoption, with the impact on our future and with its direction. Desperately, we are looking for unique signs from heaven, listening to the many Naradas, breaking our heads in helplessness.

And here comes one of these reliable recurrent kilometer stones (yes, I am back in Europe), giving us clear hints toward the future. I am talking about Grid 2005, this year's International Workshop on Grid Computing in Seattle, just before this year's Supercomputing. In its sixth year, this year's Grid workshop featured a program committee of over 100 computer scientists that reads like the "Who's Who" of the distributed computing community, selecting 30 papers and 18 posters out of almost 200 submissions. This is for sure one of the highest quality research conferences on Grid computing in the Universe, framed this year by visionaries (and veterans) like Dennis Gannon from Indiana University and his keynote on research challenges for the next generation of Grid systems, and Fran Berman (SDSC), Fabrizio Gagliardi (EGEE, now Microsoft), Carl Kesselman (UCLA/ISI) and Mark Linesch (GGF and HP) on the panel titled "What Will Grids Look Like in Five Years."

So, what ARE the trends we are taking away from these heavyweights? Listening to the scientists, obviously, you recognize the current challenges and open questions. One area is coming up again and again -- and it's data, data, data. I think I am still underestimating this single, largest challenge. How do we bring structure and order into this exploding data hurricane? Exponentially growing data, in many different formats, distributed over many sources and sites around the world, replicated for the sake of international collaborations in virtual organizations -- all this requires sophisticated, scalable and robust data management capabilities. They have to take into account very heavy loads, performance bottlenecks, denial of service attacks and different policies, while still guaranteeing minimum service response time.

Another recurring topic is the wide field of security -- protecting your core business, resources, data, applications, IP and other assets. Although authentication and

authorization are well understood and implemented in simple environments, they are still a challenge for global or virtual organizations in research and industry. Policy-based distributed access control, security credentials and credential renewal for long-running jobs, identity management, group memberships, policy management, attribute-based mutual trust and trust building, resource negotiation, suspicion levels associated with requesters, and general scalability challenges result in central versus decentralized approaches. And, finally, firewalls and conflicting security policies play an important role and obstacle in today's collaboration among multiple organizations.

One important trend is the introduction of Web services as a means for standardized communications in services-oriented and interoperable systems, to expose application functionalities to end-users and to provide wrappers for legacy applications integration. Other contributions to the workshop concentrated on semantic Web technologies and Grids, complex fault scenarios in large distributed systems, (autonomous) failure detection and self-repairing systems.

The keynote speaker, Dennis Gannon, has been involved in so many distributed and Grid projects that to mention all would go far beyond the scope of this blog. So, HIS voice counts! Extrapolating from his lessons learned in Grid projects like I-Way, IPG, NEESGrid, BIRN, GEONGrid and LEAD, he envisions that data will increasingly become the most demanding challenge. He described progress to this, through Web-form interfaces, and the fusion of data management and interoperable workflow tools. He anticipates building LIGHTWEIGHT Grid infrastructures and a "Google" for Grid data and application services discovery; challenges with scalability of Grid services; and the virtualization of data storage and computational resources. Are Web services really the right foundation for building Grids? Or even for building the galactic-scale Grid operating system?

The panel on "Grids in five years" was another highlight of the conference. According to Fran Berman, we will see more and more innovative Grid-based applications from the commercial sector, which integrate across scale -- like e-Bay's Grid-based shopping, real-time multi-media applications, and smart clothes, cars and highways. She also envisions Grids that will be easy to use, secure and efficient, supported by a large number of tools for performance analysis, optimization, debugging, administration and more. Mark Linesch from GGF described the transition from today's "old" world of static, silo, physical, manual, application-oriented IT, toward the "new" dynamic, shared, virtual, automated and service-oriented world. An important step toward success is the recently

announced closer collaboration of GGF with the Enterprise Grid Alliance on standards and their implementations in enterprise environments. Carl Kesselman described the future service-oriented architecture, decomposable and dynamically integrating, with applications as services composed into workflows; on-demand provisioning of resources delivered to services; and the creation virtual communities with the challenge to support dynamic policies and trust, understand what policies will work (and what not), and provide quality of services and agreement mechanisms. Finally, Fabrizio Gagliardi described the evolution from mainframes to mainstream personal and departmental clusters becoming powerful nodes in a Grid -- a vision obviously affected by his recent move to Microsoft. Watch out and listen carefully to Bill Gates' keynote at SC'05: we will hear more from Microsoft in the near future!

By the way … after six successful years and growing number of participants, the oversized Grid workshop will divorce from the annual Supercomputing conference next year and become a standalone event -- the 7th IEEE/ACM International Conference on Grid Computing. It will be co-locating with the 2006 Cluster Conference, which both will take place in September in Barcelona, right after the famous annual Fiesta weekend J

posted Tuesday, 15 November 2005

###

An Interview with John Powers, President of Digipede Technologies

Ken Farmer, 11/16/05

Wednesday November 16 2005 @ 12:20PM EST

We sat down this week with John Powers, CEO of Digipede Technologies in Oakland, California. As a Microsoft Gold Certified partner working with several groups in Redmond, Digipede offers the first commercial grid computing solution built on .NET. I asked John about how the company's offerings fit with Microsoft's in traditional HPC and beyond.

Q. When people talk about HPC on Windows, your name always seems to come up. Can you tell us a bit about the Digipede Network?

A. The Digipede Network™ is a distributed computing solution that delivers dramatically improved performance for real-world business applications. The software scales from a handful of computing nodes to thousands. Built entirely on the .Microsoft NET platform, the Digipede Network is radically easier to buy, install, learn, and use than other grid computing solutions.

If a user has an existing command-line process they need to distribute and run on their network, the Digipede Workbench provides a simple, graphical user interface to make that task easy. Instead of requiring any shell scripting to move applications and data around, the Workbench makes it simple for a non-developer to get into grid, quickly and easily.

Of course, the Digipede Network can also work with applications that are already installed on compute resources—again, using the powerful, flexible Digipede Workbench.

Lastly, the Digipede Framework Software Development Kit (SDK) makes it straightforward for developers to create grid-enabled applications themselves.

As for HPC, we see applications there for sure, but my partners and I come out of the enterprise software business. At our previous firms, we used to develop big enterprise applications, and we ran into scalability issues all the time. We developed application-specific solutions for those issues, but we wished for a more general-purpose set of technologies for building distributed applications. At Digipede, we've built the technologies we always wanted. Our experience – and our focus – is in bringing the benefits of distributed computing to mainstream enterprise applications. That's where most of our customers are today.

Q. Digipede is the first commercial grid computing software for .NET; why did you choose .NET?

A. While .NET is great technology, this was a market decision as much as a technology decision -- I like markets that are gigantic and lightly contested. There are plenty of distributed computing options on the Linux platform; most clusters, and most grid computing projects, are dominated by vendors selling solutions on top of Linux. That's fine – that just leaves more room for us. Windows is the dominant operating system on the planet, period. Yet the options for distributed computing on the Windows platform are quite limited.

Our enterprise customers tell us that they run Windows for a broad variety of applications, and they don't want to convert to Linux just to access the computing power of clusters or grids. They want to scale out their applications where they run today – which is on Windows.

And let's be clear that when we say .NET, it's true that the Digipede Network is built entirely on .NET, but we can distribute applications built with any technology that runs on Windows 2000 or later – including .NET, COM, Java, or for that matter, standalone executables developed in Fortran or C.

When we talk to customers, they tell us, very clearly:

- Most developers use Visual Studio to build applications. It's the tool preferred by most developers of enterprise applications, both internal developers and ISVs.

- Most computers run Microsoft Windows. For a large class of customers, the added complexity of maintaining a cross-platform distributed computing system is not worth the extra 10 or 20 percent of compute resources they could capture. It is far more important for these customers to support a heterogeneous combination of dedicated and shared resources running Windows (e.g. cluster nodes, desktops, and departmental servers) than multiple operating systems.

But our focus on the Microsoft platform goes beyond Microsoft's present technology and market share. Microsoft is a terrific partner, and their partner ecosystem is just unbelievable. As Microsoft Gold Certified partners, we've received marketing support, access to their product managers, insights into their product roadmap, and numerous other benefits. We just exhibited at the global launch of SQL Server 2005, Visual Studio 2005, and BizTalk Server 2006 in San Francisco, where we met with numerous potential partners – ISVs, OEMs, system builders, resellers, integrators – and we're committed to playing an important role in the Microsoft distributed computing ecosystem.

Q. So that's the enterprise story; what about scientific and technical computing, the traditional HPC market?

In traditional HPC, in academic and scientific computing, the story is different, of course. Most applications do NOT run on the Windows platform. But even here, we're hearing a lot of compelling stories about why customers want a Windows option. They are tired of

multiple Linux distros, of chasing down the right version of the right libraries to run with the right distro, and just the general level of tweaking that goes on to get an application running properly on a Linux cluster.
That's still the minority. There are scientists who have really gotten into the whole Linux culture, and are happy to climb the relatively steep learning curve, to compile their own OS and libraries to tune their application just so. But there are also a lot of scientists who tell us this is not what excites them, and that they are eager to run with a single, reliable platform. Microsoft has a lot to do to establish itself as that platform, but the door is open, and the Microsoft HPC group is doing a great job in starting that process. We see a lot of opportunities to work with that group.

Q. You mentioned that your customers are mostly outside of traditional HPC; can you give some examples of how they are applying your software?

A. Sure. We've got customers running simulations, statistical analysis, heavy text processing, and media transformation. We have fund managers doing heavy financial engineering, we have electric utility analysts running complex system operations models, we have software developers researching better algorithms for email filters. The breadth of applications is incredible, and growing week by week.

We've also got customers using the Digipede Network to deploy "Software as a Service" or SaaS offerings. For example, there's a company in San Diego that runs events – conferences, parties, and the like – and they've built a portal through which their event coordinators, customers, and vendors can share information. That site does a lot of report production, and churning out nicely formatted reports turns out to be more compute-intensive than you might think. Given the nature of their business, they get bursts of traffic around the time of each event (and they run thousands of events each year), so they need to support many simultaneous users. They used the Digipede Network to scale out this application, and increased the number of simultaneous users by a factor of ten with no change in quality of service.

The nice thing about this example is that there are so many similar possibilities. Think of a media company with a library of broadcast quality video that wants to make selections from that library available on cell phones or other hand-held video players. The configuration I just described above would allow them to do that on demand – without requiring the video files to be converted and stored in multiple formats on the company's own storage systems.

Q. You mentioned Software as a Service. Can you tell us some more about that?

Really, this is just the beginning for Software as a Service. We see Microsoft's increasing emphasis in this area as a huge opportunity for us. There are three infrastructure components you need to deliver really interesting next-generation services:

• A flexible, scalable, high-performance database. Microsoft has just delivered on that,

with SQL Server 2005.

• A simple, powerful way to build and deploy Web services. Microsoft has just delivered on that, with .NET 2.0 and the huge productivity improvements in Visual Studio 2005.

• A simple, flexible way to allocate your computing power to your service. That's what Digipede brings to the table. These are the components on which the most innovative and useful next-generation services will be built. And by integrating with Visual Studio 2005 and SQL Server 2005, and building entirely on .NET, we're positioned to grow as Microsoft's presence in this space grows.

Clearly these opportunities extend into the HPC arena as well, into scientific and technical computing. As organizations collaborate more, there are more opportunities to share expertise and resources. Creating a protein folding service, or a BLAST service, or astronomy services like those Jim Gray at Microsoft Research is working on – all kinds of new science will be possible with the right infrastructure behind new services.

Q. Microsoft's efforts in HPC are centered on the Compute Cluster Solution being demonstrated this week at Supercomputing in Seattle. How does the Digipede Network compare to Windows Server CCS?

We've been in touch with Microsoft's HPC team from the beginning, and we've shared our roadmap with them. By design, we've taken a complementary approach to theirs. The HPC team has been very clear about their focus: They're going to make Windows compute clusters a viable choice for high-performance computing. To do that, they've been working on a Windows-tuned MPI stack, working with the Visual Studio team on support for parallel debugging, working with the 64-bit OS team on support for fast interconnects, and building a job scheduler to support all that. They're getting support from compiler developers, from developers of numerical libraries. And they're working with ISVs in scientific and technical computing, getting them to support the Windows platform.

As experienced Microsoft partners, we know how to take advantage of what Microsoft provides, and to build new capabilities on top of that. So we're building on top of the tools provided by the HPC team -- and the SQL Server team, the Visual Studio team, the Windows Server 2003 team, the .NET Framework team, and so on.

The Digipede Network can live alongside CCS on each compute node, or it will interoperate with the CCS Job Scheduler to launch jobs on a CCS cluster. (I say "will" because they're not shipping yet; we've got the most recent CCS beta, and can show you what I mean, but we will officially support that interface when CCS ships next year.)

So – we can take advantage of their Job Scheduler without recreating those capabilities. We don't need to develop and tune an MPI stack; we can use theirs. For applications that

require support for fast interconnects, Microsoft got that covered for us. Instead, we build new capabilities on top of their platform, and add the ability to run applications on shared resources (not just dedicated cluster nodes), the ability to run distributed applications on 32-bit systems (not just 64-bit), the ability to harvest resources from 32- or 64-bit desktops, servers, or cluster nodes, and so on.
We also offer a very "enterprise-developer-friendly" programming model, with the ability to instantiate objects locally and execute them remotely without worrying about any of the distributed computing plumbing. While there is a lot of MPI talent in the scientific and technical computing community, that talent just isn't there in the enterprise computing community. So we've spent a lot of effort on our Digipede Framework SDK, through which programmers can grid-enable their applications without learning new programming techniques.

As a result, we're targeting different types of applications than the Microsoft HPC team. They're looking at shared memory MPI jobs, tightly coupled jobs with lots of inter-processor communication, things like finite element analysis and computational fluid dynamics. We're looking at loosely coupled jobs, parameter sweeps, Monte Carlo analysis, other types of scenario analysis, complex searches, and other problems that are more naturally parallel.

But to me the most exciting opportunities are with customers that have a need for a hybrid solution, what we've called an "extended cluster." In those scenarios, the ability to add shared desktops and servers to a dedicated CCS cluster dynamically really tips the economics for the customer away from a Linux cluster in favor of a Microsoft / Digipede solution. We'll have some specific stories about that soon.

Q. Grid computing, distributed computing, cycle harvesting – these ideas have been around for a while; why are you having success when others have not? Is it just low price?

As we have talked to customers in this market, the feedback we've received about grid computing has been remarkably consistent. Price per se is NOT a barrier to adoption. The number one barrier to adoption is complexity.

With that complexity has come a very high threshold for the adoption of grid solutions. When you need consultants, and long feasibility studies, and longer implementation projects – just to get in the game – well, that's a high threshold for most customers to cross. As a result, grid customers tend to be very large organizations, and they tend to spend a lot of time and money dealing with complex installation, configuration, integration, training and support issues.

Our goal is to eliminate complexity wherever possible, and to provide a much lower entry point. As you weed out complexity, you open this market to a much wider range of customers. That's how the benefits of grid computing will really start to show up – when

they reach small to medium businesses, and individual departments within larger organizations. That's where the energy is, that's where the innovation is, that's where the greatest needs are,
Our pricing is part of that strategy. When you can buy great grid computing software for under a thousand dollars, when you can get it installed and running without spending a dollar on consultants, then it's much easier to get started. You don't have to cross this high threshold of cost and complexity – you can get started today and see the benefits for yourself, in your own organization, and then add more power as you need it.

When you talk with folks in small and medium businesses, and in small and medium labs and departments within much larger organizations, you find out – they buy products that solve problems. They are not going to do big consulting projects to study how they might change their whole IT infrastructure to improve application performance. They're going to buy something they can use right away, without a lot of fuss.

We don't see our product pricing as dramatically lower than other solutions out there; our price per node is fairly similar to other offerings. The real difference is the lower entry point; the simplicity of implementation makes it possible for us to support a five node sale. Our competition can't say that.

Q. At the same time, five nodes don't deliver the performance improvements of 5,000. Who can gain real benefits from these smaller systems?

A. Well, I have no objection to selling a 5,000 node system. But we're seeing customers benefit from 5, 20, 50-node systems in many different fields. If your most important analytic application runs for an hour on a single machine, and we can cut that time to 12-15 minutes with a 5-node implementation of the Digipede Network Team Edition, that's a big deal. That means an analyst can get four or five times as many model runs done in a day as before, which changes the way he or she works every day.

And then it's time to add more nodes, which is just what we see happening. I'll give you an example. One of our customers is a company called InBoxer. They make an advanced spam filtering product, and they do a lot of statistical modeling – their classification of email as spam or not spam gets better the more they calibrate it by exposing it to huge amounts of email. Their statistical models were running all night on a single processor. They got the Digipede Network, installed the Digipede Agents on four or five machines, and saw an immediate impact; their jobs were cut down to less than two hours. Their CTO was so impressed that he went shopping – he bought several additional servers, deployed additional Digipede Agents on those servers, and they're down to about 40 minutes on those same jobs. That's huge for them – it completely changes the way they interact with their data and their models, and the pace at which they can improve and grow their business. And we made exactly zero trips out to help them with this – they got everything up and running themselves.

And on slightly larger projects, the pattern is similar; we're working with a big fund **An**

manager who is implementing a 50-processor system to replace some home-brew scripting they did themselves. They get better performance, ease of maintenance – but the biggest benefit they identified is ease of expansion. As they add new applications, add new users, add new assets to their portfolio, their processing requirements grow, and they can add more agents incrementally as needed.

Q. When you say "add more applications," is that usually done using the Digipede Framework SDK you mentioned earlier?

A. Absolutely. This is a really important area for us, and if I were to point to any single point of differentiation that has gotten customers most excited, this would be the one. The whole Digipede team comes from the enterprise software industry, and we're keenly aware of the needs of developers. We realize that application developers know their applications well, and want to focus on the functionality their applications deliver – and not on the nuts and bolts of distributed computing.

The Digipede Framework provides developers with the easiest path available for grid-enabling their applications. It fits right into Visual Studio (.NET and 2005) – and that means all the documentation is accessible the same way the MSDN documentation is, and that IntelliSense makes the Digipede Framework classes as easy-to-work-with as all of the .NET Framework. In short, we allow developers to work the way they're used to working, so developing a distributed application is no more difficult than developing any other kind of application.

The Digipede Framework API is remarkably simple for the powerful functionality it enables. For example, developers can designate the classes in their applications that they would like distributed on the Digipede Network. At run time, each object of that class that they instantiate is automatically migrated by the Digipede Network to a machine ready to execute the work; a method is called by the Digipede Agent on that machine. After the work is completed, the objects are re-serialized and streamed back. Nearly all of the developer's code is unchanged; the object-oriented nature of .NET programming is preserved completely. The developer no longer needs to worry about the difficult parts of distributed computing: how and where to move the appropriate EXEs and DLLs, determining which computers are on the network and available for work, and guaranteeing execution by monitoring processes. All of those chores are taken care of by the Digipede Network so the developer can concentrate on his or her own application.

Q. Do you have anything to add before we wrap up?

A. I want to wish you well with WindowsHPC.org – it's a great site, and it fills a real market need. I thank you for the opportunity to discuss our role in the Windows HPC community and beyond. I look forward to working with you and Microsoft to build that community.

####

Bill Gates Shares Microsoft's Vision for Technical Computing

Tuesday, Nov 15 @ 13:33 PST

Today at Supercomputing 2005, an international conference for high-performance computing (HPC), networking and storage, Microsoft Corp. Chairman and Chief Software Architect Bill Gates delivered the opening keynote address to more than 7,000 attendees. In the address, titled "The Role of Computing in the Sciences," Gates shared a vision of how the software industry can contribute to accelerating scientific research and engineering innovation, calling for broad collaboration between the computing industry, academia and government to make technical computing easier and more productive. Gates also announced that, as part of the company's steps toward realizing this vision, Microsoft is funding joint research projects at 10 academic centers worldwide and has released the beta 2 version of Windows Compute Cluster Server 2003, Microsoft's first product offering for high-performance computing.

"Technical computing is crucial to the many discoveries that impact our quality of life -- from making safer, more efficient cars and airplanes to addressing global health issues and environmental changes," Gates said. "Moreover, most sciences are becoming computational sciences, which is why advanced computing capabilities need to be seamlessly integrated into the end- to-end scientific process. We see many opportunities to collaborate with the scientific community on innovative solutions that will accelerate the pace of insight and discovery."

In addition to computational modeling and clusters -- groups of computers that run as unified systems -- Gates addressed the transformation resulting from the availability of massive amounts of real-world data from low-cost sensors. This powerful combination creates new opportunities, but also new challenges, particularly with how to manage, search, analyze and publish that data and the resulting conclusions. Improving end-to-end workflow will be pivotal in changing the world of work for scientists, engineers and researchers.

Gates urged software-makers and other technology vendors to work with the scientific community to spur the next revolutionary wave of discovery by creating software, tools and standards to help eliminate today's barriers to scientific collaboration.

Windows Compute Cluster Server 2003

Microsoft is working with the computing industry to help facilitate the next wave of discovery and deliver software that addresses some of the barriers scientists, engineers and researchers face. With the delivery of Microsoft Windows Compute Cluster Server 2003 Beta 2 today and general product availability scheduled for the first half of 2006, Microsoft will offer an HPC platform that accelerates the time to insight. Windows

Compute Cluster Server 2003 is designed to be simple to deploy, operate and integrate with existing infrastructure and tools.

To help ensure simple integration of Windows Compute Cluster Server 2003 with a broad set of applications and tools, Microsoft is working closely with software, application and hardware partners across the industry. This collaboration is designed to meet customers' unique needs by enabling them to choose among and run a variety of compatible HPC applications. Microsoft is being joined in its booth at Supercomputing 2005 by 17 software partners and a number of hardware partners, including Dell Inc., HP, IBM Corp., NEC Corp., Orion Multisystems Inc., Tyan Computer Corp. and Verari Systems Inc., to demonstrate 20 industry-specific HPC applications.

"The MathWorks has been providing engineers and scientists with technical computing tools for the last 20 years," said Cleve Moler, chairman and chief scientist at The MathWorks Inc. "The capabilities and direction of Windows Compute Cluster Server 2003 complement The MathWorks' mission to provide software to accelerate the pace of innovation and discovery in engineering and science. The combination of the Windows platform and MATLAB distributed computing tools allows customers to make better and faster progress with high-performance computing tasks and deliver meaningful results that can accelerate the time to insight."

Microsoft is working with software partners across the manufacturing, life sciences and geosciences, and other industries to make a robust set of commercial applications available on Microsoft Windows Compute Cluster Server 2003. This includes integrating applications with the Microsoft Message Passing Interface and the Microsoft job scheduler, and offering performance-tuning technologies to run on Windows Compute Cluster Server 2003. As part of this effort, Microsoft is working with Intel Corporation to provide software vendors with a robust set of resources -- including 64-bit hardware, software and tools from both companies, and joint on-site engineering assistance -- to enable broader application support on Windows Compute Cluster Server 2003 when it is released to the market.

Partnering With and Investing in the HPC Community

Gates also announced an investment in 10 Institutes for High-Performance Computing worldwide. This multiyear, multimillion-dollar investment in joint research projects at these institutes will help guide ongoing software research and product innovation at Microsoft to address the most challenging technical computing problems. These institutes are Cornell University (U.S.); Nizhni Novgorod State University (Russia); Shanghai Jiao Tong University (China); Tokyo Institute of Technology (Japan); University of Southampton (England); University of Stuttgart (Germany); University of Tennessee (U.S.); University of Texas at Austin (U.S.); University of Utah (U.S.); and University of Virginia (U.S.).

"Easier access to workgroup and departmental HPC resources, such as that provided by Windows Compute Cluster Server 2003, will play an ever-increasingly important role in furthering scientific and engineering innovation in academia and industry," said Jack Dongarra, Ph.D., distinguished professor and director of the University of Tennessee Innovative Computing Lab. "Our research project will focus on self-adapting high-performance applications and software on the Windows platform, which could help scientists and engineers reduce the time and effort to discovery. I've found that Microsoft's passion for technology innovation around an integrated development and deployment software platform will help expand the potential for many who haven't had access to high-performance computing resources in the past."

In the closing statements of his keynote address, Gates reiterated Microsoft's long-term commitment to working with the HPC community and emphasized his enthusiasm for the future role of computing in scientific discovery.

Founded in 1975, Microsoft is the worldwide leader in software, services and solutions that help people and businesses realize their full potential.

###

SuperComputingOnline.com

SilverStorm Technologies Announces Support for Windows Compute Cluster Server 20

Wednesday, Nov 16 @ 09:30 PST

SilverStorm Technologies announced it has been working with Microsoft and select early adopter customers during beta testing of Windows Compute Cluster Server 2003. With support for InfiniBand interconnect drivers, Windows Compute Cluster Server 2003 can offer mainstream commercial customers the same benefits of greater performance, lower latency and one wire interconnect that large cluster InfiniBand networks have enjoyed. SilverStorm Technologies also is powering leading high performance computing applications, such as Landmark Graphics Nexus, which will run on Windows Compute Cluster Server 2003.

SilverStorm's InfiniBand equipment will be used at the Supercomputing 2005 trade show November 14-17, 2005, in Seattle during Microsoft's live keynote demonstrations of Windows Compute Cluster Server 2003 beta 2. As an active member of the Open InfiniBand (OpenIB) Alliance, SilverStorm spearheaded Windows OpenIB development, contributing significantly to the code. SilverStorm will participate in the first multi-vendor interoperability test of InfiniBand at the Supercomputing 2005 trade show.

"SilverStorm is excited to be working with Microsoft to deliver high performance computing solutions for mainstream markets," said Reini Florin, vice president of marketing, SilverStorm Technologies. "By working with Microsoft, SilverStorm can provide mainstream markets the same benefits large cluster networks enjoy today -- high quality, reliable, scalable solutions combined with unparalleled service."

The SilverStorm family of products includes the industry's most scalable InfiniBand switching solution providing users the highest port density available in a single chassis. SilverStorm provides a low latency, multi-protocol interconnect fabric that supports cluster inter-process communication as well as virtual storage and networking I/O access.

Windows Compute Cluster Server 2003 is Microsoft's official entry into the world of high performance computing. Windows Compute Cluster Server 2003 is designed to accelerate time-to-insight by providing a high performance computing platform that is simple to deploy, operate and integrate with existing infrastructure and tools.

"Microsoft is pleased to be working with SilverStorm Technologies to deliver a solution that will enable easier access to workgroup and departmental high performance computing resources for many who haven't had access in the past," said Kyril Faenov, director of high performance computing, Microsoft Corp. "SilverStorm's products act as the InfiniBand backbone in a Windows network to provide solid cluster performance improvements."

Breakthrough Simulation Wins Prestigious 2005 Gordon Bell Prize

Wednesday, Nov 30 @ 10:22 PST

A team of scientists led by physicist Fred Streitz was awarded the 2005 Gordon Bell Prize for pioneering materials science simulations conducted on the world's fastest supercomputer at Lawrence Livermore National Laboratory. The winner was announced at the conclusion of the Supercomputing 2005 conference held earlier this month in Seattle, Wash. Other team members included James Glosli, Mehul Patel, Bor Chan, Robert Yates and Bronis de Supinski of Lawrence Livermore, and James Sexton and John Gunnels of IBM. The title of their entry was "100+ Tflop/s Solidification Simulations on BlueGene/L."

Named for one of the founding fathers of supercomputing, the prestigious Gordon Bell Prize is awarded to innovators who advance high-performance computing.

Running a newly developed, three-dimensional molecular dynamics code (ddcMD) on BlueGene/L – the IBM machine that ranks No. 1 on the list of the world's Top500 supercomputers -- the team investigated solidification in tantalum and uranium at extreme temperatures and pressure with simulations ranging in size from 64,000 atoms to 524 million atoms.

"These simulations allow us for the first time to examine the process of solid formation at high temperature and pressure from the atomistic level. We can actually watch, atom by atom, as macroscopic grains grow out of the liquid and form structures," said Streitz. "This allows us to better understand the properties of these metals and has important implications for the development of stronger metals, such as those that might be used for aircraft or automobile components as well as other applications."

In what was the largest simulation of its kind ever undertaken, the team achieved a performance rate of up to 107 teraflop/s (trillion operations per second) with a sustained rate of 101.7 teraflop/s over a seven-hour run on the IBM machine's 131,072 processors. This performance, using what will be a workhorse application for BlueGene/L to do work vital to the National Nuclear Security Administration's (NNSA) mission, exceeds the best performance by any other computer on the industry standard LINPACK benchmark.

"BlueGene/L is ideally suited for precisely the kinds of molecular dynamics studies that the program needs to do at this time to better understand the behavior of materials at high temperature and pressure," said Bruce Goodwin, Associate Director for Defense and Nuclear Technologies at LLNL. "By working with IBM to develop this machine, we created simultaneously a solution that provides both the best sustained performance and the lowest cost per floating point operation. This is an ideal combination. Usually, one gets one or the other. In this case, we got both, and while we planned it this way, we are

delighted that it worked out so well."

Simulations of the solidification of metals such as tantalum and uranium under high temperatures and pressures provide valuable insights into the properties of these materials important to the National Nuclear Security Administration's (NNSA) program to ensure the safety, security and reliability of the nation's nuclear deterrent without underground testing – stockpile stewardship.

The simulations were conducted by NNSA's Advanced Simulation and Computing (ASC) program, a tri-lab effort uniting the computing know-how of Sandia, Los Alamos and Lawrence Livermore national laboratories.

NNSA scientists need the detailed information about the properties of the materials in nuclear weapons that 3D simulations provide to understand the effects of aging on those weapons systems as they age well beyond their intended design life.

Founded in 1952, Lawrence Livermore National Laboratory is a national security laboratory, with a mission to ensure national security and apply science and technology to the important issues of our time. Lawrence Livermore National Laboratory is managed by the University of California for the U.S. Department of Energy's National Nuclear Security Administration.

###

Force10 TeraScale E-Series Enables Record During Bandwidth Challenge

Wednesday, Nov 30 @ 04:34 PST

Force10 Networks announced that its TeraScale E-Series provided the high 10 Gigabit Ethernet density that enabled the Bandwidth Challenge entries at the SC05 conference to demonstrate the latest applications for harnessing large amounts of bandwidth to advance science and solve problems. The winning entry from CalTech, Fermi National Accelerator Laboratory and the Stanford Linear Accelerator Center, which transferred particle physics data between the host sites and collaborating institutions worldwide, delivered 131 Gigabits per second of sustained IP traffic, setting a new Bandwidth Challenge record.

"The Bandwidth Challenge showcases leading edge techniques and applications for utilizing massive amounts of bandwidth to advance science and computing, and the innovative approaches used during the event point the way to how the world will take advantage of vast network resources in the future," said Debbie Montano, director of research and education alliances at Force10 Networks and the chair of the Bandwidth Challenge. "At the core of this ability to transfer large datasets between geographically diverse sites is the TeraScale E-Series with both the density and resiliency to ensure disparate locations can collaborate in real time."

The real-time particle event analysis of the winning entry required transfers of Terabyte-scale datasets and leveraged computing resources from multiple sites within the U.S. and worldwide. As the backbone of SCinet, the conference's high performance network, the Force10 Terascale E-Series switched the majority of the traffic from the winning entry. Additionally, Force10 provided many of the 22 line-rate 10 Gigabit Ethernet links into the CalTech and Fermi/SLAC booths, further enabling them to sustain throughput of 131 Gigabits per second.

The winning entry also leveraged several other external networks in which the Force10 TeraScale E-Series provides the high performance backbone, including Internet2's Hybrid Optical and Packet Infrastructure (HOPI), UltraScience Net, Teragrid and StarLight, which interconnects most of the world's major high performance research and education networks. Additionally, the TeraScale E-Series was deployed in the Fermi National Accelerator Laboratory to support its winning entry.

The Force10 TeraScale E-Series provided more than 80 line-rate 10 Gigabit Ethernet ports to build SCinet. Additionally, Force10 provided more than 40 line-rate 10 Gigabit Ethernet links to exhibitors for onsite demonstrations. As the high performance switches at the nexus of the SC05 network, the Force10 TeraScale E-Series processed nearly all traffic for the Bandwidth Challenge.

The Force10 TeraScale E-Series support an industry-leading 1,260 Gigabit and 224 Ten Gigabit Ethernet ports per system, providing long-term scalability for seamless network expansion. In addition to leading density, the unique multi-processor architecture of the TeraScale E-Series distributes switching, routing and management functionality between three processors to provide unmatched system level resiliency. With the combination of density and resiliency, the TeraScale E-Series simplifies network topologies and reduces both capital and operating costs.

###

World Network Speed Record Shattered at SC05

Wednesday, Nov 30 @ 04:28 PST

For the third consecutive year, the “High Energy Physics” team of physicists, computer scientists and network engineers led by the California Institute of Technology, the Stanford Linear Accelerator Center (SLAC), Fermilab, the University of Michigan and CERN and partners at the University of Florida, Vanderbilt and the Brookhaven National Lab, as well as international participants from the UK (University of Manchester, and UKLight), Brazil (Rio de Janeiro State University, UERJ, and the State Universities of São Paulo, USP and UNESP)., Korea (Kyungpook National University, KISTI) and Japan (the KEK Laboratory in Tsukuba) joined forces to set a new world record for data transfer, capturing first prize at the Supercomputing 2005 (SC|05) Bandwidth Challenge (BWC).

The HEP team’s demonstration of “Distributed TeraByte Particle Physics Data Sample Analysis" achieved a peak throughput of 151 gigabits per second (Gbps), and an official mark of 131.6 Gbps measured by the BWC judges on 17 of the 22 optical fiber links used by the team, beating their previous mark for peak throughput of 101 Gbps by 50%. The record data transfer speed is equivalent to downloading five full DVD movies per second, or serving 10,000 MPEG2 HDTV movies simultaneously in real time, or transmitting all of the printed content of the Library of Congress in 10 minutes. The team sustained average data rates above the 100 Gbps level for several hours for the first time, and transferred a total of 475 Terabytes of physics data among the team’s sites throughout the US and overseas within 24 hours. The extraordinary data transport rates were made possible in part through the use of the FAST TCP protocol developed by Professor Steven Low and his Caltech Netlab team, as well as new data transport applications developed at SLAC and Fermilab and an optimized Linux kernel developed at Michigan.

Professor Harvey Newman of Caltech, head of the HEP team and US CMS Collaboration Board Chair, who originated the LHC Data Grid Hierarchy concept, said: “This demonstration allowed us to preview the globally distributed Grid system of more than 100 laboratory and university-based computing facilities that is now being developed in the US, Latin America and Europe in preparation for the next generation of high energy physics experiments at CERN’s Large Hadron Collider (LHC) that will begin operation in 2007. We used a realistic mixture of streams, including the organized transfer of multi-Terabyte datasets among the laboratory centers at CERN, Fermilab, SLAC and KEK, plus numerous other flows of physics data to and from university-based centers represented by Caltech, Michigan, Florida, Rio de Janeiro and São Paulo in Brazil, and Korea, to effectively use the remainder of the network capacity. The analysis of this data will allow physicists at CERN to search for the Higgs particles thought to be responsible for mass in the universe, supersymmetry, and other fundamentally new phenomena bearing on the nature of matter and spacetime, in an energy range made accessible by the

LHC for the first time."

The largest physics collaborations at the LHC, CMS and ATLAS, each encompass more than 2000 physicists and engineers from 160 universities and laboratories. In order to fully exploit the potential for scientific discoveries, the many Petabytes of data produced by the experiments will be processed, distributed and analyzed using a global Grid. The key to discovery is the analysis phase, where individual physicists and small groups repeatedly access, and sometimes extract and transport Terabyte-scale data samples on demand, in order to optimally select the rare "signals" of new physics from potentially overwhelming "backgrounds" from already-understood particle interactions. This data will amount to many tens of Petabytes in the early years of LHC operation, rising to the Exabyte range within the coming decade.

Matt Crawford, head of the Fermilab network team at SC|05 said "The realism of this year's demonstration represents a major step in our ability to show that the unprecedented systems required to support the next round of high energy physics discoveries are indeed practical. Our data sources in the bandwidth challenge were some of our mainstream production storage systems and file servers, which are now helping to drive the searches for new physics at the high energy frontier at Fermilab's Tevatron, as well the explorations of the far reaches of the universe by the Sloan Digital Sky Survey."

Les Cottrell, leader of the SLAC team and Assistant Director of Scientific Computing and Computing Services said: "Some of the pleasant surprises at this year's challenge were the advances in throughput we achieved using real applications to transport physics data, including bbcp and xrootd developed at SLAC. The performance of bbcp used together with Caltech's FAST protocol and an optimized Linux kernel developed at Michigan, as well as our xrootd system, were particularly striking We were able to match the performance of the artificial data transfer tools we used to reach the peak rates in past years."

Future optical networks, incorporating multiple 10 Gbps links are the foundation of the Grid system that will drive the scientific discoveries. A "hybrid" network integrating both traditionally switching and routing of packets, and dynamically constructed optical paths to support the largest data flows, is a central part of the near-term future vision that the scientific community has adopted to meet the challenges of data intensive science in many fields. By demonstrating that many 10 Gbps wavelengths can be used efficiently over continental and transoceanic distances (often in both directions simultaneously), the high energy physics team showed that this vision of a worldwide dynamic Grid supporting many Terabyte and larger data transactions is practical.

Shawn McKee, Associate Research Scientist in the University of Michigan Department of Physics and leader of the UltraLight Network technical group, said: "This achievement is an impressive example of what a focused network effort can accomplish. It is an

important step towards the goal of delivering a highly capable end-to-end network-aware system and architecture that meet the needs of next-generation e-Science".

The team hopes this new demonstration will encourage scientists and engineers in many sectors of society to develop and plan to deploy a new generation of revolutionary Internet applications. Multi-gigabit/s end-to-end network performance will empower scientists to form "virtual organizations" on a planetary scale, sharing their collective computing and data resources in a flexible way. In particular, this is vital for projects on the frontiers of science and engineering, in "data intensive" fields such as particle physics, astronomy, bioinformatics, global climate modeling, geosciences, fusion, and neutron science.

The new bandwidth record was achieved through extensive use of the SCInet network infrastructure at SC|05. The team used fifteen 10 Gbps links to Cisco 6500 series switch routers provided by Cisco Systems at the Caltech Center for Advanced Computing Research (CACR) booth, and seven 10 Gbps links to a 6500 switch at the SLAC/Fermilab booth, together with computing clusters provided by Hewlett Packard, Sun Microsystems and IBM and 10 gigabit Ethernet server interfaces provided by Neterion and Chelsio. The external network connections to Los Angeles, Sunnyvale, the Starlight facility in Chicago and Florida included the Cisco Research, Internet2/HOPI, UltraScience Net and ESnet wavelengths carried by National Lambda Rail (NLR), Internet2's Abilene backbone, the three wavelengths of TeraGrid, an ESnet link provided by Qwest, the Pacific Wave link and Canada's CANARIE network. International connections included the US LHCNet links (provisioned by Global Crossing and Colt) between Chicago, New York and CERN, the CHEPREO/WHREN link (provisioned by LANautilus) between Miami and Sao Paulo, the UKLight link, the Gloriad link to Korea, and the JGN2 link to Japan. Regional connections included six 10 Gbps wavelengths provided with the help of CIENA to Fermilab, two 10 Gbps wavelengths to the Caltech campus provided by Cisco Systems and California's CENIC network, two 10 Gbps wavelengths to SLAC provided by ESnet and UltraScienceNet, three wavelengths between Starlight and the University of Michigan over Michigan Lambda Rail (MiLR), and wavelengths to Jacksonville and Miami across Florida Lambda Rail (FLR). During the test, several of the network links were shown to operate at full capacity for sustained periods.

While the SC|05 demonstration required a major effort by the teams involved and their sponsors, in partnership with major research and education network organizations in the U.S., Europe, Latin America and Asia Pacific, it is expected that networking on this scale in support of the largest science projects (such as the LHC), will be commonplace within the next three to five years. The demonstration also appeared to stress the network and server systems used, and so the team is continuing its test program to put the technologies and methods used at SC|05 into production use, with the necessary level of reliability in time for the start of the LHC research program.

As part of the SC|05 demonstrations, a distributed analysis of simulated LHC physics data was done using the Grid-enabled Analysis Environment (GAE) developed at Caltech for the LHC and many other major particle physics experiments, as part of the Particle Physics Data Grid (PPDG), GriPhyN/iVDGL, Open Science Grid and DISUN projects. This involved the transfer of data to CERN, Florida, Fermilab, Caltech, and Brazil for processing by clusters of computers, and finally aggregating the results back to the show floor to create a dynamic visual display of quantities of interest to the physicists. In another part of the demonstration, file servers at the SLAC/FNAL booth, and in Manchester also were used for disk to disk transfers between Seattle and the UK.

The team used Caltech's MonALISA (MONitoring Agents using a Large Integrated Services Architecture) system to monitor and display the real-time data for all the network links used in the demonstration. It simultaneously monitored more than 14,000 grid nodes in 200 computing clusters, as illustrated in the figure. MonALISA is a highly scalable set of autonomous self-describing agent-based subsystems which are able to collaborate and cooperate in performing a wide range of monitoring tasks for networks and Grid systems, as well as the scientific applications themselves.

The network has been deployed through exceptional support by Cisco Systems, Hewlett Packard, Neterion, Chelsio, Sun Microsystems, IBM and Boston Ltd., as well as the network engineering staffs of National LambdaRail, Internet2's Abilene Network, ESnet, TeraGrid, CENIC, MiLR, FLR, Pacific Wave, AMPATH, RNP and ANSP/FAPESP in Brazil, KISTI in Korea, UKLight in the UK, JGN2 in Japan, and the Starlight international peering point in Chicago. The demonstration and the developments leading up to it were made possible through the strong support of the U.S. Department of Energy and the National Science Foundation, in cooperation with the funding agencies of the international partners.

###

Winners of SC|05 HPC Analytics Challenge Leverage SGI

Wednesday, Nov 23 @ 17:05 PST

To understand how biomolecules like DNA move across protein pores - a problem that for decades proved too complicated to simulate - a team of renowned British scientists is turning to a wide array of technologies, including visualization and computational solutions from Silicon Graphics. Their resulting project, designed to create a Simulated Pore Interactive Computing Environment (SPICE), was honored last week at SC|05, an international supercomputing conference where the researchers tied for first place in the HPC Analytics Challenge.

Led by scientists from University College London (UCL), the global SPICE project team found a way to compute the free energy profile of a translocating biomolecule along the vertical axis of a pore. Simulating translocation of DNA, RNA and polypeptides may someday help medical researchers control that activity, potentially leading to new treatments for numerous medical conditions and illnesses.

The SPICE team's approach is so advanced that traditional means had rendered the problem unsolvable. The ongoing project requires fully atomistic simulations, which illustrate molecular and chemical interactions within the smallest unit of matter.

"The time scale for DNA translocation molecules is on the order of tens of microseconds," said Professor Peter Coveney, one of the SPICE project's lead investigators and director of the Centre for Computational Science at UCL. "Simulating timescales like these for models of 275,000 atoms or larger just isn't possible with standard approaches to molecular dynamics. Even assuming that computer power will continue to accelerate along the curve described by Moore's Law, we're still a couple of decades away from the day when this type of simulation will be routine."

To achieve the project's aims, the team is using two HPC grids - the TeraGrid in the USA and the UK's National Grid Service - along with SGI visualization resources located at UCL. Key to visualizing the simulation is a Silicon Graphics Prism system with six Intel Itanium 2 processors and four ATI FireGL graphics pipes connected via OpenGL Vizserver, which allows remote team members to collaborate by interactively viewing the simulation. SPICE team members also leveraged two Silicon Graphics Prism systems located in the University of Manchester booth at SC|05. Other SGI resources, including a multi-node SGI Altix installation, are linked via the UKLight optical network so SPICE team members have high-speed access to remote grid and visualization resources.

"Without the use of a sophisticated grid infrastructure and high-performance visualization resources, we would not be able to pursue this research," said Coveney. "Throughout all our work, SGI's hardware and software has played a key role in making our efforts so

successful."

The SPICE team is led by UCL's Coveney, who collaborates with Shantenu Jha, and Matt Harvey. The trio in turn is collaborating with Stephen Pickles of UCL, Robin Pinning of the University of Manchester, Peter Clarke of the University of Edinburgh, Bruce Boghosian of Tufts University, Charlie Catlett of TeraGrid, Charles Laughton of Nottingham University, Rob Pennington of National Center for Supercomputing Applications (NCSA) TeraGrid, Sergiu Sanielevici of the Pittsburgh Supercomputing Center, Jennifer Schopf of Argonne National Lab, and Richard Blake CCLRC Daresbury.

"The work of these researchers reveals a unique and inspiring kind of ingenuity - one that uses collaboration to push beyond the limitations of current technology to do something that has been widely viewed as impossible," said Tim Butchart, UK country manager and EMEA marketing director, SGI. "We're delighted to see SGI visualization and computation solutions play a pivotal role in the SPICE team's success, and we congratulate them on winning the HPC Analytics Challenge at SC|05."

As it does annually, SC|05 last week presented awards that recognize the innovative work of conference participants and HPC leaders. The HPC Analytics Challenge is a new award competition honoring select technical and commercial applications that use leading-edge, advanced analytics techniques to solve complex, real-world problems. The growing need to leverage large amounts of underutilized data has led to the use of sophisticated methods for analysis and high-end visualization in conjunction with high performance computing, bandwidth and networking capabilities.

The two teams that won awards were chosen from among 21 participants in the competition. Tying with the SPICE team was a team of researchers whose project is titled, "Real Time Change Detection and Alerts from Highway Traffic Data."

Other SC|05 Awards
Several other teams earned accolades at SC|05 for work completed on Columbia, NASA's 10,240-processor SGI Altix supercomputer. For instance, one of the other five finalists in the HPC Analytics Challenge leveraged Columbia's formidable computational power to create a new, more effective visualization tool for interrogating large, unsteady data sets. In another category that featured two teams tied for first place, the Best Technical Paper Award went to "High Resolution Aerospace Applications using the NASA Columbia Supercomputer," authored by Dimitri J. Mavriplis of the University of Wyoming, Michael J. Aftosmis of NASA Ames Research Center, Marsha Berger of the Courant Institute.

In addition, the SGI Altix family took top honors in the Tour de HPCycles panel session. In the SC|05 competition fashioned after the cycling world's Tour de France, a panel of HPC end users voted to award jerseys for a range of categories. SGI Altix earned the

coveted Yellow Jersey, the award for best overall supercomputer.

###

Ohio professor receives national computational science award

Tuesday, Nov 22 @ 06:18 PST

Capital University Professor Ignatios Vakalis won the Undergraduate Computational Engineering and Sciences (UCES) award in Seattle last week during the international Supercomputing 2005 conference. Vakalis holds many distinguished roles as a professor of mathematics and computer science, executive director of the Center of Computational Studies at Capital University, and educational coordinator of a statewide initiative in computational science. He was awarded for his innovation, educational impact, and breadth in developing and implementing educational materials for Computational Engineering and Sciences (CES).

"It is a great honor to win the UCES award given by the Krell Institute and sponsored by the Department of Energy," Vakalis said. "For the last seven years I've had the chance to work with many talented and passionate educators and researchers, and I was given the chance to form a consortia of institutions for the purpose of developing computational science educational materials and curricula for the undergraduate arena."

The UCES project bestows the award to promote excellence in computational science and education. Awards are given to faculty, students, or members of other organizations who make significant and innovative contributions in computational science education. This is the twelfth year for the prestigious award.

Vakalis is responsible for developing computational engineering and science course material, including the posting of web tools with cross-disciplinary impact, and for outreach through the establishment of the Keck Undergraduate Computational Science Education Consortium.

Vakalis is also a co-principle investigator on a $250,000 National Science Foundation (NSF) research project to help educators and researchers develop computational science instructional guidelines and shared educational curriculum for undergraduate students attending Ohio colleges. Educators and researchers at the Ohio Supercomputer Center (OSC), Capital University, and the Ohio Learning Network (OLN) are leading the two-year study to improve and standardize undergraduate computational sciences course curriculum.

Research on improving cyber-infrastructure is coordinated and assisted by OSC and involves faculty and students from a consortium of nine Ohio schools that are developing computational science modules for delivering course content tailored to community, four-year, and minority-serving colleges. The materials meet educational objectives for computer science courses, and Internet video technologies facilitate shared instruction over Ohio's Third Frontier Network (TFN).

Computational science is a key factor in understanding science. It is indispensable to solving complex problems in every sector from traditional science and engineering to such vital areas as national security, public health and economic innovation. Computational models capture and analyze vast amounts of experimental and observational data and address problems previously deemed intractable or beyond imagination. Computational science affords the opportunity to view and investigate phenomena like tiny atoms, large galaxies, the rapid process of photosynthesis, complex aircraft engines and much more.

Undergraduate computational science programs are critical to education because they are the technical force of the future. These programs invigorate students to pursue graduate work and prepare future mathematics and science teachers for secondary education. Therefore, the greatest impact of reform is taking place at the undergraduate levels.

"My goal for developing the computational science coursework is that students will see the beauty and the practical use of computational science," said Vakalis. "I want students to be aware of the intersection and interplay among mathematics, computing and science."

Vakalis serves as project director for the Keck consortium and is Capital University's principal investigator for the grant awarded from the W. M. Keck Foundation. The consortium includes 12 institutions across the country: Capital University, College of the Holly Cross, Harvey Mudd College, The Ohio State University, Pomona College, San Diego State University, San Diego Supercomputer Center, Shodor Education Foundation, Skidmore College, University of Wisconsin - Eau-Claire, Wittenberg University and Wofford College.

The consortium serves as a model for institutions to collaborate and develop computational science materials and curricula. It also serves as a catalyst in forming a new consortium of institutions. Consortium goals are to develop class-tested educational materials for a variety of computational science courses, infuse computational science curricula to member institutions, and prepare the next generation of scientists so they are equipped with the necessary computational tools.

###

SuperComputingOnline.com

PNNL Wins SC05 StorCloud Challenge

Saturday, Nov 19 @ 15:21 PST

For the second year in a row, the Pacific Northwest National Laboratory demonstrated its computational storage capabilities at the annual supercomputing conference. The PNNL team streamed video data from high-capacity storage on the conference show floor to Hewlett-Packard-loaned hardware in the lab's booth. The amount of data moved ranged from 30 to 60 gigabits per second. At the low end, this is the equivalent of transferring and processing a full DVD of video every second.

This direct access to data on parallel file systems across high-speed wide area links was made possible by recent network bandwidth increases at PNNL, including new fiber optic connections to Richland and a connection to the Department of Energy's UltraScience Net exchange in Seattle. UltraScience Net provides on-demand dedicated bandwidth channels at multi, single and sub lambda resolutions. PNNL used six dedicated 10 gigabit lambdas to transfer data from an HP StorageWorks Scalable File Share parallel file system to an Itanium cluster located more than 250 miles away. Following processing, the resulting data set was written back to the storage cluster

SuperComputingOnline.com

UK e-Science project wins top supercomputing award

Thursday, Nov 17 @ 21:10 PST

A UK e-Science project has won a top award at SC05, the world's premier supercomputing conference in Seattle this week. SPICE (Simulated Pore Interactive Computing Environment) achieved success in the HPC Analytics Challenge for demonstrating the use of innovative techniques in rigorous data analysis and high-end visualisation to solve a complex, real-world problem

"SPICE shows how the power of supercomputers on both sides of the Atlantic can be harnessed to simulate and visualise biological processes of unprecedented complexity. We're delighted with this award," says Professor Peter Coveney, principal investigator for the SPICE project from University College London.

The SPICE team convinced the judges with their simulation of DNA strands passing through a cell membrane. Knowledge of this important biological process is crucial for understanding the transfer of genetic information during cell division, and for applications such as the design of high-throughput DNA screening devices. However, it takes place over a much longer timescale than is possible to simulate using conventional computational methods.

"Many biological processes take longer than a nanosecond – that's what makes them so computationally difficult," says Dr Shantenu Jha, technical lead on the SPICE team.

SPICE uses technology developed under another UK e-Science project, RealityGrid, to marshal the resources of supercomputers on the UK National Grid Service (NGS) and the US TeraGrid, connected by dedicated high bandwidth optical channels. Even with resources of this grid-of-grids to hand, the simulation is too large for straight computation. SPICE has overcome this obstacle by dividing the simulation into two stages.

In the first, the researcher gets a rough "feel" for the DNA's progress from the response of a haptic device (joystick) used to pull it through a protein nanopore embedded in the cell membrane. "You try to pull the DNA through the pore and you can feel the strain on it. It's a very smart way of probing the DNA's local energetic environment – and it's fun," says Dr Jha.

In the second stage, insight gained from the first is used to set the parameters for a set of full-scale simulations. "By doing some smart exploration first, we're limiting the computation we need for a detailed, rigorous analysis," says Dr Jha.

Such complex simulations would not be possible without the use of dedicated optical

networks to connect supercomputers in the US and UK. The researchers steer the simulation in real time via the haptic device, each snapshot of the simulation requiring several hundred processors and simultaneous high-end compute and visualization resources. Standard packet-switched networks, even with high bandwidth, cannot guarantee sufficient quality of service for such interactivity.

"Without dedicated optical networks in the US, UK and across the Atlantic, SPICE would be impossible. There's no loss or re-ordering of data which means that we can steer the simulations interactively," says Professor Coveney. SPICE is one of the first demonstrations of the UK's new dedicated optical research network, UKLight.

###

SuperComputingOnline.com

The MathWorks Demonstrates Strong Presence at Supercomputing 2005

Thursday, Nov 17 @ 19:46 PST

The MathWorks today announced it demonstrated a strong leadership position at this year's premier conference on high-performance computing, networking, and storage, Supercomputing 2005 (SC|05). The event, which concludes this week in Seattle, Washington, served as the stage for The MathWorks to introduce the latest version of its high-productivity development tool, the Distributed Computing Toolbox 2. This latest version of the software enables engineers and scientists to dramatically simplify the task of developing distributed computing applications. Additionally, the Distributed Computing Toolbox 2 was featured by Microsoft chairman Bill Gates in his keynote address on November 15.

"Advancements in high-performance technical computing tools are crucial to engineers and scientists tasked with solving today's toughest computing challenges," said Kyril Faenov, director of high performance computing at Microsoft Corp. "The Distributed Computing Toolbox 2 from The MathWorks, coupled with Microsoft Windows Compute Cluster Server 2003, brings the power of high-performance computing to the departmental and workgroup levels. As a result, the administration and management of clusters can be simplified for those groups who typically had to rely on dedicated IT groups in the past."

The strong presence of The MathWorks at SC|05 demonstrates the Company's commitment to providing engineers and scientists with distributed computing tools for technical computing. In addition to announcing the release of the Distributed Computing Toolbox 2, Company representatives, including MathWorks cofounder and creator of MATLAB(R), Cleve Moler, shared their expertise throughout the event's extensive technical program by participating in tutorials, panels, and presentations at partner booths.

The MathWorks participated in the following SC|05 technical programs and demonstrations:
-- Loren Dean and Silvina Grad-Freilich of The MathWorks co-presented the tutorial, "Parallel and Distributed Computing with MATLAB," with Jeremy Kepner and Albert Reuther of MIT Lincoln Laboratory. The tutorial used customer stories to illustrate some of the most popular technologies for writing parallel and distributed MATLAB programs.
-- Ohio Supercomputing Center's booth featured a Parallel MATLAB mini-symposium.
-- Microsoft and Platform Computing featured demonstrations of the Distributed Computing Toolbox 2 in their booths.
-- Orion and ClearSpeed partner booths both featured MathWorks tools, while partner Hewlett Packard demonstrated how MathWorks tools are integrated with the Enfuzion(R)

scheduler from Axceleon.
-- The Birds-of-a-Feather session, "2005 HP Challenge Award," included a presentation by Cleve Moler.
-- Cleve Moler was also a member of the panel "Return of HPC Survivor--Outwit, Outlast, Outcompute." The panel was structured as a series of "rounds," each posing a specific question about system design, philosophy, implementation, or use.

###

Internet2 Demonstrates Optical Networking Firsts

Wednesday, Nov 16 @ 15:30 PST

For the first time, three radio telescopes distributed around the world will be connected via dynamically provisioned dedicated optical circuits for an electronic Very-Long-Baseline Interferometry (e-VLBI) observation. Internet2 announced this scientific and networking achievement at the first major demonstration of its nationwide Hybrid Optical and Packet Infrastructure (HOPI) testbed, during the SC|05 conference held in Seattle, Washington this week. The demonstration marks a critical milestone in dynamic or "on demand" optical networking that can support even the most extreme applications used by the global research and education community today.

Leveraging the HOPI infrastructure together with the NSF-funded DRAGON testbed, the telescopes located in Westford Mass. US; Greenbelt, Md. US; and Onsala, Sweden will be dynamically linked via dedicated low-latency optical circuits to a central data correlator and simultaneously transmit multiple gigabits-per-second of data during a 20-minute observation. The team will also attempt to connect to a fourth telescope in Kashima, Japan during the demonstration. Historically, radio astronomy data was recorded on magnetic tape or disk at each site and shipped to the central processing location for the analysis.

"VLBI is one of the most powerful techniques available for the high- resolution imaging of distant radio sources in the universe and for making accurate measurements of the motion of the earth in space," said Alan Whitney, principal scientist at the MIT Haystack Observatory in Westford, Massachusetts. "These capabilities also allow scientists to measure such things as continental drift and to calibrate the orbits of GPS satellites to enable more accurate position measurements on the surface on the Earth. Advanced optical networks like HOPI and DRAGON, will undoubtedly open new doors for radio-astronomy observations and important science."

As a part of Internet2's mission to design and deliver an advanced network infrastructure to meet the emerging needs of the research and education community, Internet2 has built the HOPI nationwide testbed to investigate next-generation network architectures that combine the best qualities of optical and packet technologies. The testbed is a model for the future of Internet2's high performance Abilene network which serves as a platform for both experimental networking applications as well as stable production IP services.

"The HOPI testbed has far-reaching applications in the scientific, engineering, and medical arenas which have come to require far more sophisticated network and resources than those previously available," said Rick Summerhill, co-chair of the HOPI design team and Internet2 director of network research, architecture and technologies. "Internet2's HOPI investigation represents a new paradigm in networking that goes well

beyond traditional production services of today. In doing so, we hope to catalyze a new era of advanced applications which at this point have only been imagined."

Not only will the demonstration highlight the capability to provision on- demand light paths within an administrative domain, but it also proves for the first time, the ability to provision those optical circuits across multiple network administrative domains for global data transmissions. Utilizing DRAGON-developed inter-domain Generalized Multiprotocol Label Switching (GMPLS) capabilities, which provides control plane capabilities, automated end-to-end circuit provisioning, and management of network resources, the optical routes were seamlessly connected across scientific, HOPI and DRAGON domains. The paths also crossed UKLight, SURFnet, NorthernLight, Nordunet, SUnet, JGN2, StarlLight, GIG-EF, and BOSnet.

"We believe the control plane technologies DRAGON has developed and integrated into HOPI pulls together a number of efforts within the R&E community and the international Internet standards bodies to show that these dynamic hybrid network architectures are indeed viable and of great value to the scientific and academic communities," said Jerry Sobieski, lead coordinator of the HOPI Testbed Support Center and project manager for the DRAGON Project. "This demonstration opens the door for both significant advances in radio astronomy and geodesy as well as establishes a foundation on which the global networking community can expand the scope and availability of these capabilities."

Designed to model future optical networking infrastructures, the HOPI testbed utilizes facilities from Internet2's Abilene Network, the National LambdaRail (NLR) infrastructure, The MAN LAN exchange point, and regional optical networks. HOPI nodes, deployed in Seattle, Chicago, New York City, Los Angeles and Washington D.C., each consists of a 10-Gigabit Ethernet switch provided by Force10 Networks, a fiber cross connect switch provided by Glimmerglass, and HP's servers to measure network performance and manage control plane capabilities. For more information on the HOPI project visit its Web site

The radio astronomy demonstration will be held during SC|05 at the Internet2 booth at the following times:

Internet2 BOOTH#2435:
Wednesday, November 16, 2005 - 11am-2pm PT
Thursday, November 17, 2005 - 11am-2pm PT

###

SuperComputingOnline.com

Terrascale Technologies introduces the fastest clustered storage

Wednesday, Nov 16 @ 15:19 PST

Terrascale Technologies, developers of next generation data serving solutions for enterprise-class cluster and grid applications, today announced that it will demonstrate its new Storage Brick hardware platform - for the first time ever - at Supercomputing 2005 (SC|05).

Storage Bricks are data servers loaded with Terrascale's revolutionary TerraGrid software. This hardware/software solution forms virtual network disk drives that offer a single data repository (known as a 'global name space file system') up to 18 Exabytes (18 million Terabytes) in size with input/output (I/O) rates that exceed 10 Gigabytes per second. This unprecedented performance makes TerraGrid-enabled Storage Bricks the fastest clustered storage solution available today.

Previously, system architects had to choose between A) Direct attached storage resources (DAS or SAN) to individual computer systems - costly because a single copy of data could not be accessed by more than one system, or B) Sharing data via legacy network attached storage (NAS) protocols (for example, NFS) - reduces costs by increasing access but limits overall performance. By providing previously unattainable levels of shared data throughput, TerraGrid- enabled Storage Bricks eliminate the need to choose.

"Bundled with our ground-breaking TerraGrid 2.0 software, Storage Bricks epitomize the second generation of storage networking solutions," said Gautham Sastri, President and CEO, Terrascale Technologies. "Enterprises now have increasingly faster networks and extremely powerful computer systems that demand data delivery rates that exceed the capabilities of first generation storage networking solutions. TerraGrid-enabled Storage Bricks help ensure that these computing environments are optimized by moving data at rates that are literally a measured order of magnitude above that of first generation products, and are ideally suited to a variety of data-intensive applications. Our customers have found that they can easily deploy Storage Bricks and eliminate legacy solutions, like NFS servers and other complex parallel file system products, to greatly improve productivity and accelerate ROI," Mr. Sastri added.

Visit Terrascale at SC|05 - Washington State Convention and Trade Center, Seattle, WA, Booth 2031 through Friday Nov. 18 at 4 PM.

SuperComputingOnline.com

Starbridge and SGI Accelerate FPGA Computing Solutions

Tuesday, Nov 15 @ 17:24 PST

Starbridge Systems today announced a strategic relationship with Silicon Graphics to further accelerate algorithm computation using field- programmable gate arrays, or FPGAs. SGI will provide Starbridge's Viva development software under a temporary license as part of its RASC (Reconfigurable Application-Specific Computing) development kit. The combined solution is capable of drastically decreasing compute time for computationally intensive applications over non-optimized systems. Starbridge will hold demonstrations of the collaborative system at its booth during SC|05 International Conference for High Performance Computing, Networking and Storage, November 12 - 18 in Seattle.

"We are very pleased to be teaming up with Starbridge to provide the tools to make the power of FPGAs accessible to software developers, researchers, and hardware designers," said Ron Renwick, Configurable Computing Product Manager, SGI. "Our relationship with Starbridge makes high-performance reconfigurable computing accessible to the entry-level user as well as the power user."

SGI RASC technology, combined with Viva, can drastically increase application performance over conventional systems, a solution that Starbridge has coined Hypercomputing. Viva unlocks the power of SGI RASC technology by allowing those that are not familiar with low-level FPGA development techniques to easily develop highly efficient, high-performance FPGA applications that run directly on SGI Altix servers and Silicon Graphics Prism visualization systems equipped with SGI RASC technology.

Research institutions and government agencies have been searching for a way to rapidly accelerate large, computationally intensive applications or algorithms in order to decrease the time and expense required. FPGAs are recognized in the industry for their ability to handle millions of instructions per second in parallel fashion. However, up until now, FPGA development has required the use of low-level, text-based design languages, such as VHDL or Verilog, that were originally designed for circuitry layout rather than for high-performance or embedded computing.

Starbridge's Viva development software's graphical user interface simplifies the development of complex, supercomputer-realm applications, providing a platform for creating reusable and expandable library objects written for numerical methods.

"We are very excited that SGI chose Viva for its RASC development kit," said Kurt Dobson, CEO of Starbridge. "This is another proof-point that shows the industry has a need for a solution that accelerates the solving of complex applications and algorithms."

SuperComputingOnline.com

Texas Memory Systems' demonstrates new InfiniBand interface

Tuesday, Nov 15 @ 16:13 PST

Texas Memory Systems, makers of the World's Fastest Storage, announced it has developed solid state storage technology that utilizes an InfiniBand interface, an open I/O architecture that provides scaleable performance of 2.5 Gigabits per second to 120 Gigabits per second. Texas Memory Systems will be demonstrating their 4x InfiniBand interface this week at SC|05, the Super Computing conference in Seattle, Washington. Utilizing 4x InfiniBand a single RamSan solid state disk will offer up to 3-Gigabytes per second of sustained random data access over just four interface ports.

The InfiniBand interface will allow Texas Memory Systems' RamSan solid state storage to natively connect to the high-bandwidth, low latency servers used in high performance computing and Oracle grid computing environments. Currently, InfiniBand enables increased network bandwidth of up to 10 Gigabits, provides redundant connectivity support, allows for shared resources, and offers lower CPU utilization.

"High performance computing environments are increasingly adopting InfiniBand as the network of choice for server to server interconnects due to its low latency," said Woody Hutsell, Texas Memory Systems' Executive Vice President. "Conveniently, this same low latency network has enormously high bandwidth which perfectly accommodates our high bandwidth solid state disk systems and is critically important to our plans for a next generation disk storage system that we will announce next year."

Texas Memory Systems is working closely with other InfiniBand technology providers to ensure flawless interoperability across every component in the system. Host adapters from one such provider, Mellanox, will be used by Texas Memory Systems as part of their demonstration at the SC|05 conference.

"Texas Memory Systems' native InfiniBand solid state storage opens storage I/O bottlenecks for bandwidth-hungry applications bypassing gateways or separate Fiber Channel SANs," said Thad Omura, Vice President of Product Marketing for Mellanox Technologies. "In addition, the convergence of computing and storage traffic on the same InfiniBand fabric simplifies the network and eases the management of the entire cluster."

InfiniBand-based RamSan solid state disk systems are expected to be generally available in early 2006.

HP Makes High-performance Computing More Affordable and Productive

Tuesday, Nov 15 @ 16:10 PST

HP today introduced visualization technology designed to make high-performance computing (HPC) more affordable and productive for a broader range of customers. Launched at the SC/05 Supercomputing Conference, the HP Scalable Visualization Array (SVA) uses industry-standard components and leverages advances in clustering, graphics and networking technology to allow scientists and engineers to visualize and analyze complex sets of data. HP SVA is a key part of the HP Unified Cluster Portfolio, which is designed to integrate computation, data management and visualization in a single cluster environment, greatly enhancing productivity in HPC environments. Driven by this portfolio, HP continues to lead the HPC market in revenue with more than 30 percent market share, according to second quarter 2005 figures released by IDC.

Customers in aerospace, oil and gas, automotive, manufacturing, defense, medical and scientific research organizations using the HP SVA can match their growing computational capability with powerful visualization tools. SVA technology can produce a vast, high- resolution display wall of 100 million pixels or more at approximately half the cost of competitive proprietary products.

"Our ability to generate large data sets has outstripped our ability to visualize them, both for deriving science and verifying correctness," said Hugh Couchman, scientific director, SHARCNET, a prominent HPC institute in Ontario that conducts leading-edge research in science and engineering. "With HP SVA, integrated with the HP Scalable File Share and XC software, we will now be able to visualize and manipulate huge data sets in real time as an extension of computing and data management in the cluster."

Along with the new array, HP announced additional enhancements to the HP Unified Cluster Portfolio, including:

-- Version 3.0 of the HP XC System Software, including support for workstations operated as SVA nodes, large node options and new distributed management and provisioning features;

-- Enhancements to HP StorageWorks Scalable File Share, including support for dual-gigabit Ethernet and Advanced Data Guard protection;

-- Support for Microsoft Windows Compute Cluster Server 2003;

-- Updates to HP ClusterPack for HP-UX, with new tools for easier custom configuration and console management;

-- Expanded options for HP Cluster Platforms, including systems with dual-core Intel Xeon processors and HP Cluster Platform Express;

-- Novell Validation Suite HPC and Grid Project updates.

The HP SVA consists of a cluster of HP workstations running Linux, industry-standard graphics cards and network adaptors, and an integrated software system. System software includes HP XC System Software for cluster infrastructure and workload management, and HP Scalable Visualization Array software for configuration and job management. HP StorageWorks Scalable File Share and HP Remote Graphics software are optional.

"The HP Scalable Visualization Array reduces the cost of visualization solutions through the use of industry standards and interoperable products," said Winston Prather, senior vice president and general manager, High Performance Computing Division, HP. "By increasing scalability and flexibility, and integrating computation, data management and visualization with our Unified Cluster Portfolio, HP is meeting the needs of more customers in more application areas and fostering faster scientific discovery around the globe."

HP SVA with HP SFS and XC System Software

In concert with the SVA launch, HP announced a new release of HP XC System Software and updates to HP SFS.

XC version 3.0 supports HP workstations as SVA visualization nodes and large 16-processor node options, providing a single, managed cluster system that can support distributed and SMP applications. New features also include distributed management services, which allocate tasks to multiple nodes to improve system performance and availability, and expanded smart provisioning, which automatically determines cluster topology and the optimal configuration of system services.

HP SFS updates include support for dual-gigabit Ethernet, providing a 50 percent increase in input/output bandwidth for Ethernet connected clusters; high-resiliency Advanced Data Guard protection, which protects against even double-simultaneous disk failures; support for double-density disks, providing up to 1,000 terabytes of storage per file system; and improved resiliency to better detect network failures.

As part of a SC/05 StorCloud demonstration, HP plans to showcase a 64- terabyte, 40 gigabit-per-second HP SFS server connected over a Cisco Systems-powered wide-area-network to a cluster of HP Integrity servers located at Pacific National Labs 200 miles away.

HP Unified Cluster Portfolio enhancements

HP also announced that it will support Windows Compute Cluster Server 2003, which is scheduled for release in the first half of 2006. With the addition of Windows Compute Cluster Server 2003 to its Unfied Cluster Portfolio, HP will offer customers strong choices of HP-UX, Windows and Linux to meet their HPC operating environment requirements. Windows Compute Cluster Server 2003 is expected to be available on HP Cluster Platform 3000 and 4000 systems and Cluster Platform 4000 blade systems.

The HP Message Passing Interface (HP-MPI), a leading implementation of the MPI standard and critical for creating and porting parallel applications, will be ported to Windows Compute Cluster Server 2003. In addition, HP and Microsoft are working with software vendors to port and optimize 64-bit applications to Windows Computer Cluster Server 2003.

In support of HP-UX, HP ClusterPack 2.4 offers new tools for easier custom configuration and console management. It also includes many enhancements for easier deployment of HP-UX clusters in production environments where interruptions from adding or deleting compute nodes must be minimized. It is also now integrated with HP System Insight Manager software.

The HP Cluster Platform 3000 now offers dual-core Intel Xeon processors as an option in the HP ProLiant DL380 control nodes. Dual- core AMD Opteron processors are available with both the HP Cluster Platform 4000 and the blade-based HP Cluster Platform 4000BL. HP Cluster Platforms are available in clusters of up to 1,024 nodes.

HP Cluster Platform Express is a faster, easier way to configure and order single-rack cluster solutions compared to traditional customization methods. Customers can select from a menu of popular cluster components, which are then factory assembled to reduce integration costs while shortening the time to deployment.

More information about HP SVA and the HP Unified Cluster Portfolio is available at its Web site.

Novell Validation Suite - HPC and GRID Project updates

In other HPC news, HP announced that the infrastructure software for the Novell Validation Suite - High Performance Computing and Grid Project, announced in March 2005 by HP and Novell, has been validated.

To date, HP and Novell and participating software vendors have validated a range of infrastructure applications, including Altair, PolyServe Matrix Server and Cluster Volume Manager, Scali Connect and Scali Manage, as well as Grid resource managers from Axceleon, Data Synapse, TurboWorx and United Devices. CAE applications from ABAQUS, ANSYS and MSC NASTRAN have also been validated for the project.

The Novell Validation Suite program focuses on defining and certifying integrated multi-application and platform stacks on Novell SUSE Linux to reduce the risk of deploying complex Linux solutions in HPC environments.

###

Myricom and Extreme Networks Target High Performance Computing

Tuesday, Nov 15 @ 15:57 PST

Myricom and Extreme Networks are teaming up to help customers with High-Performance Computing (HPC) and mainstream networking applications with jointly tested network solutions that achieve 10-Gigabit performance. The two companies will demonstrate interoperability in booth #436 at the SC|05 event being held in Seattle this week. Upon entering an agreement where Myricom has joined Extreme Networks'(R) Go Purple Partner Program, the companies conducted a comprehensive series of tests combining high-availability switches from Extreme Networks with Myricom's Myri-10G Network Interface Cards (NICs) and achieved wire-speed performance at 10-Gigabit speeds.

"Demonstrating the performance results of our combined solution provides customers with information to confidently select Extreme Networks and Myricom," said Tim Nolan, Extreme Networks' vice president of business development. "Now users can obtain the high availability and performance of our network switches for High Performance Computing clusters."

Wire-speed Performance at 10-Gigabit Speeds

Joint testing encompassed multiple product sets including the latest 10-Gigabit solutions from both companies. Extreme's modular BlackDiamond 8800 switch series features the low latency and high densities demanded in high-end-enterprise and HPC applications, while Myricom's newly available Myri-10G solutions extend the price and performance advantages of Myrinet interconnect technology to 10-Gigabit speeds and, for the first time, to native Ethernet.

Tests featuring Extreme Networks' BlackDiamond 8810 modular switch running with the newly available Myri-10G NICs commenced with demonstrating interoperability using both 10GBase-SR and 10GBase-CX4 links. In performance tests of the switch, traffic throughput exactly matched theoretical 10-Gigabit line rates for standard 1500-Byte Ethernet frames and 9-KByte jumbo frames. The same wire-rate results were seen in unidirectional tests, bidirectional send/receive tests, and multiple-port patterns, showing that the BlackDiamond 8810 switch is truly a wire-rate switch. For these tests of switch throughput, the NIC firmware was operating in a "raw" Ethernet mode in which the data bypasses the host's IP protocol stack. Additional performance tests of the Myri-10G 8-lane PCI-Express NICs were performed between dual-Opteron servers running Red Hat Enterprise Server 4.0, kernel 2.6.9-22.0.1.ELsmp. With the standard Myricom driver and firmware, the netperf benchmark showed TCP/IP data rates in a range from 9.6 to 9.8 Gigabits/s, and UDP/IP data rates in excess of 9.9 Gigabits/s.

"We've found Extreme Networks' switches very easy to configure, and their wire-rate performance is well suited to the rigorous requirements of High Performance Computing," says Dr. Chuck Seitz, founder and CEO of Myricom. "Working together and certifying the performance of joint high-speed solutions positions the companies well as the needs of traditional HPC users and mainstream data centers converge."

As part of their joint efforts, Myricom and Extreme Networks have launched programs to train their respective sales and marketing teams on each vendor's products.

###

BlueArc to Demonstrate World's Fastest Network Storage System

Tuesday, Nov 15 @ 15:54 PST

BlueArc Corporation, provider of the world's highest performance network storage systems, today announced that the company's Titan Storage System, measured as the world's fastest storage system and file system, according to multiple third-party benchmarks, will be integrated with SuperComputing 2005's StorCloud, the high performance computing industry's collaborative effort to deliver an unprecedented heterogeneous "petabyte-scale" storage farm to enable showcase applications for demonstration purposes demonstrating interoperability and massive performance. For the StorCloud exhibit, Titan will be deployed as a NAS gateway, front-ending Sun Microsystems' StorageTek disk, as part of a joint solution delivered by the two companies.

For the fifth consecutive year, BlueArc will be exhibiting at the high performance computing industry's largest annual conference, anticipated to draw more than 9,000 attendees. The SuperComputing conference is unique in its combination of a state-of-the-art technical program and premier industry and research exhibition. For the past 17 years, the event has brought together scientists, engineers, systems administrators, programmers, and system managers to showcase innovative developments that are sparking new ideas and new industries. BlueArc's Titan Storage System, the NAS system of choice for high performance computing environments, both commercial and academic, will be on display in booth #707. Titan will also be a key component of StorCloud, located in booth #1116.

After a successful debut at SuperComputing 2004, StorCloud continues as a special initiative for SuperComputing 2005 to build a high performance computing storage capability showcasing HPC storage technologies and applications. Portions of StorCloud will comprise state of the art heterogeneous devices and technologies to build a virtual on-site "storage on request" capability to support researchers and demonstrate high bandwidth applications at the conference. Other portions will be dedicated to supporting specific data intensive applications with predetermined configurations and requirements.

Since the product's introduction in early 2004, Titan has gained significant traction in the world's most challenging computing environments -- delivering network storage performance and throughput capable of enabling massive compute clusters to achieve their full potential, and eliminating bottlenecks commonplace with traditional processor-bound NAS architectures. High performance compute clusters are seeing rapid deployment in all industries, ranging from entertainment, life sciences and oil and gas, to the largest commercial enterprises. As the perfect complement to companies' HPC clusters, Titan achieves throughput and I/O performance that eclipses competitive

systems, bound by antiquated architectures.

"Titan is uniquely designed for the high performance computing market, delivering unparalleled performance and capacity, enabling researchers, scientists and project managers the ability to deliver monumental innovations without requiring monumental costs or storage management," Steve Daheb, vice president of marketing for BlueArc, said. "BlueArc is pleased to work together, as a TekAlliance partner, with Sun Microsystems, to provide a collaborative solution for this year's StorCloud initiative at SuperComputing 2005 -- enabling the conference's 9,000 attendees the opportunity to see the world's most-unique storage platform in action."

BlueArc's Titan Storage System is backed by a unique hardware-accelerated architecture that allows customers to scale their storage systems far beyond any other servers available on the market today. Titan's inherent ability to support primary, nearline, and archive storage within the same storage system enables customers to upgrade the system, utilizing the latest disk technology, when needed. Automatically storing data on the most appropriate storage dramatically reduces cost while maintaining access times consistent with application requirements. Titan allows a single file system to grow up to 256 terabytes and delivers throughput of up to 20 Gigabits per second (Gbps). The architecture adapts to changing application or capacity needs, thus simplifying management, accelerating productivity and protecting a customer's long-term investment. Alongside the system's hardware benefits, Titan also features an industry-leading software suite including virtual servers, policy-based data migration, iSCSI support, remote mirroring and a WORM file system.

###

Joint Industry and Government Initiative to Demonstrate Long Distance InfiniBand

Tuesday, Nov 15 @ 15:53 PST

Cisco Systems, Intel Corporation, Lawrence Livermore National Laboratory, Microsoft, Naval Research Laboratory, Obsidian Research, the OpenIB Alliance and Qwest Communications today announced they are demonstrating extended computing resources using InfiniBand technology as part of SCinet at SC05. Sponsored by the IEEE and ACM, SC05 is the premier international conference on high performance computing, networking and storage.

These leading telecommunications and technology organizations are jointly providing the equipment, software and applications for industry and research partners to demonstrate the value of high-performance, low-latency direct access networking at 10 Gigabits per second over a long distance infrastructure. InfiniBand is a high performance, switched fabric interconnect standard for servers and OpenIB is an industry Alliance that supplies open source InfiniBand software. Both are quickly becoming the preferred standard in high performance computing, grid and enterprise data centers.

The demonstration will link servers, clusters, storage systems, switches and optical service platforms over circuits provided by Qwest originating at the Intel data center in DuPont, Wash. and extending over 50 miles to the Washington State Convention Center via the Pacific Wave NorthWest GigaPop at the University of Washington.

During the conference several demonstrations will highlight some of the possible applications for long distance InfiniBand:

-- Remote data center replication

-- High performance interfaces to the WAN for IB-based clusters+supercomputers

-- GRID computing

-- High performance media streaming

-- Campus area InfiniBand: aggregating departmental clusters into super-clusters

-- Campus and metro InfiniBand storage

At each endpoint, InfiniBand over optical (either DWDM or SONET OC192) is converted by Longbow XRs from Obsidian Research. The Longbow enables globally distributed InfiniBand fabrics to seamlessly cross connect by encapsulating 4X **Joint**

InfiniBand over OC-192c SONET, ATM or 10GbE WANs at full InfiniBand data rates. The conversion is totally transparent to the InfiniBand fabric and is interoperable with OpenIB's software stack and subnet manager. The Naval Research Laboratory (NRL) initiated and supported early development of this capability.

"Microsoft's keynote demonstration at Supercomputing 2005 showcased the improved productivity for scientists made possible by seamless access from the workstation to structured data stores, personal desk-side clusters for interactive analysis and large heterogeneous pools of computing resources for detailed studies," said Kyril Faenov, director of high performance computing, Microsoft Corp. "The high-bandwidth connectivity to Intel's Dupont location allowed us to seamlessly incorporate a 256-core Intel Xeon cluster running Windows Compute Cluster Server 2003 to the mix of computing resources."

"The goal of this demonstration is to show that InfiniBand and the OpenIB software can support advanced simulation, computing and visualization across wide area networks," said Bill Boas vice-chair of OpenIB and a computer scientist at Lawrence Livermore National Laboratory. "Our researchers need such capabilities to enable the range of simulations they perform for our stockpile stewardship mission for the National Nuclear Security Agency of the Department of Energy."

"InfiniBand is an industry standard that is currently deployed worldwide, and this demonstration highlights the benefits that are driving so many organizations to make InfiniBand their interconnect of choice," noted Jim Pappas, director of technology initiatives, Server Platforms Group, Intel Corporation. "Companies that need high performance, low latency, enterprise-level communication, particularly over significant distances, will be particularly interested in the demonstration."

Cisco Systems has contributed the Cisco ONS 15454 SONET Multiservice Provisioning Platform (MSPP) chassis with Cisco ONS 15454-OC192-LR2 Line cards and Cisco ONS 15454 DWDM Multiservice Transport Platform (MSTP). Intel's contribution includes use of 128-nodes (512-cores) from the Intel Dual-Core HPC Cluster physically located in DuPont, Wash. Based on off-the-shelf technologies, including the next-generation dual-core Intel Xeon processor and an InfiniBand interconnect, the cluster represents a new era that rapidly increases performance while reducing or holding steady the requirements for power, heat and floor space. Industry collaborators and end users can access the machine through the Intel Remote Access Service and use it to test drive their codes and accelerate their move to Intel multi-core computing. Alone, this cluster delivers theoretical peak performance of 3.2 teraflops. NRL and Obsidian Research donated the Longbow XR InfiniBand range extenders and Qwest provided the optical fiber infrastructure.

###

SuperComputingOnline.com

Fakespace acquires visualization software company

Tuesday, Nov 15 @ 15:31 PST
Today, at Supercomputing 2005, Mechdyne Corporation announced it will acquire VRCO Inc., a developer of interactive 3D visualization software. The addition of VRCO's rich history and expertise, in conjunction with Mechdyne's acquisition of Fakespace Systems in 2003, reaffirms the company's position as the world's largest and most comprehensive dedicated visualization company, with a reputation for exceptional services and support, the industry's most recognizable brands and the world's largest installed base of advanced visualization users.

"In the 1990's, Fakespace and VRCO commercialized the technologies that introduced the world to immersive and interactive visualization. Today, our products help designers, engineers and researchers work with highly complex computer-generated data in new and intuitive ways," said Chris Clover, President and CEO of Mechdyne Corporation and Fakespace Systems Inc. "The deal brings together two pioneering and innovative companies to form the world's premier provider of consulting services, software and visual systems development."

VRCO's software toolkits, end-user applications, and professional services are a synergistic fit with ongoing efforts by Fakespace to develop enabling software which targets interactive, large-scale visualization systems. For example, Fakespace recently introduced Conduit, a graphics distribution middleware that can make virtually any graphics based software application VR-aware and cluster computing compatible. The benefits of the software products from both companies include shortened design cycles, improved collaboration and increased data insight. Fakespace Systems' broad portfolio of visualization systems includes the highly popular FLEX, reconfigurable immersive environment, PowerWall high resolution displays and the Beacon line of cost effective projection technologies. Major Fakespace and VRCO customers include Los Alamos National Labs, Boeing, Purdue University, Petronas, NASA, Chevron, and numerous universities and research laboratories throughout the world.

"VRCO and Fakespace have been strategic partners for almost a decade and this merger was the logical next step for enhanced growth of our companies," said Cathy Lascara, Chief Operating Officer of VRCO. "We are anxious to apply our collective capabilities to meet the real world challenges of our customers in an unprecedented manner."

Mechdyne will maintain the VRCO brands and its offices in Virginia Beach. Fakespace and VRCO staff will be integrated under the Mechdyne parent company with the Fakespace division focused on large-scale visualization system integration and VRCO focused on software products and services. Together, the two companies will form the foundation of "Mechdyne Solutions", an organization dedicated to providing its customers with world-leading turnkey systems, software, and services. Senior executives from VRCO will retain their roles in the VRCO management team.

SuperComputingOnline.com

Clusters and Storage Revolution Brings Transformative Business Benefits

Tuesday, Nov 15 @ 15:22 PST

Supercomputing 2005 marks a key inflection point in the broad adoption of high-performance computing and storage in the mainstream enterprise market. Isilon Systems, the leader in clustered storage, today announced that Microsoft and Isilon will be demonstrating the use of the Isilon IQ clustered storage systems with Microsoft Windows Compute Cluster Server 2003. Isilon and Microsoft will demonstrate the powerful role clustered computing and storage can play in accelerating digital workflow, fostering deeper collaboration within the enterprise and significantly reducing IT and infrastructure costs. Isilon also announced today that Digital Dimension, ResearchChannel and ZOIC Studios are among a number of Isilon customers benefiting from clustered storage solutions in Microsoft Windows computing environments.

"Microsoft believes the clustered computing and storage revolution is currently underway as we watch cluster-based computing expand rapidly from primarily academic and governmental deployments to broad and diverse adoption across the enterprise market," said Kyril Faenov, director of high-performance computing, Microsoft Corp. "We are pleased to be working with Isilon to bring the benefits of high-performance computing and clustered storage to the broader marketplace."

"There are three macro-trends driving the clustered storage revolution -- the explosive growth of unstructured data and digital content, a fundamental technology paradigm shift to clustered computing architectures and the strong desire to pair industry-standard hardware with intelligent software to achieve dramatic cost savings and productivity gains," said Sujal Patel, co-founder and CTO, Isilon Systems. "Enterprise customers across a broad range of industries are benefiting from the shared vision of Microsoft and Isilon and will harness the power of these next-generation clustered computing and storage platforms to transform their businesses."

"Applications that demand rapid access to information in large file sets are among the most aggressive consumers of storage capacity in today's enterprises. As IT managers deploy extensive clustered computing systems in more general business applications, they need storage solutions that can keep pace with performance needs," said Richard Villars, Vice President of Storage Systems research at IDC. "Solutions like Isilon's clustered storage products will play a critical role in meeting the capacity and throughput requirements of leading applications."

ZOIC Studios, an award-winning visual effects and computer graphics animation studio for commercials, music videos, feature films and episodic television, uses more than 75 terabytes of Isilon IQ clustered storage to support one of the world's largest, high-

performance Windows computing clusters -- with more than 1,000 computing nodes. By combining the power and simplicity of Windows computing with Isilon IQ's high performance and scalability, ZOIC has created a complete clustered computing and storage solution that enables them to realize the full value of their digital workflow, advancing the state of the art for visual effects and animation.

ResearchChannel, a nonprofit media and technology organization that connects a global audience with its participant research and academic institutions, is available to more than 21 million households through cable and satellite distribution. Its website, with users in over 70 countries worldwide, provides programs on demand and through a live webstream over the Internet and Internet-II backbones. ResearchChannel has deployed Isilon IQ as the heart of its storage infrastructure to support its Windows-based digital workflow. During Supercomputing 2005, ResearchChannel will demonstrate its ultra-high bit-rate, high-definition (HD) streaming video technologies, which use Isilon IQ clustered storage and run on the Microsoft Windows operating system. The company will showcase live and on-demand video between Amsterdam, Australia, Tokyo, the University of Washington and its booth at the Seattle Convention Center. While conventional HD streams at less than 20 megabits per second, ResearchChannel will stream its HD video over the Internet-II at up to 1500 mbps - which is more than 20 times the size of broadcast HD signals. These video and networking technologies promise to pave the way for a wide range of applications, including advanced video conferencing and telemedicine.

Digital Dimension, a two-time Visual Effects Society and five-time Emmy award-winning visual effects studio specializing in 3D animation, motion graphics and visual effects for film, television and interactive media, has deployed Isilon IQ as the primary storage for its wide-array of Windows-based, high-performance computing projects. The combination of Isilon's state-of-the-art clustered storage systems and Microsoft's ubiquitous Windows platform has enabled Digital Dimension to realize dramatic performance benefits for its rendering and animation projects, leading to significantly increased productivity and collaboration as well as reduced costs.

Windows Compute Cluster Server 2003 is Microsoft's official entry into the world of high-performance computing (HPC), and is designed to accelerate time-to-insight by providing an HPC platform that is simple to deploy, operate, and integrate with existing infrastructure and tools.

Isilon recently announced the next generation of its award-winning Isilon IQ family of clustered storage systems -- the Isilon IQ 1920, 3000, 4800 and 6000 -- which are powered by the latest version of Isilon's award-winning OneFS distributed file system software and use high-speed, low-latency InfiniBand for intracluster communication.

###

Force10 Networks Demonstrates Secure Line-Rate 10 Gigabit Ethernet Networking

Tuesday, Nov 15 @ 15:19 PST

Force10 Networks today announced that it will demonstrate the industry's first line-rate 10 Gigabit per second intrusion prevention and detection technology along with industry-leading 10 Gigabit Ethernet density this week at SC05, the premier international conference on high performance computing, networking and storage. The Force10 TeraScale E-Series family of switch/routers will provide more than 80 line-rate 10 Gigabit Ethernet ports to build out the backbone for SCinet, the conference's production-quality network, while its intrusion prevention and detection technology will secure the high performance network.

"SC05 is one of the premier events to showcase the future of high performance computing and the increasing role data security will have for governments, education and businesses," said Debbie Montano, director of research and education alliances at Force10 Networks. "For the high performance community, building a 10 Gig network now also means that it must be secured at line-rate speeds."

Today, Force10 announced that it has acquired MetaNetworks Inc., a pioneer in flexible intrusion detection and prevention technology. With the acquisition of the industry's only line-rate 10 Gigabit intrusion prevention system (IPS), Force10 can both build and secure high performance networks like SCinet.

For this year's conference, Force10 is providing more than 80 line-rate 10 Gigabit Ethernet links to build a high performance network on the show floor. As the foundation of the network, the TeraScale E1200 will process nearly all traffic on the network, which is expected to exceed 400 billion bits per second. With the TeraScale E-Series providing scalable 10 Gigabit Ethernet density at the core of SCinet, the conference's high performance network that enables real-time demonstrations, and Force10's line-rate intrusion prevention technology deterring malicious traffic, SC05 is providing its exhibitors with security that does not compromise network performance.

In addition to providing the high performance network core, the Force10 TeraScale E-Series is delivering 40 line-rate 10 Gigabit Ethernet ports to exhibitor booths for demonstrations and the annual Bandwidth Challenge, which highlights and rewards the latest advances in applications that dynamically utilize multiple computing resources to solve problems. Participants in the Bandwidth Challenge use the SCinet infrastructure to demonstrate practical high performance computing applications that require a significant amount of throughput.

Several other Force10 customers are also leveraging their networks for the Bandwidth

Challenge. As a key exchange point for the world's research and education networks, the StarLight optical network exchange, which leverages the Force10 TeraScale E-Series at the core of its network, will be called upon for several of the demonstration and challenge entries. Additionally, TeraGrid will use the collective bandwidth from its eight sites, six of which rely upon the TeraScale E-Series, for its entry into this year's challenge.

Among the booth demonstrations that will take place at the conference is one that leverages Internet2's nationwide Hybrid Optical and Packet Infrastructure (HOPI) testbed, in which the Force10 TeraScale E-Series is deployed. In the first major demonstration of the testbed, Internet2 will, for the first time, attempt to connect four radio telescopes distributed across three continents via dynamically provisioned dedicated optical circuits for real-time analysis of radio astronomy data.

The TeraScale E-Series, which provides industry-leading Gigabit and 10 Gigabit Ethernet densities, will be on exhibit in Force10's booth (#636). Additionally, Force10 will also have on display the industry's only line-rate 10 Gigabit intrusion detection and prevention technology.

SC05 brings representatives from many technical communities together to exchange ideas, celebrate past successes and plan for the future. The conference takes place Nov. 12-18, 2005, at the Washington State Convention and Trade Center in Seattle.

###

SuperComputingOnline.com

BioTeam to Port its Award-winning Software to Microsoft Compute Cluster Server

Tuesday, Nov 15 @ 15:14 PST

The BioTeam, an internationally recognized consulting collective dedicated to delivering vendor-agnostic informatics solutions to the life science industry, today announced that it will be amongst the first solution providers to offer a Life Science solution on the Microsoft's new Windows Compute Cluster Server 2003. The BioTeam will be demonstrating a beta version of its iNquiry software for Windows Compute Cluster Server 2003 at Supercomputing 2005 this week. iNquiry is BioTeam's ready-to-use cluster and web portal optimized for use in life science informatics. The portal installs more than 150 applications to instantly deliver a core of applications that scientists and researchers rely on. iNquiry simultaneously provides access the power required to perform complex computing computation together with simple access to the applications scientist routinely use.

"Accelerating scientific discovery is an important application of Microsoft's new Window Compute Cluster Server 2003," said Kyril Faenov, Microsoft's Director of High-Performance Computing. "We see the BioTeam's iNquiry software as a good example of a turnkey, high-performance computing solution on the scientific desktop."

"Windows Computer Cluster Server 2003 has everything needed to quickly deploy a Windows-based cluster. The combination of Microsoft's product and iNquiry will be a solid and compelling compute solution platform for researchers and scientists," said Michael Athanas, Ph.D., founding partner of the BioTeam. "I was pleasantly surprised by the solid core cluster software components in the product."

The BioTeam plans to release iNquiry shortly after the general availability of Windows Compute Cluster Server 2003, which is targeted for the first half of 2006.

Department of Defense Achieves Supercomputing Performance Milestone

Tuesday, Nov 15 @ 13:37 PST

Linux Networx, The Linux Supercomputing Company, announced today that a 2,048-processor cluster system it delivered last year to the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) has been optimized and achieved performance of 10.65 trillion operations per second (teraflops). The system is the most powerful supercomputer at the Army Research Laboratory Major Shared Resource Center (ARL MSRC) and it ranks as the 24th fastest supercomputer in the world according to the Top500.

The Linux Networx system, known as the John Von Neumann (JVN), was deployed as part of the Technology Insertion 2004 (TI-04) program, an initiative to modernize the Department of Defense high performance computing (HPC) capabilities. After a successful acceptance with a theoretical peak of 14.7 teraflops, the ARL and Linux Networx team continued to work together to fine tune the system for optimal performance.

"What we really have with Linux Networx is a partnership, and we have been working together to optimize the JVN system over the past year for remarkable performance and reliability," said Thomas Kendall, lead systems engineer at the ARL MSRC. "We are impressed that we have been able to achieve an actual performance of 10.65 teraflops and 72 percent efficiency by tuning this machine, but the real value lies in the fact that the system is now a more effective tool for work on DoD's applications."

Since it went into production in May 2005, the ARL has utilized JVN to support various applications including electromagnetic, penetration mechanics, and aerodynamics. Electromagnetic codes such as ICEPIC and Monocode scale extremely well on JVN, and a project utilizing ICEPIC has used more than one million processor hours in the last five months. The ICEPIC project supports the design and development of high power microwave sources for the disruption/destruction of electronics. Another project using Monocode studies the improvement of millimeter-wave radar signature prediction for ground vehicles.

"JVN provides an increase in computing capability that gives DoD scientists and engineers the ability to solve complex physics-based problems in a shorter timeframe -- thus providing our soldiers with the best technology and weapons systems available and giving them more time for training, enhancing their readiness and potentially saving lives," said Charles J. Nietubicz, director of the ARL MSRC.

"Reaching this milestone in partnership with the DoD demonstrates the Linux Networx

commitment to helping our customers achieve their supercomputing goals," said Dean Hutchings, president and COO of Linux Networx. "The JVN system was tested extensively at our Solutions Center, and after it was delivered we worked with the ARL to tune it to their exact specifications, which allowed us to achieve this landmark performance. This system helps the ARL solve complex, mission-critical problems faster than ever before."

The HPCMP provides the supercomputer services, high-speed network communications, and computational science expertise that enables the U.S. Defense laboratories, such as ARL, to conduct a wide range of focused research, development, and test activities. The Linux Networx system is a part of this partnership to put advanced technology in the hands of U.S. armed forces more quickly, less expensively, and with greater certainty of success.

The Linux Networx system includes 2,048 Intel Xeon 3.6 GHz processors and Intel 64-bit extension technology. The system also uses Myrinet high-speed interconnect from Myricom and Gigabit Ethernet technology from Foundry Networks, Inc.

###

SuperComputingOnline.com

PathScale InfiniPath-Powered Cluster Shatters HPC Challenge Benchmark

Tuesday, Nov 15 @ 13:25 PST

PathScale, the developer of innovative software and hardware solutions to accelerate high performance computing, today announced that its InfiniPath HTX InfiniBand Adapters have been deployed by AMD in its Developer Center to maximize application scaling on the newly installed Dual-Core AMD Opteron processor-based 'Emerald' cluster. The combination of Dual-Core AMD Opteron processors and the InfiniPath interconnect is already demonstrating unprecedented performance, enabling Emerald to outperform traditional supercomputers in several critical benchmarks in the latest High Performance Computing (HPC) Challenge.

The HPC Challenge, sponsored by DARPA, the National Science Foundation (NSF), and the U.S. Department of Energy (DOE), consists of nine benchmarks that evaluate how HPC systems handle real-world applications. Based on the most recent benchmarks, the AMD Emerald cluster with the InfiniPath InfiniBand interconnect outperformed much larger supercomputer systems. For instance, the 512-core AMD Opteron processor-based Emerald cluster configuration outperformed the highest-end systems from the three leading supercomputing suppliers in the Random Access (GUPs), Random Ring Latency, and Natural Ring Latency benchmarks. These benchmarks are highly sensitive to memory update performance and the speed of network communications, and showcase the clear performance advantages of the InfiniPath interconnect and Dual-Core AMD Opteron processors with Direct Connect Architecture.

The AMD Emerald cluster, supplied by Rackable Systems, is comprised of 144 nodes, each with two 2.2 GHz Dual-Core AMD Opteron processors, for a total of 576 processing cores. Each node is equipped with a single PathScale InfiniPath HTX InfiniBand Adapter connected to a SilverStorm 9120 144-port InfiniBand switch. The AMD Emerald system, which achieved 2.1 TFLOPs on the Linpack benchmark, has been submitted to the Top500 supercomputing list, which ranks the 500 largest supercomputers in the world. The AMD Emerald system is one of the largest publicly accessible Dual-Core AMD Opteron processor-based InfiniBand clusters in the world.

"When matching up the low-latency PathScale InfiniPath interconnect with the low-latency Direct Connect Architecture of AMD Opteron processors, the Emerald cluster produces phenomenal results," said Pat Patla, director, server/workstation marketing, Microprocessor Solutions Sector, AMD. "We have worked with PathScale to showcase this incredible performance through Emerald – the most powerful cluster ever implemented at the AMD Developer Center."

Located at the AMD Developer Center in Sunnyvale, Calif., Emerald is designed to provide AMD's development collaborators and customers a way to benchmark and test

performance-sensitive computing applications using the company's Dual-Core AMD Opteron processor technology. PathScale's InfiniPath is a cluster-interconnect that plugs directly into the HyperTransport interface on AMD Opteron processor-based servers, and designed to dramatically improve communications within the cluster.

"The HPC Challenge benchmark results prove that InfiniPath can scale AMD Opteron processor-based clusters to performance levels that exceed systems from some of today's supercomputing giants," said Scott Metcalf, CEO of PathScale. "These benchmarks further validate the performance advantages of InfiniBand, and should demonstrate to the scientific and engineering communities that they no longer have to rely on proprietary technologies from the traditional, high priced supercomputing suppliers. They can now have an advantage by leveraging AMD Opteron processor-based Linux clusters and the InfiniPath interconnect to build cost-effective systems for their most demanding applications."

The PathScale InfiniPath interconnect helps deliver on the promise of Linux cluster computing by significantly lowering communications latency, helping to improve the performance of complex applications. The technology is enabling scientists, engineers and researchers to more effectively solve a whole new class of computational challenges, from weather modeling and aerospace design to drug discovery and oil and gas research. Today, the InfiniPath interconnect is used by leading scientific and engineering organizations in both the private and government sectors.

###

PathScale InfiniPath Software Release Achieves New Levels of Scalability

Tuesday, Nov 15 @ 13:19 PST

PathScale, the developer of innovative software and hardware solutions to accelerate high performance computing, today announced, at SC05, the release of a low-profile form factor InfiniPath HTX InfiniBand Adapter and the release of version 1.1 of its InfiniPath software. This latest InfiniPath software release fully supports the OpenIB Gen 2 software stack and is designed to maximize application scaling and performance on InfiniBand-based Linux clusters.

In addition to supporting the OpenIB stacks, the new InfiniPath 1.1 software release includes a number of improvements in overall performance and scalability on MPI and TCP applications. PathScale has tested the scalability of dozens of well-known HPC applications on the new 576 Processor AMD Emerald cluster with outstanding results. Using the newly released software, record-breaking performance has been achieved on the HPC Challenge, Pallas, and Topcrunch benchmarks. The 1.1 release also includes support for a broader range of Linux distributions, including support for RHEL4, ROCKS 4.0.0, SUSE 10, Fedora Core 4 and CentOS 4.0. The 1.1 release is immediately available to all existing customers and can be downloaded from the PathScale support web site. OpenIB was developed by the OpenIB Alliance, an industry association chartered to deliver a single, open-source software stack for deploying InfiniBand solutions. The OpenIB Alliance, founded in June 2004, is comprised of the leading computer technology vendors and end-user organizations. OpenIB provides a single, open-source standard that ensures interoperability between software stacks and hardware products. This promotes competition in the market and enables PathScale to help drive broad adoption of InfiniBand, while delivering the most scalable and efficient solutions for high performance computing.

“With the release of the InfiniPath 1.1 software with OpenIB support, PathScale is proving that InfiniBand scales to levels only previously accessible with proprietary interconnects,” said Scott Metcalf, CEO of PathScale. “PathScale is showing its leadership position in the InfiniBand world by showcasing the scalability of InfiniBand and being the first vendor to bring a supported version of OpenIB to market. PathScale is giving the scientific and engineering communities more choices in the marketplace, enabling them to realize the benefits of a truly open standards-based high-performance interconnect.”

PathScale is also introducing a half-height, or low profile, InfiniBand adapter card at a reduced price of $795. Because the InfiniPath interconnect improves the efficiency of the entire cluster, the reduced price makes it the obvious choice for end-users sensitive to both price/performance and scalability. The new card, which can be mounted either

horizontally in a 1U server or vertically in a 2U server, gives cluster users more flexibility for deployment options.

“Voltaire is excited to see Pathscale’s commitment to OpenIB as demonstrated in its latest software release,” said Patrick Guay, vice president of marketing at Voltaire. “Coupled with the availability of the new low-profile card, we feel that the choices our customers can enjoy have once again been expanded.”

The PathScale InfiniPath interconnect delivers on the promise of InfiniBand and cluster computing by significantly lowering communications latency while eliminating bottlenecks that can slow down complex applications. The technology is enabling scientists, engineers, and researchers to more effectively solve a whole new class of computational challenges, from weather modeling and aerospace design to drug discovery and oil and gas research.

###

Absoft Announces Availability of IMSL Fortran Numerical Library

Tuesday, Nov 15 @ 10:47 PST

Absoft Corporation, a leading provider of software development tools, announced today at the Supercomputing 2005 “SC|05” Conference in Seattle availability of Visual Numerics' IMSL Fortran Numerical Library version 5.0 optimized for Linux. The IMSL Fortran Numerical Library is the gold standard mathematical and statistical function library for Fortran programmers developing high performance computing applications. The IMSL Fortran Numerical Library saves development time by eliminating the need to write code from scratch.

Using only two or three of the IMSL Fortran Library's mathematical or statistical routines will more than pay for the product in application development time-savings alone. With this new offering optimized for Linux, the complete IMSL Fortran Numerical Library is now available to all Absoft customers regardless of whether they are using Linux, Windows, or Macintosh OS X platforms. “Absoft and Visual Numerics have been collaborating on high quality Fortran solutions for almost a decade.” said Phil Fraher, CEO of Visual Numerics. “Now, Linux Fortran developers worldwide can leverage the software expertise of our two companies through the reliable and robust combination of the Absoft compiler and the IMSL Libraries.” Visual Numerics' latest IMSL Fortran Numerical Library, version 5.0, includes new powerful and flexible interface modules for all applicable routines.

These new interface modules allow users to utilize the fast, convenient optional arguments of the modern Fortran syntax for 100% of the relevant algorithms in the library, allowing for greater control and faster, simpler development of code. The IMSL Fortran Numerical Library provides full depth and control via optional arguments for experienced programmers, reduces development efforts by checking data-type matches and array sizing at compile time, and gives developers a simple and flexible interface to the library routines as well as speeds programming and simplifies documentation. At the heart of the IMSL Numerical Libraries lie the comprehensive and trusted IMSL mathematical and statistical numerical algorithms. The IMSL Fortran Numerical Library combines over 1000 fully tested and highly optimized algorithms and includes the complete IMSL F90 Library, the IMSL FORTRAN 77 Library, and the IMSL parallel processing features. Utilities are included to simplify large-scale computing with the ScaLAPACK library. Full documentation detailing each routine is included in PDF format.

“As Absoft continues to lead the market in platform availability and choices for deploying applications using the IMSL Fortran Numerical Libraries, the popularity and market demand for this outstanding product continues to grow.” said Absoft executive vice president Jeff Livesay. “Over the past year we have expanded our relationship with

VNI and look forward to introducing to our mutual customers other powerful product combinations from Absoft and VNI in the future." The IMSL Fortran Numerical Library is available from Absoft as a standalone product, or bundled with Absoft's Fortran 95 v9.0 Compiler for Linux on AMD or Intel 32-bit & 64-bit processors. The IMSL Fortran Numerical Library can also be optionally bundles with Absoft's High Performance Computing Software Development Kits (HPC SDK) for AMD Opteron or Intel-based systems running Linux.

###

Sun Microsystems Delivers Leading-Edge Technologies and Tools

Tuesday, Nov 15 @ 10:35 PST

At Supercomputing 2005, Sun Microsystems today announced significant performance upgrades to its Sun Fire x64 (x86, 64-bit) servers powered by the high-performance AMD Opteron processor, new high performance computing (HPC) tools and applications for x64 and SPARC processor-based Sun Fire servers running the Solaris 10 Operating System (OS), alliances with HPC industry leaders, and breakthrough storage technologies to enable rapid deployment of power-efficient Terascale compute clusters. With the help of Sun's innovative HPC solutions and industry-leading alliances unveiled today, customers can maximize productivity, lower risk, reduce cost and complexity in the data center, and speed time to market, while embracing deployment-ready Terascale HPC and grid solutions to unleash the power of innovation.

"Our presence and announcements at Supercomputing this week demonstrate that Sun is a leading player in high performance computing and focused on aggressively growing our HPC offerings," said John Fowler, executive vice president, Network Systems Group, Sun Microsystems. "Sun's industry-defining x64 and SPARC processor-based Sun Fire servers -- combined with our information management solutions and engineering brilliance -- give us a competitive advantage in this market, allowing us to bring no-risk, leading- edge HPC capabilities to customers. Moreover, Sun's legacy of building innovative alliances with companies such as Luxtera, Inc. and the Defense Advanced Research Projects Agency continue to be a driving force behind next-generation compute offerings."

Customers Around the World Leverage Sun's x64 Solutions for Large-scale HPC Clusters

In a separate announcement today, the Tokyo Institute of Technology (Tokyo Tech) has purchased Sun Fire x64 servers to build Japan's largest supercomputer. Tokyo Tech's project will use Sun Fire x64 servers with 10,480 AMD Opteron processor cores, Sun storage technologies and NEC's integration expertise to build the Tokyo Tech supercomputer. The system will help provide researchers with compute power for a wide range of scientific applications, such as analysis of the complex molecular structure of proteins, simulated blood flow diagnosis in human brains, and clarification of the generation mechanism of Earth and planetary magnetic field. This marks Sun's largest HPC win to date.

Furthermore, Sun recently announced that Atlantic Computational Excellence Network (ACEnet) Canada will build the largest HPC network in Atlantic Canada based on the Sun Fire x64 servers. ACEnet's grid will be a service for a wide range of research, including oil and gas, marine engineering, pharmaceuticals, underwater vehicles, physics,

physical oceanography, and computer science throughout Atlantic Canada.

In addition, Sun last week opened its Sun Solution Center for HPC, a testing and benchmarking facility located in Hillsboro, Ore., that is designed to make HPC practical and attainable for a wide array of customers and partners. The facility offers customers access to world-class scientists and algorithm experts who specialize in developing and deploying large-scale HPC solutions based on best practices, and also provides them access to some of the highest-performance x64 and SPARC processor-based systems on the planet. In this unique environment, customers and partners have a low-risk method of deploying and running their applications on a variety of environments, including the Solaris 10 OS, and standard distributions of Linux and Windows. To learn more about the Sun Solution Center for HPC, please visit its Web site.

Sun Fire Servers Outdistance the Competition and Continue to be the Fastest in their Class

The Sun Fire X4100 and Sun Fire X4200 servers are now available with AMD's newest, highest-performing multi-core AMD Opteron Model 285 SE processor, outperforming similar competitor's systems based on dual-core Intel processors. The Sun Fire X4100 server with one AMD Opteron Dual Core processor Model 285 SE comes equipped with the Solaris 10 OS, 36GB SAS disk drive, 1 GB memory, and KVMS enabled service processor, and is competitively priced from $4,161 (USD). More information on the Sun Fire x64 servers can be found at its Web site.

The combination of Sun Fire x64 servers, Solaris 10 OS, and Sun Studio 11 software delivers new world record performance. In addition to more than 50 world record results posted to date, the Sun Fire x64 systems, powered by the latest multi-core processors, have established the following new performance records:

-- World Record SPEC OMPM2001 result for systems with up to 2 sockets and 4 threads
The Sun Fire X4100 server beats IBM and HP systems running compute-intensive applications. The 4-way Sun Fire X4100 server, with two AMD Opteron 285SE processors, surpasses the most recent IBM POWER5+-based p5 520 result by over 34% on the medium size set.

-- World Record among all 2-socket x86-compatible systems on SPECfp_rate2000 benchmark
Based on the real world user applications, the floating point throughput component of SPEC CPU2000 benchmark highlights the superior performance of multi-core Sun Fire X4200 server, when compared to the single or dual-core Xeon-based HP ML570 server, by over 70% and 56% respectively.

-- World Record SPEC OMPM2001 performance for all single socket systems running 2

threads

The Sun Fire X2100 server, the ultimate low-cost platform for developing and running scalable compute intensive applications, demonstrates the multi-core advantage of the AMD Opteron Model 175processor and highlights the superiority of the Solaris 10 OS and Sun Studio 11 compiler software as well as the underlying system's design. The SPECompM2001 results for the Sun Fire X2100 dual-core server scales virtually linear (1.99 times) when compared to the Sun Fire X4100 4-way server result, that was achieved at the same clock frequency.

-- World Record among all 8-core x86-compatible systems on SPECfp_rate2000 test
The Sun Fire V40z server, powered by the latest multi-core AMD Opteron 880 processors, demonstrates linear scalability with the processor frequency, for the floating point throughput component of the compute-intensive SPEC CPU2000 benchmark, surpassing the record previously held by HP by over 6%.

More information on the Sun x64 server and workstation benchmarks, including results for Sun's widely acclaimed Sun Fire X2100, Sun Fire X4100 and Sun Fire X4200 servers can be found at its Web site.

Sun Delivers Advanced Tools for HPC and Terascale Clusters

Sun recently announced several new tools and solutions that further accelerate its position and momentum in the HPC market:

-- Sun customers can now build and order large-scale computing solutions from Sun by the TeraFLOP through its "TeraFLOP to Go" solution, which gives customers a low-risk method to easily build or expand compute clusters by providing innovative, HPC-ready building blocks. The "TeraFLOP to Go" solution is pre-defined and tested to scale from .25 to 2.0 TeraFLOPS, with larger configurations available upon request. Sun will build, test and deliver the ready-to-deploy system to the customer using the Sun Customer Ready Systems (CRS) program. For more information on the "TeraFLOP to Go" program, please visit its Web site.

-- Sun unveiled a sneak-peek and customer early access to the Sun HPC ClusterTools 6 software, a major update to its ClusterTools portfolio. This powerful software solution enables customers to build and deploy large-scale Message Passing Interface (MPI) applications on the entire range of Sun Fire servers, powered by either x64 or SPARC-based processors running the Solaris 10 OS.

-- Sun has significantly expanded its portfolio of HPC information management solutions with the acquisition of StorageTek, and is now uniquely positioned to help enterprises efficiently manage the complexities of data management that are driven by increased

compute power. The UltraSPARC IV+ servers are ideal for HPC applications that require large-scale shared memory to run effectively. SPARC processor-based Sun Fire servers running on the Solaris OS are up to twice as fast as competitive solutions and offer extensive compute, input/output and storage capacity with the lowest power consumption on the market. This expanded Sun portfolio provides customers with added scalability, flexibility and choice in storage and data management with solutions for heterogeneous HPC environments. On display at Supercomputing will be the newly announced T10000 tape drive that offers the highest throughput of any tape drive available today and initial capacity of 500 GB of uncompressed data. Additionally, the industry-leading FlexLine 380 enterprise disk storage system is supporting three different HPC applications at the show with both Fibre Channel and InfiniBand connected storage -- highlighting the commitment to emerging technologies in this space.

Sun and its Alliances Paving the Way to Build the Next Generation of Supercomputers

Sun's powerful and strategic alliances with HPC industry leaders are key to its efforts in designing a productive Petascale system with massive compute bandwidth and ultra-low latency. Sun announced today that it has selected Luxtera Inc., a fabless semiconductor company and world leader in silicon photonics, to collaborate on the development of dense wavelength division multiplexing (DWDM) optical links. Luxtera's DWDM-on-a-chip technology will serve as the building block of future Terabit links in Hero, Sun's High Productivity Computing Systems (HPCS) project, part of the United States' Defense Advanced Research Projects Agency (DARPA) High Productivity Computing Systems program. Sun and Luxtera are demonstrating the world's first DWDM 40 Gigabit per second optical link built entirely in CMOS in the HPCS section of Sun's booth (#1416).

Pursuing further breakthroughs in productivity, Sun is also leading the development of Fortress, a next-generation programming language, to achieve highly productive development and run-time environments on ultrascale systems. Dr. Guy Steele, Sun fellow, Sun Microsystems Laboratories, will deliver an "Introduction to Fortress" seminar on Nov. 15 at 3 p.m. in the Grand Ballroom A of the Sheraton Seattle Hotel (across the street from the Seattle Convention Center).

###

SuperComputingOnline.com

Absoft Announces Enhanced Edition of the High Performance Computing Software

Tuesday, Nov 15 @ 10:40 PST

At today's Supercomputing 2005 conference, Absoft Corporation, a leading manufacturer of advanced software development tools, announced it is working with Intel Corporation on an Enhanced Edition of Absoft's newly released High Performance Computing Software Development Kit (HPC SDK) for 64-bit Intel Xeon processor-based servers and clusters running Linux. The Enhanced Edition will begin shipping in December, 2005. Absoft announced general availability for the Basic Edition of the HPC SDK in June at the International Supercomputing Conference 2005 in Heidelberg.

“We are gaining substantial traction with our increasingly comprehensive line of HPC SDKs, which have essentially the same look and feel across multiple platforms,” said Absoft CEO Wood Lotz, “but the Intel-based SDKs are leaping ahead both in market volume and in the range of powerful software tools included in the new Enhanced Editions for developing the highest performing computing codes on clusters, servers, and grids.”

The Absoft HPC SDK is a complete solution allowing software developers to easily compile, run, debug, and optimize high performance software applications. The HPC SDK includes some of the industry's leading software tools optimized for 64-bit extensions and is preconfigured for easy installation. It includes Fortran and C/C++ compilers from both Intel and Absoft, debuggers, math libraries, pre-built message passing (MPI) libraries, and other development tools. Also included is a comprehensive set of example HPC programs to help new users quickly learn how to take full advantage of the HPC SDK features. The HPC SDKs from Absoft are the only commercially available HPC SDKs that combine tools from multiple vendors to provide customers with true best-of-class solutions. With just three easy mouse clicks, customers can install and begin using the HPC SDK to take advantage of their existing hardware and software investments. Absoft Corporation is the single point of contact, simplifying all pre-sale, post-sale and technical support issues related to the HPC SDK.

The HPC SDK Enhanced Edition includes all of the features of the increasingly popular HPC SDK Basic Edition for Intel Xeon processor-based systems along with additional tools from Intel, including: Intel VTune™ Performance Analyzer to easily streamline code, and Intel Thread Checker and Intel® Thread Profiler for OpenMP and POSIX threads debugging and bottleneck analysis. Also included are open source tools such as Marmot, an MPI thread checker. Future versions of the Enhanced Edition of the HPC SDK will include the FxP parallel MPI debugger from Absoft.

“Absoft's decision to include our products, with our exceptionally strong support for

threading, enables Absoft customers to develop applications today which best exploit the capabilities of current and future platform offerings," said James Reinders, Director of Marketing for the Software Products Division of Intel Corporation. The HPC SDK was developed for 64-bit Intel Xeon processor-based servers and clusters running Linux, such as the IBM xSeries eServer and the IBM eServer Cluster 1350, IBM's leading-edge Linux cluster solutions. The HPC SDK is also available for systems such as the new IBM Grid and Grow™ and the IBM BladeCenter.

###

IBM Continues Global Supercomputing Lead

Monday, Nov 14 @ 15:11 PST

The world's foremost supercomputing authority, the TOP500 Organization, today named an IBM supercomputing system as the most powerful supercomputer in the world. IBM's Blue Gene/L tops the list with an unprecedented sustained performance of 280.6 Teraflops, or trillions of floating point calculations per second.

Joining Blue Gene/L in the TOP500 list's top three supercomputers are IBM's own Blue Gene Watson system at 91.29 Teraflops, and the recently unveiled ASC Purple supercomputer at Lawrence Livermore National Laboratory with 63.39 Teraflops.

IBM supercomputing systems lead the list, delivering for the first time in the history of the TOP500 an aggregate performance of over one Petaflop.

The Blue Gene/L and ASC Purple systems were developed with the Department of Energy's National Nuclear Security Administration and are installed at Lawrence Livermore National Laboratories in California. The ranking of ASC Purple at Lawrence Livermore National Laboratory marks the first time a POWER5-based system has claimed a top-three placement on the list.

"ASC Purple and Blue Gene/L mark the completion of a ten-year challenge to develop supercomputers for entry-level simulations that support a future free from nuclear testing," said Dr. Dimitri Kusnezov, Director of the NNSA Advanced Simulation and Computing program. "Reaching this point is not only a reflection of the importance of our industrial partnerships, but also of the sustained support of the Department of Energy, NNSA and Congress to the mission. With these new tools, the national laboratories are pushing the envelope in computational science, bringing improved confidence to guide decisions on our national nuclear security."

"The complement of simulation capabilities BlueGene/L and Purple bring to the three national labs under the National Nuclear Security Administration's Advanced Simulation and Computing program is of critical importance," said Dona Crawford, associate director for Computation at Lawrence Livermore National Laboratory. "These systems put us on the threshold of a new era in high performance computing, where simulation is the integrating element of the scientific discovery triad. We're excited to have simulation capabilities that will help us to better understand the complex physical phenomena necessary to ensure the safety, security and reliability of the nation's aging nuclear deterrent."

IBM is the leading provider of both installed supercomputing systems with 219 systems as well as total aggregate supercomputing power, with a record total of 1.214 Petaflops.

IBM has five of the systems in the Top 10, including Mare Nostrum, Europe's most powerful supercomputer which is powered by IBM's Power Microprocessor and eServer BladeCenter JS20 -- the only top-ten supercomputer based on blade server technology.

According to numbers compiled by the TOP500 List of Supercomputers, IBM is the overwhelming leader in global supercomputing with 52.77 percent of the total processing power, nearly three-times the power of its closest rival. In addition, IBM is debuting three new Blue Gene systems on the Top500 List.

Princeton University, MIT and Zurich Research Center have installed IBM eServer Blue Gene systems and appear for the first time on the Top500 list.

A unique partnership between Princeton University scientists and information technology administrators has brought one of the world's fastest supercomputers, IBM's Blue Gene, to the University to spur advancements in research. In conjunction with IBM, the University's Office of Information Technology (OIT) collaborated with Princeton researchers to purchase and install the Blue Gene system to aid current and future research solving complex problems in areas including astrophysical sciences, engineering, chemistry and plasma physics.

Since IBM announced the commercial availability of the IBM eServer Blue Gene Solution, a commercial version of the research project, in November 2004, a record number of 19 Blue Gene systems appear on the list. Based on IBM's Power architecture, the IBM eServer Blue Gene Solution is optimized for bandwidth, scalability and the ability to handle large amounts of data while consuming a fraction of the power and floor space required by today's fastest systems. IBM and its partners are exploring a growing list of high performance computing (HPC) applications including life sciences, financial modeling, hydrodynamics, quantum chemistry, molecular dynamics, astronomy and space research and climate modeling for eServer Blue Gene.

Along with Lawrence Livermore National Laboratory, ASTRON, AIST, NIWS, NCAR, University of Edinburgh, San Diego Supercomputing Center, Argonne National Lab, The Ecole Polytechnique Fédérale de Lausanne (EPFL) and the IBM Deep Computing Capacity on Demand Center, these research institutions make up a growing ecosystem of early collaborators dedicated to harnessing Blue Gene's power to advance research.

"By giving our clients access to innovative, affordable and flexible supercomputing power like Blue Gene, POWER5-based p5 575 systems, JS20s and the Deep Computing Capacity on Demand Center, we are providing new resources to drive breakthroughs in business, science and industry," said Dave Turek, vice president, Deep Computing, IBM. "Whether we are talking about improving the accuracy of weather forecasts, designing better automobiles or improving disease research, we are seeing the advent of a new supercomputing age."

Other key indicators of IBM supercomputing leadership:

-- #1 BlueGene/L -- doubling the performance achieved in the June 2005 Top500
-- #2 Blue Gene Watson (BGW)
-- #3 ASC Purple p5 575
-- 219 total systems -- the most by any vendor, and more than Hewlett-Packard, Cray and SGI combined
-- First-ever one Petaflop of aggregate performance, representing 52.77% of total supercomputing power on the Top500 list
-- More systems of any vendor in the Top 10 with five
-- More systems of any vendor in Top 20 with eight
-- More systems in Top 100 with 49
-- Europe's most powerful two supercomputers are IBM systems -- Mare Nostrum in eighth place overall, and ASTRON Netherlands at ninth
-- IBM provides the most commodity clusters on the list, with 158 out of 360
-- 73 systems are based on IBM's Power Architecture technology
-- Power processors total 327,622 out of 731,069 on the list -- over 13% more than from any other single competitor
-- Blue Gene taken by itself would be third highest number of entries, more than Cray, SGI, Dell, and all others except HP and IBM

The "TOP500 Supercomputer Sites" is compiled and published by supercomputing experts Jack Dongarra from the University of Tennessee, Erich Strohmaier and Horst Simon of NERSC/Lawrence Berkeley National Laboratory and Hans Meuer of the University of Mannheim (Germany). The entire list can be viewed at its Web site.

###

SuperComputingOnline.com

Foundry Networks Convenes Inaugural Meeting of HPC Ethernet Alliance

Monday, Nov 14 @ 15:04 PST
Foundry Networks today announced that its HPC Ethernet Alliance partners will be meeting on November 16, 2005 during the SC2005 Conference this week. SC2005 is the premier international conference on high performance computing, data networking and storage being held at the Washington State Convention and Trade Center in Seattle.

In recognition of the growing interest in cluster and high performance computing, Foundry has formed the HPC Ethernet Alliance Partners Program to accelerate and promote the deployment of high-performance Ethernet solutions in HPC and cluster computing environments.

"The HPC Ethernet Alliance brings together industry leaders in Ethernet switching and high performance computing to promote open and mainstream technologies for HPC and cluster computing solutions," said Mark Danckert, Foundry's director of alliances. "During this first meeting of our alliance partners we will set the program's future course discussing a range of topics including: alliance charter, membership, organization, marketing programs and technical initiatives."

The HPC Ethernet Alliance meeting will be held in Room 209 at the Washington State Convention and Trade Center from 8:00 am to 10:00 am on Wednesday, November 16, 2005. This meeting is open to Foundry's HPC Ethernet Alliance Partners.

SuperComputingOnline.com

Extreme Networks Showcases Networking Solutions

Monday, Nov 14 @ 14:58 PST

Extreme Networks is showcasing its high performance network solutions that help provide a resilient, available and secure communications infrastructure for High Performance Computing applications at this week's Super Computing Conference (SC05), booth # 610.

At its booth, Extreme Networks will demonstrate switch and NIC technology that delivers end-to-end Ethernet message passing latency under 15ms, challenging specialized interconnects for mainstream cluster computing applications.

"With our long history of high performance switching platforms leveraging low latency, high availability and wire-speed performance, Extreme Networks is well positioned for the HPC market," said Suresh Gopalakrishnan, vice president of worldwide marketing for Extreme Networks. "Show attendees will have an opportunity to see how our Ethernet solutions are well suited to meet their needs."

Speaking Engagement at SC|05

At SC|05, Extreme Networks' Director of Product Management, Dave Hawley, will present technologies that support the next generation of Ethernet in high performance clusters, including computation, storage, and control data transport.

Topic: The Next Generation of Ethernet for HPC
Session: Interconnect Technologies
Place: Washington State Convention Center, Room: 617
Date and Time: Wed., Nov. 16, 11:30am - 12:00pm

Enterprise Adoption of High-Performance InfiniBand Clustered Storage on the Rise

Monday, Nov 14 @ 14:51 PST

Isilon Systems today announced that leading enterprise customers across a broad range of industries have deployed InfiniBand-enabled clustered storage systems from Isilon using Cisco InfiniBand Server Fabric Switches (SFS). Customers such as Kelman Technologies Inc., DNA Productions and the Incorporated Research Institutions for Seismology (IRIS) are among the more than 90 percent of all Isilon customers who have chosen to deploy InfiniBand- based versions of Isilon's products since they were first introduced to the market in April 2005. Isilon is the first company to implement InfiniBand in a clustered storage solution. Over the last three months, enterprise customers have deployed more than 2,500 ports of Cisco SFS InfiniBand switches as the high-bandwidth, low-latency interconnect for their Isilon IQ clustered storage systems. By integrating Cisco SFS switches with Isilon's state-of- the-art clustered storage architecture, Isilon provides its customers with a high-performance, high-capacity, scalable data storage solution while ushering next-generation data center connection technologies into mainstream enterprise environments.

"We have seen rapid growth in the enterprise adoption of InfiniBand for high-performance server clusters and grid computing solutions," said Stu Aaron, director of marketing and product management, Cisco Server Networking and Virtualization Business Unit. "Now, with the Isilon solution, we're excited to see the high performance benefits of InfiniBand extended to storage clusters as well. The success Isilon is having with this offering is testament to the compelling customer value proposition of InfiniBand."

"Isilon remains one-hundred percent focused on delivering game-changing business benefits to the enterprise storage market, and we view InfiniBand as a key new technology component of that value proposition for our customers," said Sujal Patel, co-founder and CTO, Isilon Systems. "Cisco continues to push the frontier of networking technologies and is laying the foundation for what we believe is the next generation of enterprise computing and storage infrastructures."

Rapid Adoption Across A Broad Range of Industries

Kelman Technologies, Inc., a leader in the field of seismic data processing for use in oil and gas exploration, has deployed a 25 terabyte Isilon IQ 1920i cluster coupled with a 24-port Cisco SFS 7000 InfiniBand server switch at its Houston location to meet the massive data throughput requirements of its Intel Linux compute cluster. With the increased demand for seismic analysis worldwide, Kelman Technologies' use of Isilon IQ clustered storage with InfiniBand allowed the company to achieve extremely high concurrent data

throughput -- in excess of 50-60 megabytes per second per node -- while removing the burdensome management constraints of traditional SAN and NAS storage systems. The speed of Isilon enabled Kelman to implement its leading-edge depth migration technology and win contracts with large international oil and gas companies that were previously unobtainable.

"We rigorously evaluated a number of storage solutions and ultimately chose Isilon's because it offered the most compelling combination of high performance and low capital acquisition costs," said Pat McKenny, V.P. Geophysical Processing, Kelman Technologies. "The performance and manageability of the Isilon IQ storage system has provided Kelman with a huge business advantage."

DNA Productions, Inc., a full-service animation company, has deployed a 42-node, 80 terabyte Isilon IQ 1920i cluster coupled with a 96-port Cisco SFS 7008 InfiniBand server switch, enabling them to harness the sheer computing muscle of their render farm to achieve new levels of realism in animation. With its networked storage interconnected through low-latency, high- performance Cisco SFS InfiniBand switches, Isilon IQ provides more than 3 Gigabytes per second of total aggregate throughput running through one single pool of shared digital content, dramatically accelerating DNA's mission-critical film production.

"We looked at nearly every SAN and NAS storage system and ultimately selected Isilon IQ because it allows us to deliver better results, more quickly and with less cost," said Brian Chacon and Rich Hirneise, DNA Productions. "Isilon's use of InfiniBand enables us to achieve the high- performance and concurrent throughput requirements we need to create the incredible animation audiences have come to expect."

The Incorporated Research Institution for Seismology (IRIS) is deploying a 60 terabyte Isilon IQ 6000 cluster coupled with a 24 port Cisco SFS 7000 InfiniBand switch to accelerate the collection and delivery of its exploding seismology data stores. IRIS, supported by the National Science Foundation, other federal agencies, universities and private foundations, is dedicated to exploring the Earth's interior through the collection and distribution of seismographic data. In deploying an InfiniBand-enabled Isilon IQ clustered storage system, IRIS creates one single pool of seismology data, allowing them to move to an all digital workflow and leave behind the scalability restraints and resources requirements associated with their former tape-based storage system.

The Isilon IQ 1920, 3000, 4800 and 6000 products are the first clustered storage systems to implement low-latency, high-performance InfiniBand architecture-based communications for intracluster connectivity without requiring any changes to a customer's existing IT environment. By combining Cisco SFS InfiniBand switches with the power of Isilon's OneFS distributed file system and enterprise-class, industry-standard

hardware, Isilon IQ creates a complete, cost-effective storage solution that delivers high-performance and highly concurrent throughput more than 10 times that of traditional storage systems. The Isilon IQ clustered storage system offers the InfiniBand switch option at the same price as its GigE-enabled products, making InfiniBand a compelling solution for its performance, linear scalability and cost advantages.

Cisco SFS InfiniBand server switches combine InfiniBand's unprecedented bandwidth and latency characteristics with enterprise-class high availability, delivering interconnects for server clusters that provide high performance at an affordable cost. The Cisco SFS 7000 series InfiniBand server switches provide 24 to 96 ports of 10 Gigabytes per second server connectivity in platforms designed for best-in-class reliability, high availability and serviceability -- delivering the highest-bandwidth, lowest-latency clustered storage interconnect available in the marketplace.

###

Visual Numerics Participating in SC|05

Monday, Nov 14 @ 22:37 PST

Visual Numerics, celebrating 35 years of producing leading numerical analysis and visualization software, today announced its participation at the annual SC|05 supercomputing conference in Seattle. This year, the company is co-chairing the SC|05 High Performance Computing (HPC) Analytics Committee, demonstrating its high performance computing-compatible IMSL and Java-based JMSL Libraries, and presenting two topics related to analytics and high performance computing.

SC|05 HPC Analytics Committee

As the first software vendor to develop analytical applications in Java for 64-bit platforms, Visual Numerics is on the forefront of development environments for grid computing, parallel computing and supercomputing. As a chair of the SC|05 HPC Analytics Committee, the company helped to introduce the inaugural HPC Analytics Challenge -- a brand new forum showcasing computational intensive applications that solve complex, real-world problems through the use of rigorous and sophisticated methods of data analysis and high-end visualization. For more information about the Challenge, which debuted this year, see its Web site.

Visual Numerics Presentations at SC|05

Visual Numerics will demonstrate its high performance computing-compatible IMSL and Java-based JMSL Libraries in booth #2006, and discuss the application of these technologies in industries ranging from banking to education to scientific research.

In addition, Visual Numerics is offering two data-intensive presentations at the conference:

-- "High Performance Numerical Analysis in Computational Biology Using the IMSL Numerical Libraries" given by Fred Doolittle, IMSL Product Manager, on Tuesday, November 15 from 10:30-11:00 a.m. in Room 615-616. This presentation will provide examples of a variety of computational environments and solutions that have been developed in this field using the IMSL Numerical Libraries.

-- "The IMSL and JMSL Libraries in Action" given by Tim Leite, Director of Educational Programs, on Wednesday, November 16 from 1:30-2:00 p.m. at the SC|05 Education Booth. Leite will demonstrate the IMSL and JMSL Libraries and will also discuss Visual Numerics' Education Programs (Student Licensing, Knowledge in Motion, etc...).

"Our commitment to building advanced analytical development environments for supercomputing began more than 30 years ago," said Phil Fraher, president and CEO of Visual Numerics. "As high performance computing becomes more accessible to a wider range of companies, Visual Numerics will continue to provide numerical analysis solutions that are compatible with higher processing power. Our goal is to help students, scientists and researchers use the most up-to-date data analysis applications today to solve the problems of tomorrow."

###

ClearSpeed Technology Announces Record Performance Per Watt

Monday, Nov 14 @ 12:22 PST

ClearSpeed Technology, a developer of high-performance, low-power, programmable co-processor solutions, today announced its first LINPACK benchmark score for its Advance accelerator board to coincide with the Supercomputing 2005 (SC05) conference in Seattle. The score of 30.2 Gigaflops (GFLOPS) was achieved in a mainstream dual processor workstation with a single ClearSpeed Advance accelerator board. The Advance board operates at a very low power dissipation of under 25 watts, weighs 9 oz. and achieved the Linpack score without program changes. The score appears at the Top500 Web site, in the "Performance of Various Computers Using Standard Linear Equations Software" report maintained by Professor Jack Dongarra, University of Tennessee.

By itself, the workstation is capable of achieving a LINPACK score of 8 GFLOPS and has a power consumption of over 300 watts. Adding a single 25 watt Advance board yields an almost 4x improvement in performance, increasing the performance per watt ratio for the system considerably. In trials, the Advance board has also been shown to lower overall system power consumption at the same time as significantly increasing performance. As the dense linear algebra workload is undertaken on the power efficient accelerator board, the host CPU uses much less power, resulting in a reduction in the total system power consumption

"The Advance board's performance results on the LINPACK benchmark represent ClearSpeed's potential in revolutionizing the HPC industry," said John Gustafson, ClearSpeed CTO of HPC. "Our performance per watt advantage allows us to exceed previous performance levels while addressing the problems of power consumption and heat dissipation."

LINPACK is a collection of Fortran routines that analyze systems of linear equations such as those that arise in physics simulations and financial models. One routine in the collection has become the standard benchmark for measuring the performance of scientific computers and is most notably used by the Top 500 organization in ranking the world's 500 most powerful computer systems. While it is common to optimize the benchmark program for specific architectures, ClearSpeed achieved the 30.2 GFLOP score without changing a single line of the standard benchmark code.

"ClearSpeed appears to be a very interesting technology for high performance with very low power consumption," said Jack Dongarra, maintainer of the LINPACK benchmark.

ClearSpeed's announcement of 30.2 GFLOPS sustained on the LINPACK benchmark coincides with its presence at SC05, the premier international conference on high performance computing. In booth number 802 ClearSpeed will demonstrate its

workstation LINPACK performance and power consumption. Additional demonstrations will show the Advance board in a variety of manufacturers systems, including those from Compusys, IBM and Sun Microsystems. It will also be present in a number of other booths, including those of AMD, Cray, IBM, Intel, NCSA and Sun Microsystems.

###

SuperComputingOnline.com

ASPEED Software Demos at SC05

Monday, Nov 14 @ 12:19 PST

ASPEED Software will demo its ACCELLERANT 3.3 in the HP booth at the high-performance computing industry's premier event - SC05 - this week in Seattle. ACCELLERANT 3.3 is the first solution that enables high-performance computing applications to quickly exploit new multi-core, grid and cluster-based platforms - without re-engineering or reworking the underlying algorithms or logic that drive key business decisions. It provides ongoing real-time dynamic optimization to speed the time to decisions - without compromising the integrity of the application.

ASPEED is being adopted by application owners in banking, finance, pharmaceuticals and energy in their efforts to reduce both run times and the cost of faster response. Its customers have time-critical applications that need to run more frequently and include many more variables; achieving the desired responsiveness often forces them to re-write or re-engineer these applications for distribution, diverting their focus and putting significant pressure on their staffing budget.

ASPEED is an HP Partner providing HPC software development and run-time software. Attendees can experience the demo in HP's booth at the following dates and times:

-- Mon., Nov. 14 7 p.m. - 9 p.m. PT

-- Tues., Nov. 15 2 p.m. - 6 p.m. PT

-- Wed., Nov. 16 10 a.m. - 2 p.m. PT

-- Thurs., Nov. 17 1 p.m. - 4 p.m. PT

"Until now, the focus on performance has centered around faster and cheaper hardware and networks, but organizations are quickly learning that an application-centric approach has higher potential for improving performance at a significantly better cost," said Kurt Ziegler, executive vice president of development for ASPEED. "That's where ASPEED comes in - our ACCELLERANT software dynamically optimizes the way applications use resources to more effectively exploit existing and new hardware as well as cluster and grid fabric, giving them more responsive applications with full application integrity and significantly improved availability."

SuperComputingOnline.com

Mintera Providing 40 Gigabit per Second Optical Transport for SC05

Monday, Nov 14 @ 12:17 PST

Mintera Corporation, the technology leader enabling migration of optical transport to 40 Gigabits per second (Gbps) in metro-core, regional, and ultra long-haul networks was selected to provide 40 Gbps transport solutions for SCinet at the Supercomputing 2005 conference (SC05) sponsored by ACM and IEEE Computer Society.

SCinet is a collection of high performance networks built to support the annual International Conference for High Performance Computing and Communications. SC05 will be conducted in Seattle, WA, from November 12-18, 2005; Bill Gates, Microsoft Chairman and Chief Software Architect, will deliver the keynote address.

SCinet chair Barry Hess, Deputy CIO of Sandia National Laboratories stated, "Each year we sponsor a High-Performance Bandwidth Challenge supporting innovative and bandwidth-intensive application demonstrations that stress the SCinet network infrastructure and capacity while delivering innovative application value on OC-192 or higher links. Hess added, "Mintera's 40 Gbps system will help our infrastructure capacity exceed 500Gbps this year."

SCinet is providing connectivity to many national and worldwide networks through peering relationships with different service providers. Mintera's system, is being used to multiplex four OC-192's up to 40 Gbps for transport across an existing optical line system infrastructure which was originally designed for 10 Gbps transport thus providing a significant band-width boost.

"We are excited to support the SC05 Conference, working with SCinet to show the viability of Mintera's 40 Gbps technology over existing line systems", said Mintera's CEO, Terry Unter. "Mintera's 40 Gbps products offer unique adaptive technology, which allows straight forward use on a variety of existing infrastructures with little or no re-engineering. Using our product, carriers can install 40 Gbps wavelengths seamlessly, and without traffic interruption on their deployed infrastructure, helping to maximize their return on investments by providing a very significant and cost-effective service and capacity upgrade of the existing infrastructure".

SuperComputingOnline.com

Mintera Providing 40 Gigabit per Second Optical Transport for SC05

Monday, Nov 14 @ 12:17 PST

Mintera Corporation, the technology leader enabling migration of optical transport to 40 Gigabits per second (Gbps) in metro-core, regional, and ultra long-haul networks was selected to provide 40 Gbps transport solutions for SCinet at the Supercomputing 2005 conference (SC05) sponsored by ACM and IEEE Computer Society.

SCinet is a collection of high performance networks built to support the annual International Conference for High Performance Computing and Communications. SC05 will be conducted in Seattle, WA, from November 12-18, 2005; Bill Gates, Microsoft Chairman and Chief Software Architect, will deliver the keynote address.

SCinet chair Barry Hess, Deputy CIO of Sandia National Laboratories stated, "Each year we sponsor a High-Performance Bandwidth Challenge supporting innovative and bandwidth-intensive application demonstrations that stress the SCinet network infrastructure and capacity while delivering innovative application value on OC-192 or higher links. Hess added, "Mintera's 40 Gbps system will help our infrastructure capacity exceed 500Gbps this year."

SCinet is providing connectivity to many national and worldwide networks through peering relationships with different service providers. Mintera's system, is being used to multiplex four OC-192's up to 40 Gbps for transport across an existing optical line system infrastructure which was originally designed for 10 Gbps transport thus providing a significant band-width boost.

"We are excited to support the SC05 Conference, working with SCinet to show the viability of Mintera's 40 Gbps technology over existing line systems", said Mintera's CEO, Terry Unter. "Mintera's 40 Gbps products offer unique adaptive technology, which allows straight forward use on a variety of existing infrastructures with little or no re-engineering. Using our product, carriers can install 40 Gbps wavelengths seamlessly, and without traffic interruption on their deployed infrastructure, helping to maximize their return on investments by providing a very significant and cost-effective service and capacity upgrade of the existing infrastructure".

SuperComputingOnline.com

AIC Showcases the World's First 1.5U Form-Factor Rack-Mount Server Enclosure

Monday, Nov 14 @ 12:14 PST

Advanced Industrial Computer, Inc. (AIC) will be showcasing the world's first 1.5U rack-mount server, RMC15, at SC05, slated for November 14-17, 2005, at Washington State Convention and Trade Center, Seattle, Wash. As the leading rack-mount server and storage enclosure designer and manufacturer, AIC continues to provide cutting-edge enclosure solutions by unveiling the world's first 1.5U rack-mountable server enclosure, RMC15, to the industry. Owing to the ever-increasing CPU speed and consequential thermal issue faced by 1U form-factor, RMC15 is the perfect solution -- a ground-breaking invention with superior cooling strategy, more expansions, lower vibrations and acoustic noise over 1U form-factor. Research shows that RMC15 provides 31% more airflow, which results in 23% lower CPU temperature compared to standard 1U form-factor server chassis with same MB and CPU.

With the storage capacity of 2U form-factor, RMC15 is proven to be space and cost effective. By having 0.5U less space, 7 more rack-mount servers can be fit into a 42U rack cabinet. The deployment of RMC15 enables more servers to fit into the rack cabinet and relatively higher I/O performance per rack, which increases the efficiency of hardware management while lowering the overall maintenance cost.

"AIC is constantly committed to delivering innovative solutions since the beginning. While everyone is struggling on how to cool 1U systems, AIC is working on two separate solutions for high performance computing applications. The most obvious approach is to increase the heat-exchange surface area and at the same time increase the airflow over the surfaces. To improve even more, we need to lower the air intake temperature (T-ambient) and ensure hot exhaust does not get re-circulated within the chassis. A 1.5U is just the right answer to all of the above," said Gene Lee, Vice President of Advanced Industrial Computer, Inc. (AIC).

"1.5U form factor is indeed a ground-breaking invention providing advantages from both 1U and 2U form factors. The PCI expansion, dedicated cooling strategy and space-efficiency make it a cost-effective solution to the market today," said Adam Chou, General Manager of Iwill USA, Inc. "As the leading motherboard solution provider, Iwill USA will continue to partner with AIC to accelerate the trend of providing innovative solutions."

Calient Networks Teams with MCNC and LSU CCT to Advance Scientific Research

Monday, Nov 14 @ 11:57 PST

Calient Networks, MCNC, the Louisiana Optical Network Initiative, and the Center for Computation & Technology at LSU announced today at Supercomputing 2005 a partnership to drive optical networking and grid computing towards a new realm of advanced scientific applications. Calient, a leading carrier-class photonic switching systems and software provider, and MCNC, a network research facility and nonprofit provider of network services to North Carolina universities and state government, have partnered to integrate optical control plane and Grid computing technologies for research purposes. Working with the Center for Computation & Technology (CCT) at LSU in conjunction with the Louisiana Optical Network Initiative (LONI), MCNC and Calient are providing new optical network capabilities to drive extreme-scale scientific applications. At SC 05, CCT and LONI will demonstrate how the control of high-speed networks can be coupled to computing resources, data storage archives, and visualization services for enhancing the study of highly complex and computation intensive scientific problems such as the modeling of black holes and other sources of gravitational waves.

"Network infrastructures are no longer commodity 'plumbing' tools but integral components for 'virtual' research teams that rely on resources, both human and physical, residing in various geographic locations," said John Crites, president and CEO of MCNC. "The combined technological expertise between MCNC and Calient together with other research partners will increase innovation from National Research & Education Networks (NRENs) and foster more scientific collaboration."

"Calient was an early pioneer in delivering optical switching and GMPLS technologies, and has developed great expertise in NRENs demonstrated by their successful involvement in the SuperSINET and JGN-2 programs in Japan," said Mark Johnson, chief technology officer at MCNC. "We are pleased to partner with them, having already seen positive results."

"MCNC operates the nationally-recognized North Carolina Research and Education Network (NCREN) and has a world-class team with a history of pioneering advanced network research," said Charles Corbalis, CEO and president at Calient. "Not only are they leading the Experiment Support Services effort for the National LambdaRail but they are also involved in many other advanced network projects. We are committed to helping them achieve their goals of supporting next-generation network initiatives."

Calient's DiamondWave(R) PXC optical switch has been deployed at MCNC's site in Raleigh, and will be the base of a new optical exchange to spur research activity throughout North Carolina universities and facilitate national and international research.

A working relationship already existed between MCNC and LONI based on shared experiences in operating regional networks dedicated to research and education. “We recognized the value of bringing Calient switches to both LONI and NCREN to enable better provisioning of our optical network resources,” said Mark Johnson. LONI has already achieved early success with integrating the Calient DiamondWave PXC product deployed in Baton Rouge with Grid middleware developed by the CCT. “The ability to dynamically allocate bandwidth, coordinate network intelligence, and vertically integrate Grid computing with network services is a linchpin for worldwide NREN collaboration,” added Gigi Karmous-Edwards, principal scientist at MCNC and also the chair of the Control Plane and Grid Integration Middleware Working Group at the Global Lambda Integrated Facility (GLIF).

“Since our research required real-time cooperation and dynamic connectivity between supercomputer sites in Louisiana and with European collaborators we needed new transport protocols and dynamic optical network configurations,” said Ed Seidel, CCT director and research scientist. With the network and attached resources, researchers can also share huge data sets of information around the world instantly, such as complex models of storm surge or path predictions. “More accurate projections can now be realized,” said Seidel. “Network provisioning allows for faster, complex real-time problem resolution.”

“Despite the impact of recent hurricanes in Louisiana, together with Calient and MCNC resources and researchers, we were able to showcase this new network capability in a relatively short time period. These recent disasters proved the value of collaborative support,” said Charlie McMahon, director of LSU Office of Telecommunication. As soon as communications were re-established after the storms, MCNC and Calient personnel were in touch with LSU CCT and doubled their efforts to help LONI get ready for networked computer simulation demonstrations at the October iGRID 2005 and SC 05.

Calient products and solutions can be seen at Supercomputing 2005 in Seattle, November 14-17 at booth 657.

###

Level 5 Networks Works with Cisco to Offer Ethernet Interconnect Option

Monday, Nov 14 @ 11:50 PST

Level 5 Networks, a leading developer of high-speed server interconnect solutions, today announced that it has completed testing of a solution that accelerates applications running on high-performance computing (HPC) clusters using Cisco Gigabit Ethernet switches and Level 5 Networks Ethernet interconnect products. The tests confirmed that Level 5 Networks Ethernet interconnect cards met interoperability criteria with Cisco Catalyst 6500 Series switches, the company's flagship LAN switching platform. The two companies will demonstrate the performance of the combined solution at the Super Computing 2005 (SC05) show in Seattle, Washington from November 14 through 17 in Level 5 Networks booth, #2213.

Additionally, Cisco will be showcasing Level 5 Networks high-performance Ethernet interconnect technology available in its High Performance Computing Lab at its corporate headquarters in San Jose to demonstrate the accelerated application performance speed available when Cisco and Level 5 combine forces in an HPC. The demo will be running an approximate 100-node HPC cluster, and customers will be able to see the immediate performance gains when applications run on the combined power of EtherFabric interface cards and Cisco switches.

"Our customers are demanding low latency Ethernet switching for their data centers and High Performance Cluster environments," said Walt Blomquist, director of product management in the Internet Systems Business Unit at Cisco. "Combining Level 5 Networks EtherFabric technology as the server's network interface with a Cisco switched Ethernet infrastructure furthers our commitment to providing customers with intelligent networking solutions. This helps our customer to run their HPC applications significantly faster than is possible with conventional Ethernet interconnects."

Demand for high-performance computing environments is rising dramatically. Clusters provide a scalable, modular computing environment that can easily be upgraded as server technologies improve while increasing the overall availability of the computing infrastructure. As HPC systems move to a clustered environment, the need grows for high-bandwidth, low-latency server interconnections. Until now, interconnects have been expensive, required significant expertise to implement, and could not scale easily to the high port densities needed today. Because EtherFabric is standards based, it dramatically accelerates application performance without changing the existing IT infrastructure, giving customers significant gains in application performance without requiring changes to the network or the applications.

Companies that choose the Ethernet HPC interconnect based on technology from Level 5

Networks and Cisco Systems benefit from:

* Combined state of the art Ethernet switching and server technologies that dramatically accelerates applications on HPC clusters

* Low-risk solutions that are standards compliant and completely compatible with existing network equipment and application software

* Taking advantage of existing Ethernet networking expertise and management tools

- Accelerating the performance of the applications running within their HPC environment without drastic changes to the servers, applications or to the network

“Level 5 Networks has started a movement to accelerate application performance, liberate server compute power and elevate business productivity by optimizing IT infrastructure,” said Dan Karr, chief executive officer of Level 5 Networks. “The combined power of our high performance Ethernet interconnects and Cisco Ethernet switches makes our vision a reality for HPC customers, and we look forward to expanding our work with Cisco to continue bringing best-in-class solutions to the market.”

###

Level 5 Networks to Accelerate High Performance Computing with Sun Microsystems

Monday, Nov 14 @ 11:46 PST

Level 5 Networks, a leading developer of high-speed Ethernet server interconnect solutions, today announced that Sun Microsystems will make Level 5 Networks EtherFabric cards available to their customers through their Customer Ready Systems (CRS) program. Level 5 Networks also announced support for the Sun Solaris 10 Operating System (OS) and is demonstrating EtherFabric for the Solaris OS running on an x86 platform at the SCO5 trade show, in booth #2213.

Level 5 Networks previewed their high performance Ethernet-base solution to Sun customers at the recent Sun HPC Consortium. Now, Sun customers can immediately increase the performance of their Sun servers and accelerate the applications running on them without the risk of implementing non-standard niche interconnect solutions.

"Sun and Level 5 Networks share a common commitment to deliver high performance network computing solutions that reduce cost and complexity for our customers," said Bjorn Andersson, director, HPC and Grid Computing, Sun Microsystems, Inc. "Coupled with Sun's high performance servers, Level 5 Networks EtherFabric cards can deliver significant performance gains while preserving our customers' investment in Ethernet infrastructure."

The decision to make EtherFabric available directly via Sun's Customer Ready Systems program addresses the market need and subsequent growth of high performance computing. HPC clusters process large amounts of data and consume server resources to solve complex, numerically intensive problems. Traditionally the domain of engineers, physicists, astronomers and other scientists solving grand challenge problems, recent trends in technology and in the enterprise are increasing the scope of application of HPC tremendously. By supporting EtherFabric in the Sun-based HPC clusters, Sun customers are better able to utilize their existing computer resources, resulting in increased productivity and improved return on their investment in IT infrastructure.

"Level 5 is driving the future of HPC by delivering dramatic application performance while maintaining full compatibility with existing IT infrastructures," said Dan Karr, chief executive officer of Level 5 Networks. "This initiative with Sun Microsystems makes our EtherFabric solution available to Sun users and immediately brings the full potential of high performance computing to Sun data centers."

EtherFabric accelerates the performance of networked applications on HPC clusters using standard Ethernet and TCP/IP networking protocols. The standards-based EtherFabric solution is compatible with existing applications and increases application performance

and computational throughput without expensive new hardware and risky software changes. Other non-standard high performance networking solutions, by contrast, require changes to both hardware and software. EtherFabric enhances data center performance and lets Sun customers optimize their existing IT infrastructure to maximize the investment in their high-performance servers that support the computing-intensive demands in industries such as life sciences, automotive, financial services and others.

Advantages of deploying Sun servers with EtherFabric:

* Enables applications to fully leverage multi-CPU/core Sun servers

* Accelerates the performance and improves efficiency of applications running on Sun Server clusters and storage systems

* Increases productivity and cuts time to market

* Provides a total solution to fully comply with industry standards; is a simple upgrade to existing IT operations

* Optimizes investments in Sun infrastructure, lowers the total cost of ownership and improves capital lifecycle

"This agreement is a great opportunity for Level 5 Networks," said Anne MacFarland, director of infrastructure architectures and solutions at The Clipper Group, a New England-based IT analysis firm. "Level 5 Networks offers a way to use existing Ethernet connectivity, but alleviates the high CPU loading that can constrain the performance of many traditional Ethernet systems. Level 5 brings more performance to compute-intensive applications, and more scalability to the server clusters on which such applications run."

###

SuperComputingOnline.com

Juniper Networks Routing Platforms Enable Core of SCinet at SC05

Monday, Nov 14 @ 11:45 PST

For the fifth straight year, Juniper Networks T-series routing platforms will participate in the core of SCinet, a collection of high-speed networks built to support the annual International Conference for High Performance Computing, Networking and Storage (SC05). SCinet provides the infrastructure for both an extremely advanced high-capacity experimental network, and the high-performance, production-quality network supporting exhibitor's networking needs.

As part of the core of SCinet, Juniper Networks T-series routing platforms will help support all SC05 network traffic, including the supercomputer applications being run by researchers participating in the conference.

The T-series delivers the unique combination of advanced routing features, best-in-class IP/MPLS capabilities, and unmatched reliability and scale that makes it the ideal platform for both production networks and advanced research and experimental networks. Shipping since April 2002, over 1,000 T-series units are deployed in networks worldwide. With the multi-chassis TX Matrix technology, the T-series provides a seamless migration path to multi-terabit speeds.

For more information, please visit the Juniper Networks booth (#530) at SC05. In the booth, Juniper Networks will also be highlighting a range of core routing, security and application acceleration solutions.

DOE/LLNL BlueGene/L and IBM gain Top Positions in the TOP500

Sunday, Nov 13 @ 21:29 PST

In what has become a closely watched event in the world of high-performance computing, the 26th edition of the TOP500 list of the world's fastest supercomputers was released today at the Supercomputing Conference (SC05) in Seattle, WA. The new TOP500 list, as well as the previous 25 lists, can be found on the Web at its Web site. The No. 1 position was again claimed by the BlueGene/L System, a joint development of IBM and DOE's National Nuclear Security Administration (NNSA) and installed at DOE's Lawrence Livermore National Laboratory in Livermore, Calif. BlueGene/L also occupied the No. 1 position on the last two TOP500 lists. However, the system was doubled in size during the last six months and reached a new record Linpack benchmark performance of 280.6 TFlop/s ("teraflops" or trillions of calculations per second). No other system has yet exceeded the level of 100 TFlop/s and this system is expected to remain the No. 1 Supercomputer in the world for the next few editions of the TOP500 list.

The pace of innovation and performance improvements seen at the very high end of scientific computing shows no sign of slowing down. This time, four of the TOP10 systems on the June 2005 TOP500 list were displaced by newly installed systems, and the last 221 systems on the list from June 2005 are now too small to be included.

The new No.3 system, also installed at LLNL, is the ASCI Purple system, built by IBM and based on their p575 server. It reached 63.4 TFlop/s.

Two new systems at DOE's Sandia's National Laboratories captured the No. 5 and 6 spots, with a Dell PowerEdge-based system slightly outperforming Cray's Red Storm system. The NEC-built Earth Simulator, which has a Linpack benchmark performance of 35.86 TFlop/s and had held the No. 1 position for five consecutive TOP500 lists before being replaced by BlueGene/L last November, has slipped within one year to No. 7.

The No. 10 spot was captured by a Cray's XT3 system at DOE's Oak Ridge National Laboratory with 20.53 Tflop/s. This is also the new entry level for the TOP10, up from just under 10 TFlop/s Linpack performance one year ago.

IBM has now established itself as the dominant vendor of supercomputers with almost half of the list (43.8 percent) carrying its label. Also, five of the TOP10 systems are from IBM. Hewlett-Packard (HP) remains unchallenged at the second position in this survey with 33.8 percent of all systems.

Intel microprocessors are at the heart of two-thirds (333) of all 500 systems. Intel's new EM64T-based processors are very successful in the high performance computing (HPC)

market place, with 81 systems using them already. AMD's Opteron processors are also steadily gaining ground, now with 55 systems using them compared to only 25 systems six months ago.

The U.S. is clearly the leading consumer of HPC systems with 305 of the 500 systems installed there, while the European (100 systems) and Asian share (66 systems) is slowly decreasing over time.

Here are some highlights from the newest Top 500:
Only systems exceeding the 1.64 TFlop/s mark on the Linpack benchmark were qualified to make the list this time, compared to 850.6 GFlop/s one year ago. The last system on the latest list would have been listed at position 173 just one year ago.

The entry level for the TOP10 exceeds 20 TFlop/s and the entry point for the top 100 moved from 2.026 TFlop/s one year ago to 3.98 TFlop/s. Total combined performance of all 500 systems on the list is now 2.30 PFlop/s ("petaflops" or thousand "teraflops"), compared to 1.127 PFlop/s one year ago.

Other trends of interest:
A total of 333 systems are now using Intel processors, with 81 one these are already using the new EM64T processors. The second most-commonly used processors are the IBM Power processors (73 systems), ahead of AMD Opteron processors (55).

There are 360 systems now labeled as clusters, making this the most common architecture in the TOP500. Of these, 249 cluster systems are connected using Gigabit Ethernet and 70 system using Myricom's Myrinet.

At present, IBM and Hewlett-Packard sell the bulk of systems at all performance levels of the TOP500. IBM remains the clear leader in the TOP500 list with 43.8 percent of systems and 52.8 percent of installed performance. HP is second with 33.8 percent of systems and 18.8 percent of performance. No other manufacturer is able to capture more than 7 percent in any category.

The U.S. is clearly the leading consumer of HPC systems with 305 of the 500 systems installed there (up from 267 one year ago). The European (100 systems) and Asian share (66 systems) is slowly decreasing.

Dominant countries in Asia are Japan with 21 systems (down from 23 six months ago) and China with 17 systems (down from 19 six months ago).

In Europe, Germany (24 systems) lost the No.1 spot again to UK again (41 systems). Six months ago Germany was in the lead with 40, compared to UK's 32 systems.

The TOP500 list is compiled by Hans Meuer of the University of Mannheim, Germany; Jack Dongarra of the University of Tennessee, Knoxville; and Erich Strohmaier and Horst Simon of NERSC/Lawrence Berkeley National Laboratory.

###

OSC's Blue Collar Computing Concept to Spark Imaginations at SC|05

Friday, Nov 11 @ 12:51 PST

The Ohio Supercomputer Center (OSC) will showcase its Blue Collar Computing and Third Frontier Network (TFN) initiatives by highlighting these projects and more at Supercomputing 2005 (SC|05). A week-long international supercomputing conference, SC|05 is the world's leading event for high performance computing and research, and is a hotbed of collaboration opportunities. The conference will be held at the Washington State Convention and Trade Center in Seattle, Washington, next week.

OSC's showcase will feature staff and guest presentations and demonstrations, including the popular MATLAB mini-symposium. Leading experts from Procter & Gamble, the Council on Competitiveness, MathWorks, OpenFPGA, GridChem and computational science initiatives will speak on groundbreaking topics affecting the quickly evolving world of high performance computing (HPC) .

"In seven years, OSC has cultivated a strong exhibiting presence in the world's premiere supercomputing conference," said Stan Ahalt, OSC Executive Director. "We take this opportunity to showcase the programs that have established OSC as a national leader in high performance computing, networking, and research."

During the conference, OSC will focus on both Blue Collar Computing and the TFN. Blue Collar Computing – "HPC for the Rest of Us" – is a national initiative to make HPC easily accessible across the full spectrum of industry and research by rethinking HPC software, hardware and training. OSC's Computational Science Initiative is a related program that creates dynamic partnerships among K-20 educational institutions to train a diverse population of students for the leading-edge technical careers of the future.

TFN, the nation's most advanced statewide fiber-optic network designed to enhance education, research and economic development regionally, will feature a live demonstration of its collaborative measurement and testing research projects. The Live Performance Monitoring of TFN-Abilene-ADECnet Network Paths demonstration features TFN (OARnet Network Backbone), Abilene (Internet2 Network Backbone) and ADECnet (American Distance Education Consortium Network Backbone) network paths using the "ActiveMon." ActiveMon is an open-source software developed at OARnet, OSC's networking division.

Other demonstrations include:

* iWarp , a remote direct memory access (RDMA) technique to create high-end custom network fabrics to increase performance;
* A face recognition application using parallel MATLAB matching live video camera

images of SC|05 attendees with a database of more than a thousand faces;
* A StorCloud demonstration from The Ohio State University's Department of Biomedical Informatics featuring oil reservoir management and biomedical imaging.

Another highlight is OSC-Springfield, a Department of Energy (DOE) center focused on data-intensive computing issues such as bioinformatics, computational biology, data warehouse applications, and earth science and materials modeling. An interactive touch monitor addresses OSC's outreach programs, bioinformatics and biomedical groups, homeland security applications, and training areas.

###

Open Science Grid Technology and Applications Featured at SC|05

Friday, Nov 11 @ 07:23 PST

The Open Science Grid, a nationwide community grid built by research groups from United States universities and national laboratories, will showcase advanced grid technologies and innovative scientific applications at SC|05, the premier international conference on high performance computing, networking and storage.

Presentations and demonstrations at 13 booths will show how scientists from diverse fields contribute manpower and resources to the OSG and benefit from easy access to local and remote resources, testing and production environments for middleware and applications, and a common computing infrastructure. Over 20 member organizations representing more than 50 institutions and hundreds of researchers contribute to the OSG and benefit from access to shared resources worldwide, including over 10,000 CPUs and many terabytes of data storage.

The U.S. Department of Energy's Fermi National Accelerator Laboratory and Stanford Linear Accelerator Center exhibit, Booth 302, is dedicated to the OSG as well as the high-performance networking, computing and scientific advances from the two national laboratories. OSG-related demonstrations and presentations at SC|05 will include:

a.. Astronomy and astrophysics: Sloan Digital Sky Survey astronomers are running several applications on the OSG, including a search for near Earth objects and the processing of tens of thousands of spectra to determine the properties of quasars. LIGO physicists use the OSG to search for gravitational waves from the binary inspiral of neutron star systems.
b.. Physics: The CMS experiment uses the OSG for computing and data-intensive particle physics simulations, and will use the grid for data analysis when the Large Hadron Collider begins operating at CERN in Geneva, Switzerland. OSG collaborators from the CMS and ATLAS experiments participated in the LHC Service Challenges, which earlier this year sustained a continuous data flow of 600 MB/s on average for 10 days from CERN to sites worldwide. Particle and nuclear physics experiments use grid technology to simulate, process and transfer data in their search for the nature of matter and energy.
c.. Biology: OSG resources are used for genome analysis at Argonne National Laboratory, and for bioinformatics and biology through the Grid Resources for Advanced Science and Engineering (GRASE) VO at the State University of New York at Buffalo.
d.. Bandwidth Challenge: OSG collaborators will participate in the SC|05 Bandwidth Challenge, aiming to transfer data internationally at a higher rate than ever before.
e.. Grid technology: OSG collaborators have advanced many areas of grid technology. Demonstrations will highlight research in the areas of security, storage, data distribution, collaboration, networking and monitoring.
f.. Interoperability: OSG collaborators successfully running jobs on multiple grid **Open**

infrastructures, including the OSG, the TeraGrid, the LHC Computing Grid, the ACDC Grid and the Grid Laboratory of Wisconsin, will demonstrate this interoperability at several SC|05 exhibits.
g.. Education: The QuarkNet cosmic ray study uses grid tools and techniques in a model problem-based, student-driven and technology-dependent student investigation. QuarkNet collaborators will demonstrate how high school students reach beyond classroom walls to explore data with other students and experts and share results, publishing original work to a worldwide audience.

The research and accomplishments of the OSG collaboration will be featured at the following SC|05 exhibits: Argonne National Laboratory (Booth 1629); Brookhaven National Laboratory (2238); Caltech, Center for Advanced Computing Research (428); Center for Computational Research, University at Buffalo (459); Fermilab/SLAC (302); Indiana University (202); Lawrence Berkeley National Laboratory (1828); National Center for Supercomputing Applications (1639); Purdue University (206); Texas Tech University, HPCC (2162); University of Florida (6319); University of Iowa (128); and Vanderbilt University (6013).

SC|05 takes place November 12-18 at the Washington State Convention and Trade Center in Seattle.

The Open Science Grid makes innovative science possible by bringing multidisciplinary collaborations together with the latest advances in distributed computing technologies. This shared cyberinfrastructure, built by research groups from U.S. universities and national laboratories, includes over 10,000 CPUs and many terabytes of data storage. The OSG receives support from the National Science Foundation and the U.S. Department of Energy's Office of Science. For more information about the OSG, visit its Web site.

###

Highlight during SC|05: Fraunhofer Visualization on an IBM Cell Cluster

Friday, Nov 11 @ 05:26 PST

One highlight of the International Conference on High Performance Computing SC|05 in Seattle will be the visualization software PV-4D of the Fraunhofer Institute for Industrial Mathematics ITWM. The software has been installed on a prototype of an IBM cell blade cluster and will be presented to a large public for the first time. A successful test run was performed during the Fraunhofer Annual Meeting in Magdeburg, Germany on October 19 this year.

The visualization and analysis of extremely large amounts of data from the fields of scientific computing, image-rendering methods in medical technology, or oil exploration can be represented by easily understandable realtime animations thanks to PV-4D. The cell processor is now rendering things even faster: it is equipped with nine cores and currently represents the chip with the highest performance available for the mass market.

The professional visualization environment has been developed at the Competence Center High Performance Computing of the ITWM by Dr. Carsten Lojewski and his team and has been awarded the Joseph-von-Fraunhofer-Preis 2005. PV-4D is a pure software solution and explicitly works without the support of graphic cards. Time-critical functions use the vector units and the multi-pipes of modern processors for parallelization. The software works especially effectively if it is running on a cluster system. The network support (Gbit-Ethernet, Myrinet, Infiniband) is directly integrated into the parallel software, thus guaranteeing minimum latencies and large bandwidths.

A further central element of PV-4D is the decentralized software control. A server process is responsible for the computation of images and allows for the integration of simple viewers, which are installed on a laptop within the office network or on front-end computers of a virtual reality environment. The image composi-tion, which is usually done by special hardware, is also completely realized within the software.

The software is the tool with the highest performance available worldwide for a fast and interactive representation of gigantic amounts of data. The performance of the current cluster systems even exceeds the strongest special graphics computers, thus al-lowing for the interactive visualization of more than one terabyte of data. This performance is also convincing industrial partners, such as DaimlerChrysler, Shell AG, or the Berkeley Labs in California.

During the Supercomputing in Seattle (November 12-18), the new PV-4D Ray Tracing Kernel will be shown. It allows for the first time volume rendering of complex seismic data in real time. This opens up a new world of interaction for geophysicists, facilitating

the analysis of data for the exploration of new oil wells.

The current implementation on the cell processor exclusively uses the vector units of the individual SPEs for the so-called SIMD-Shaft Ray Tracing; the available PEs are responsible for the parallel or-ganization of data. In such a way, 128 parallel rays can be com-puted at first-order coherence on each cell processor, and at sec-ond-order coherence 32 parallel rays with direct hardware sup-port. At the IBM booth in Seattle, 16 cell processors will be com-bined to form of a high-performance ray-tracing cluster.

###

SuperComputingOnline.com

Panasas to Demonstrate Success in Bridging the Clustered Storage Needs

Thursday, Nov 10 @ 06:44 PST

At this year's international conference on high-performance computing, networking and storage, SC|05, Panasas, Inc. will demonstrate the company's continued success in bridging the clustered storage needs of the scientific and commercial high-performance computing (HPC) markets. Show attendees visiting Panasas at booth #1006 or attending one of several Masterworks presentations will learn about the results achieved with Panasas' clustered storage. Visitors will also receive details about the newest additions to the Panasas ActiveScale Storage Cluster designed to meet the diverse needs of both scientific and commercial HPC customers. The SC|05 show will be held at the Washington State Trade and Convention Center in Seattle, Washington, November 12-18, 2005.

"The Linux cluster revolution is expanding supercomputing beyond the scientific markets into mainstream commercial applications, and with this expansion comes the need to bridge the unique IT requirements of both the scientific and commercial communities," said Victor Perez, president and chief executive officer at Panasas. "Panasas is providing this bridge for high-performance clustered storage through a solution that increases performance while driving out the cost and complexity of upgrading to a next-generation storage infrastructure. This helps customers achieve an immediate return on investment from their Linux clusters."

Panasas Everywhere at SC|05
Linux clusters powered by Panasas are used by customers in a growing number of mainstream applications that touch every aspect of life, from helping accelerate the development of life-saving drugs, to meeting global demands for oil and gas, to enabling the next wave of fully animated feature films. Panasas is setting the industry standard for high-performance clustered storage by helping customers more quickly achieve their objectives and successfully compete in global markets. The power of Panasas will be on display throughout the week at SC|05 at several locations and venues.

Object-Storage Masterworks Presentation -- On Tuesday, November 15, 2005 from 3:30 -4:15 p.m., Garth Gibson, co-founder and CTO for Panasas, will deliver a Masterworks session titled "Object Storage Architecture for Scalable File Systems." Gibson will review the principles of RAID and object storage and discuss how Panasas combines RAID and object storage into a new breed of storage virtualization and clustering, enabling advances in highly scalable, high performance and high availability to meet the scaling demands of cluster applications.

Masterworks Presentations By Panasas Customers -- On Thursday November 17, 2005

from 1:30 - 2:15 p.m., representatives from Walt Disney Feature Animation will discuss their history of using high-end computing technology such as object-storage and the 57th site on the Top500 Supercomputer Sites list(11/2004) in making films. On the same day from 2:15 - 3:00 p.m., Panasas customer LSU will be presenting details of how HPC modeling was used in emergency response to Hurricane Katrina and outline sophisticated approaches under development to forecast such catastrophic events in the future.

Panasas Booth #1006 -- Customers, partners and Panasas executives will deliver sessions nearly every hour. Customer presentations will be delivered by LSU, Los Alamos National Labs and Walt Disney Feature Animation, and partner presentations by AMD, Fluent, Force10 Networks, Intel, and Verari Systems. Show attendees can also learn more about Panasas' next generation of hardware and software to be announced at the show. The new releases will further provide customers with the maximum return on investment through increased Linux cluster performance, increased management efficiency, improved business continuity and seamless integration into heterogeneous server environments.

Object-Storage Tutorial -- On Monday, November 14, 2005 from 1:30 - 5:00 p.m., David Nagle and Brent Welch of Panasas will host a tutorial on object-based cluster storage systems. This tutorial will present the fundamentals of object-based storage including the underlying architectural principals and how various products have adapted those principals into their product designs.

###

SuperComputingOnline.com

SDSC Showcases How Scientists Move From Data to Discovery

Wednesday, Nov 09 @ 12:12 PST

The data-intensive computing driving cutting-edge science will be featured at the San Diego Supercomputer Center (SDSC) exhibit at SC2005, the annual conference of high performance computing November 13-17 at the Washington State Convention and Trade Center in Seattle. Visitors to the SDSC booth #1838 can get a glimpse of how the center's focus on data is enabling discoveries that have the potential to change our world.

Dr. Fran Berman, Director of the Center, will deliver the keynote speech to educators from around the country at the SC05 Education Program. Dr. Berman will focus on information technology as an enabler for modern science and engineering research and education. She will be speaking on Monday, November 14 from 9 - 10:30 a.m. in room 3AB of the conference center.

SDSC's leadership and expertise in data technologies will be detailed in a series of one-hour demonstrations and presentations held Tuesday, November 15 through Thursday, November 17. Demonstrations are scheduled from 10 a.m. to 6 p.m. on Tuesday and Wednesday and from 10 a.m. to 3 p.m. on Thursday.

Presentation highlights include the following:

* Researchers from a PRAGMA project team will demonstrate the analysis of the Avian flu genome using a system developed for distributed proteome annotation. The system utilizes a computational data grid developed by PRAGMA partners in the international testbed.
* Reagan Moore, head of the Knowledge Storage Resource Broker (SRB) Lab, will detail the Terashake Project, a visual simulation showing peak velocities of surface ground motion for earthquake waves of a 7.7 earthquake on the San Andreas fault in Southern California. Two additional demonstrations will detail how this 240 Terabyte seismic visualization was created and the latest findings by the Southern California Earthquake Center (SCEC), who sponsored the original program.
* Jerry Rowley, director of the Advanced Database Projects Lab, will discuss how the center worked with the Red Cross and others to develop an amalgamated survivor lists for victims of Hurricane Katrina
* Greg Quinn, project lead, will introduce the Collaboration Notebook, a new online data management application funded by Microsoft Research.

Other demonstrations and presentations include an overview of the newly launched Synthesis Center and Data Central sites, the Integration of Digital Library Technology with SRB, scalable installation with Rocks software and the TeraGrid Cyberinfrastructure. For more information or a schedule of presentations, visit SDSC booth #1838 or contact Ashley Wood at 760-207-3795.

NCSA's Innovative Systems Research To Be Featured At SC|05

Wednesday, Nov 09 @ 11:39 PST

The National Center for Supercomputing Applications (NCSA) will feature its research into innovative high-performance computing systems in Booth 1639 at SC05 in Seattle Nov. 12-18. Central to NCSA's mission is the push to research and deploy innovative systems that decrease the cost and/or extend the range of computational science and engineering. Tomorrow's advanced computing systems will include processors with multiple processor cores on each chip, reconfigurable computers based on rapidly advancing field-programmable gate arrays (FPGAs), and heterogeneous systems with fast communications fabrics interconnecting various types of processors. These new systems have the potential to dramatically increase the fidelity and range of simulations and the scope and speed of analysis, data mining, and visualization. They also pose significant technical challenges. NCSA Chief Technology Officer Rob Pennington will talk about NCSA's leadership in the instrumentation, measurement, modeling, testing, and evaluation of innovative computing systems—both hardware and software—during a presentation at 1 p.m. on Tuesday, Nov. 15 in the center's booth.

Other presentations detailing the research and deployment activities of NCSA's Innovative Systems Laboratory will include:

* From 8:30 a.m. to 5 p.m. on Sunday, Nov. 13, NCSA principal engineer David Pointer will be one of four presenters at a Reconfigurable Supercomputing tutorial. The tutorial will introduce reconfigurable supercomputing and its advances in systems, programming, applications, and tools. Reconfigurable systems from SRC, Cray, SGI, and Star Bridge as well as COTS-based efforts will be considered. Application developments and performance studies will be presented. A comparative case study will be demonstrated from development to compilation and running across the aforementioned platforms. In addition to Pointer, presenters will be Tarek El-Ghazawi of George Washington University, Duncan Buell from the University of South Carolina, and Kris Gaj of George Mason University.
* At 2 p.m. on Tuesday, Nov. 15 in NCSA's booth, Pointer and NCSA researchers David Raila and Craig Steffen will give a talk on "Reconfigurable Systems Applications Programming," describing how researchers in the center's Innovative Systems Laboratory are working with application scientists to explore the potential of FPGA systems. Insights, lessons learned, and results gained from porting the NAMD, BLAST, and MATPHOT applications to reconfigurable architectures will be presented.
* Pointer, Raila, and Steffen also will be available at NCSA's booth during exhibit hours to discuss the center's reconfigurable systems research and will conduct informal demonstrations using reconfigurable hardware installed in the booth.

For more information on innovative systems research at NCSA, go to its Web site.

NCSA (National Center for Supercomputing Applications) is a unique state-federal partnership to develop and deploy national-scale cyberinfrastructure that advances science and engineering. Located at the University of Illinois at Urbana-Champaign, NCSA is one of the leading National Science Foundation-supported supercomputing centers. Additional support comes from the state of Illinois, the University of Illinois, private sector partners, and other federal agencies. For more information, see its Web site.

###

Video to Premiere before Tuesday Keynote by Microsoft's Bill Gates

Wednesday, Nov 09 @ 11:03 PST

SC|05, the premier international conference on high performance computing, networking, storage and analysis, will make public a special video that explores the vital influence of supercomputing in solving the grand challenges of social and commercial problems affecting the quality of life for people throughout the world. This informative video kicks off SC|05's opening session on Tuesday, November 15, 8:30 a.m., at the Washington State Convention and Trade Center. A grant from ACM and IEEE Computer Society, cosponsors of the SC|05 conference, allows supercomputing centers, research labs and higher education to freely use the video as one of their key educational pieces for years to come. "Many people outside of high performance computing, networking, and storage do not realize the long process that is necessary from initial concept and prototype to when an application impacts our lives down the road," said Bill Kramer, SC|05 General Chair. "Supercomputers play a critical role in jumpstarting many major applications that we have come to rely upon."

The video will explore the influence of supercomputing technology and research applications that ultimately arrive to the general consumer market as low-cost technologies.

"As a key component of information technology, supercomputing offers advances in natural disasters, transportation, environmental monitoring, and health and medicine," said David Patterson, ACM President. "We want this video to be used by educational and research institutions to broaden awareness and understanding of the essential role of supercomputing for computer science professionals, public policy officials and the public that benefits from these exciting advances."

The video tracks how the components of supercomputing - hardware, software, and algorithms - combine to produce world-altering discoveries. Extensive financial, labor, and time investment from government and industry produce supercomputing applications that "trickle-down" to commercial products and services. These applications result in major improvements in the production of automobiles, airplanes, pharmaceuticals, financial markets, and even animated movies.

Examples of initial supercomputing investments resulting in industrial and other major innovations include:

-- Flow analysis programs in the 1970s were created at national laboratories; by the early 1990s, these programs were being used for cell phone plastic molds

-- Vibration analysis (NASTRAN) was developed for government projects and is now

available to all mechanical engineering firms to design airplane wings, auto bodies, and the Tacoma Narrows Bridge

-- A Cray-2 was loaned by the National Aeronautics and Space Administration (NASA) to the National Institutes of Health (NIH) to do the first analysis of X-ray defraction crystallography data of the AIDS Virus

-- By the early 70's, graphics-based finite element analysis software, such as LS-DYNA, was limited to expensive mainframe computers generally owned by the aeronautics, automotive, defense, and nuclear industries. Now, FEA is used for safer and cleaner cars (crash analysis), lighter milk cartons, and running shoe materials

Supercomputing algorithms have resulted in human genome mapping that detects birth defects and Alzheimer's disease. Even the design of sharper and longer-lasting disposable razor blades was achieved by using computational algorithms. With the aid of computer infrastructure models, researchers are able to accurately predict the effects of the recent hurricanes and plan for future natural disasters.

This video demonstrates how supercomputers have closed a gap between scientific theory and its application. Computer science professionals and their use of supercomputers have changed our world for the better.

###

Cross-Continental InfiniBand Cluster Staged for Deployment at SC|05

Wednesday, Nov 09 @ 10:58 PST

The OpenIB Alliance today announced that nearly thirty organizations are confirmed to showcase the world's largest, cross-continental InfiniBand data center in conjunction with SCinet at next week's Supercomputing 2005 conference (SC|05) in Seattle, Washington (Washington State Convention and Trade Center). The SCinet InfiniBand cluster will host over 6TFlops of supercomputing performance, which would rank the cluster as high as 50 on the TOP500 list of the world's most powerful supercomputers. In addition, the SCinet InfiniBand fabric will have direct access to native attached InfiniBand storage solutions hosted by the StorCloud initiative.

The strong industry mix of InfiniBand component vendors, infrastructure equipment suppliers, software vendors, and research and university end-users that are participating include:

AMD, Ames Laboratory, Appro, ASUS, Cisco, Cornell Theory Center, Dell, Emcore, Hewlett Packard, IBM, Intel, IWILL USA Corp., Lawrence Berkeley National Laboratory, Mellanox, Microsoft, NCSA, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Obsidian Research, Pathscale, Pittsburgh Supercomputing Center, Rackable, RedHat, Sandia National Laboratories, Silicon Graphics, SilverStorm, Sun Microsystems, Tyan Computer, and Voltaire.

SC|05 participants with the OpenIB Alliance will use the OpenIB software stack for both Linux and Windows host systems, and the OpenIB subnet manager (OpenSM) will manage the entire Infiniband network. The availability of the OpenIB stack from all vendors is the industry's solution to widespread interoperability and mass deployment of InfiniBand fabrics in data centers for enterprise and performance computing.

The SCinet InfiniBand fabric locally interconnects participating booths on the SC|05 exhibit floor via InfiniBand over MPO multi-mode fiber using Emcore SmartLink modules.

The three remote InfiniBand clusters are at Lawrence Livermore National Laboratory in California, the Intel Dupont facility in Washington, and the Naval Research Laboratory in Virginia. These locations are connected through the wide area network (WAN) point of presence for SCinet, the Pacific Northwest Gigapop at Washington University, and utilize optical long-haul equipment from Ciena, Cisco, and Juniper running over networks provided by Abilene, Esnet, Internet2, National LambdaRail, and Qwest. At each endpoint, InfiniBand over optical (either DWDM or SONET OC192) is converted by Longbow XRs from Obsidian Research. The Longbow enables globally distributed InfiniBand fabrics to seamlessly cross connect by encapsulating 4X InfiniBand over OC-

192c SONET, ATM, or 10GbE WANs at full InfiniBand data rates. The conversion is totally transparent to the InfiniBand fabric and is interoperable with OpenIB's software stack and subnet manager.

The SCinet InfiniBand network will also be directly connected to native InfiniBand storage solutions using standard block-level, file-level, and clustered-file system technologies. Hosted by the StorCloud initiative at SC|05, this cluster of native InfiniBand storage solutions highlights the wide availability and end-user demand for higher performance, lower latency storage solutions.

"The significant vendor participation in the SCinet InfiniBand network is a testament to the remarkable progress by the developers of the OpenIB open-source software stack which was substantially funded by the DOE NNSA ASC PathForward program," said Bill Boas, Vice-Chair of OpenIB. "InfiniBand is the lowest latency, highest bandwidth fabric for data center and performance computing applications today. The multi-vendor interoperability at SC/05 validates its success and verifies the widespread acceptance of OpenIB by customers."

Additional information about the OpenIB network at SC|05 is available at its Web site.

###

Cross-Continental InfiniBand Cluster Staged for Deployment at SC|05

Wednesday, Nov 09 @ 10:58 PST

The OpenIB Alliance today announced that nearly thirty organizations are confirmed to showcase the world's largest, cross-continental InfiniBand data center in conjunction with SCinet at next week's Supercomputing 2005 conference (SC|05) in Seattle, Washington (Washington State Convention and Trade Center). The SCinet InfiniBand cluster will host over 6TFlops of supercomputing performance, which would rank the cluster as high as 50 on the TOP500 list of the world's most powerful supercomputers. In addition, the SCinet InfiniBand fabric will have direct access to native attached InfiniBand storage solutions hosted by the StorCloud initiative.

The strong industry mix of InfiniBand component vendors, infrastructure equipment suppliers, software vendors, and research and university end-users that are participating include:

AMD, Ames Laboratory, Appro, ASUS, Cisco, Cornell Theory Center, Dell, Emcore, Hewlett Packard, IBM, Intel, IWILL USA Corp., Lawrence Berkeley National Laboratory, Mellanox, Microsoft, NCSA, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Obsidian Research, Pathscale, Pittsburgh Supercomputing Center, Rackable, RedHat, Sandia National Laboratories, Silicon Graphics, SilverStorm, Sun Microsystems, Tyan Computer, and Voltaire.

SC|05 participants with the OpenIB Alliance will use the OpenIB software stack for both Linux and Windows host systems, and the OpenIB subnet manager (OpenSM) will manage the entire Infiniband network. The availability of the OpenIB stack from all vendors is the industry's solution to widespread interoperability and mass deployment of InfiniBand fabrics in data centers for enterprise and performance computing.

The SCinet InfiniBand fabric locally interconnects participating booths on the SC|05 exhibit floor via InfiniBand over MPO multi-mode fiber using Emcore SmartLink modules.

The three remote InfiniBand clusters are at Lawrence Livermore National Laboratory in California, the Intel Dupont facility in Washington, and the Naval Research Laboratory in Virginia. These locations are connected through the wide area network (WAN) point of presence for SCinet, the Pacific Northwest Gigapop at Washington University, and utilize optical long-haul equipment from Ciena, Cisco, and Juniper running over networks provided by Abilene, Esnet, Internet2, National LambdaRail, and Qwest. At each endpoint, InfiniBand over optical (either DWDM or SONET OC192) is converted by Longbow XRs from Obsidian Research. The Longbow enables globally distributed InfiniBand fabrics to seamlessly cross connect by encapsulating 4X InfiniBand over OC-

192c SONET, ATM, or 10GbE WANs at full InfiniBand data rates. The conversion is totally transparent to the InfiniBand fabric and is interoperable with OpenIB's software stack and subnet manager.

The SCinet InfiniBand network will also be directly connected to native InfiniBand storage solutions using standard block-level, file-level, and clustered-file system technologies. Hosted by the StorCloud initiative at SC|05, this cluster of native InfiniBand storage solutions highlights the wide availability and end-user demand for higher performance, lower latency storage solutions.

"The significant vendor participation in the SCinet InfiniBand network is a testament to the remarkable progress by the developers of the OpenIB open-source software stack which was substantially funded by the DOE NNSA ASC PathForward program," said Bill Boas, Vice-Chair of OpenIB. "InfiniBand is the lowest latency, highest bandwidth fabric for data center and performance computing applications today. The multi-vendor interoperability at SC/05 validates its success and verifies the widespread acceptance of OpenIB by customers."

Additional information about the OpenIB network at SC|05 is available at its Web site**.**

###

SuperComputingOnline.com

NEC Solutions America to Showcase Technologies at Supercomputing 2005

Wednesday, Nov 09 @ 10:55 PST

NEC Solutions America, a premier provider of integrated solutions for the Connected Enterprise in North America, will showcase its complete suite of high performance computing (HPC) and supercomputer technologies at SUPERCOMPUTING 2005 (SC|05), the premier international conference on high performance computing on November 14-November 17, at the Washington State Convention and Trade Center in Seattle.

In booth #714, NEC will highlight its capabilities in developing, integrating, implementing and managing large scale HPC solutions by displaying its wide range of HPC platforms, including the SX-series vector supercomputers, NEC's EM64T blade server and NEC's Intel Itanium Express5800/1000 server series, which combines the power of Intel Itanium2 CPUs with NEC's leading technology in high-end computing to deliver outstanding scalability and computing power. Additionally on display will be NEC's Fault Tolerant servers and its Storage series.

In order to best answer scientific needs of large scale and high-speed calculations, NEC has made consistent improvements to its HPC technology. NEC further underscores its commitment to the market by delivering the most advanced vector architecture, LSI technology, optical interconnect cabling and low loss PCB technology in a condensed format -- NEC's new generation of SX-8 vector supercomputers. These technologies are NEC's latest solution to create new opportunities for scientific and engineering developments.

"NEC's vector supercomputers have been at the forefront of the world's latest developments in scientific research and engineering," said Larry Sheffield, executive vice president, Solutions Platform Group of NEC Solutions America. "NEC is proud to exhibit its HPC solutions at SC|05, as we see the importance of supercomputing technology as a way to provide a strong contribution to improve quality of life and benefit society."

NEC has been providing HPC solutions based on state-of-the-art technologies for more than two decades. Currently, NEC's supercomputing technology is used to simulate many real-world applications, including weather forecast, hurricane tracking, earthquake prediction, oil and gas exploration, automotive design and nuclear energy safety, to name a few.

Berkeley Lab to Showcase HPC and Networking Leadership

Wednesday, Nov 09 @ 04:49 PST

Computing and networking experts from the U.S. Department of Energy's Lawrence Berkeley National Laboratory will share their leadership expertise via talks, technical papers and demonstrations at the SC05 conference to be held Nov. 12-18 in Seattle.

Booth Talks and Demonstrations

As part of the series of presentations in Berkeley Lab's booth (1828), representatives from DOE's three INCITE (Innovative and Novel Computational Impact on Theory and Experiment) projects will discuss their latest work into how stars and solar systems form, how to make combustion more efficient and how proteins express genetic information.

Other talks in the LBNL booth will focus on what's new at Berkeley Lab and NERSC, high performance computing benchmarking and performance analysis, large-scale calculations for nanostructures, cyber security, 10 gigabit testing on ESnet, query-driven visualization and the ACTS Collection of HPC tools. In addition to the three INCITE projects which compute at NERSC, another group of users will discuss "Core Collapse Supernovae: Trying to Explode Stars at NERSC." LBNL is also hosting a talk by two Louisiana State University researchers whose planned booth fell victim to Hurricane Katrina. A complete schedule can be found at its Web site.

Berkeley Lab will also present demonstrations of a number of tools and techniques developed to advance scientific computing and networking. Booth demonstrations will include the following:

- Large-Scale File Replication using DataMover Technology
- Integrated Performance Monitoring of HPC Workloads
- Integrating System-Wide Global Parallel File System with HPSS
- The Warewulf Cluster Toolkit
- The Advanced CompuTational Software (ACTS) Collection
- Serial, Parallel and Distributed Checkpoint/Restart for Linux
- Enabling Reliable, Secure Group Communication
- ViCE: Visual Component Editor
- Weather via Open InfiniBand
- Drosophilia Gene Expression Visualization
- Interactive Single Nucleotide Polymorphism Visualization
- SciDAC Visualization: Accelerator Modeling Pilot Program
- Visualization of NERSC Computational Science in Astrophysics, Combustion and Structural Biology

SC05 Challenges

Berkeley Lab scientists will be participating in all three challenges at SC05: HPC Analytics, HPC Bandwidth and StorCloud, which push the boundaries of high performance analytics, storage and networking applications.

A team led by Will Baird is vying for the Tri-Challenge crown by compiling the highest combined score of the three challenges. The entry is called TRI Data Storm and will be demoed in the LBNL booth.

Another Berkeley Lab team led by Kurt Stockinger is competing in the HPC Analytics with an entry called "Network Traffic Analysis with Query Driven Visualization." This entry will also be shown in the LBNL booth.

Special Award

John Bell, a senior staff mathematician and head of the Center for Computational Sciences and Engineering at LBNL, has been named as the recipient of the 2005 Sidney Fernbach Award. The Fernbach Award is given by the IEEE Computer Society for an outstanding contribution in the application of high performance computers using innovative approaches. The award will be presented at the SC05 conference in Seattle. In conjunction with the award, Bell will give a special talk in a session starting at 8:30 a.m. Wednesday, Nov. 16.

Technical Program Presentations

Berkeley Lab is also well represented in the SC05 technical program, with LBNL staff authoring or co-authoring the following five technical papers:
• "Analyzing Ultra-Scale Application Communication Requirements for a Reconfigurable Hybrid Interconnect," by John Shalf, Shoaib Kamil and Leonid Oliker of the Computational Research Division (CRD) and David Skinner of NERSC.
• "Leading Computational Methods on Scalar and Vector HEC Platforms," by Leonid Oliker, Michael Wehner and Andrew Canning of CRD, Jonathan Carter of NERSC and others from Princeton Plasma Physics Laboratory, Lawrence Livermore National Laboratory, NEC, Oak Ridge National Laboratory and JAMSTEC.
• "Apex-Map: A Global Data Access Benchmark to Analyze HPC Systems and Parallel Programming Paradigms," by Erich Strohmaier and Hongzhang Shan of CRD.
• Kathy Yelick, leader of CRD's Future Technologies Group, is an author of "Making Sequential Consistency Practical in Titanium" along with her UC Berkeley students Amir Ashraf Kamil and Jimmy Zhigang Su. Their paper has been nominated for the Best Student Paper Award.
• Additionally, Erich Strohmaier is a coauthor with UC San Diego researchers of a paper on "Quantifying Locality in the Memory Access Patterns of HPC Applications."
• Mike Welcome of the Future Technologies Group, along with scientists from LLNL, coauthored a paper on "Tera-Scalable Algorithms for Variable-Density Elliptic

Hydrodynamics with Spectral Accuracy."

On Sunday, Nov. 13, an LBNL team will present a tutorial on "Keeping Ahead of the Bad Guys: High Performance Computing Protection." Presented by LBNL's Stephen Lau, Scott Campbell and Bill Kramer, this session will focus on what can sites do to combat attackers while maintaining an open scientific environment.

Conference Organization

Finally, Berkeley Lab's involvement in this year's conference literally starts at the top — NERSC General Manager Bill Kramer is the general conference chair for SC05, and his executive assistant, Zaida McCunney, has served as vice chair of the conference committee. Other LBNL staff members have been active in the networking, communications and technical program aspects of the conference.

Berkeley Lab is a U.S. Department of Energy national laboratory located in Berkeley, California. It conducts unclassified scientific research and is managed by the University of California.

###

PSC at SC05: Networking, TeraGrid, Biomed, Big Ben

Tuesday, Nov 08 @ 15:49 PST

At SC05 in Seattle, the international conference for high-performance computing, networking and storage, the Pittsburgh Supercomputing Center will highlight scientific results achieved on Big Ben, its 10-teraflop Cray XT3. PSC's networking group will demonstrate innovative new technologies they have developed. PSC's biomedical group, the National Resource for Biomedical Supercomputing, will show a number of important life-science applications, and researchers will demonstrate PSC software that allows remote interactive control of simulations on Big Ben live from the show floor via the TeraGrid backbone.

Big Ben, which became a TeraGrid production resource on October 1, has already produced significant scientific computations that will be discussed at the PSC booth. PSC scientist Yang Wang will discuss advances in the locally self-consistent scattering method that are making possible direct quantum simulation of nano-structured materials. PSC scientist Shawn Brown will describe large-scale simulations that have elucidated a crucial step in the catalytic process of the enzyme aldehyde dehydrogenase.

With PSC-developed software that enables real-time access to simulations on Big Ben from remote locations, Paul Woodward, David Porter and colleagues at the University of Minnesota will use Big Ben via the TeraGrid backbone to simulate turbulent fluid dynamics. The researchers rely on a PSC capability - Portals Direct I/O (PDIO) - that routes simulation data from Big Ben's processors to users anywhere on the network in real time. With PDIO, Big Ben has the ability to deliver sustained data rates of 200 megabits per second, with bursts nearing 800 megabits per second.

Other prominent computational scientists will present talks describing large-scale applications running via the TeraGrid and exploiting the superior performance of Big Ben. These include Peter Coveney of University College London, Nick Nystrom of PSC, Omar Ghattas of the University of Texas, Austin, George Biros of the University of Pennsylvania, Tiankai Tu of Carnegie Mellon University, George Karniadakis of Brown University, Bruce Boghosian of Tufts University, Yoshio Tanaka of AIST, Japan, Hugh Nicholas of PSC, and Thomas Quinn of the University of Washington.

For more details of these talks and demonstrations, see: its Web site.

PSC's networking group is actively involved in networking for SC05, including support for the Open Infiniband network for SCinet. They will demonstrate the Network Path and Applications Diagnostics project, which addresses problems of path delay inherent in wide-area networks. They'll also show HPN-SSH, a high-performance secure shell that eliminates a known SSH bottleneck by allowing flow-control buffers to be defined at

runtime.

PSC's National Resource for Biomedical Supercomputing will show a variety of applications, including MCell for modeling inter-cellular microphysiology, DReAMM for model-building and visualization, the PSC Volume Browser for volumetric visualization and analysis of massive
datasets, and the Dynamo molecular simulation program for hybrid quantum mechanical/molecular mechanics simulations.

The Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Company. Established in 1986, PSC is supported by several federal agencies, the Commonwealth of Pennsylvania and private industry, and is a leading partner in the TeraGrid, the National Science Foundation's cyberinfrastructure program.

###

SuperComputingOnline.com

SilverStorm to Showcase InfiniBand Technology Leadership at Supercomputing 2005

Tuesday, Nov 08 @ 12:51 PST

SilverStorm Technologies today announced its plans for exhibiting at the Supercomputing conference November 14-17, 2005 in Seattle. In addition to showcasing several new industry-first products and technologies, SilverStorm will provide equipment and resources in support of SCinet's OpenIB multi-vendor InfiniBand backbone network. SCinet is the most advanced technology network built on site to enable world-class demonstrations and to support the annual Supercomputing conference.

In addition to the live product demonstrations running in SilverStorm's booth, SilverStorm solutions will also be showcased in an extensive set of customers and technology and channel partner booths at Supercomputing. These include Apple, Appro, Fujitsu, Intel, Mellanox, Microsoft, PathScale, Rackable Systems, RAID Inc., Sun Microsystems, Verari Systems, and Virginia Tech. SilverStorm is also partnering with Hewlett Packard who has provided a full rack of HP servers for use in SilverStorm's booth application demonstrations.

"SilverStorm is excited about the extensive presence our solutions will have across the show floor this year," said Reini Florin, vice president of marketing, SilverStorm Technologies. "This presence is a reflection of both the success of SilverStorm in the market over the last year as well as the continued rapid acceptance of InfiniBand as the interconnect solution of choice for high performance clustered computing."

SilverStorm's portfolio of director-class InfiniBand based switching and virtual gateway solutions provide users with a low latency, multi-protocol interconnect fabric that supports cluster inter-process communication and shared storage and networking I/O access. Cluster performance is significantly improved and the data center is dramatically simplified.

SC|05 Announces HPC Analytics Challenge Finalists

Tuesday, Nov 08 @ 06:17 PST

HPC Analytics, SC|05's newest initiative, has announced six finalists for its challenge that honors top technical and commercial applications developed from leading-edge, advanced analytics techniques to solve complex, real-world problems. The HPC Analytics Challenge is a unique opportunity for researchers, engineers, and analysts to showcase innovative techniques in rigorous data analysis and high-end visualization to develop and demonstrate applications showcasing powerful analytics techniques for solving complex, real-world problems. Finalists of the HPC Analytics Challenge will convene at SC|05, the premier international conference on high performance computing, networking, storage, and analysis on November 12-18.

"We were pleasantly surprised by the number of participants for this first year," said Tim Leite, HPC Analytics Committee Chair. "The numbers, quality and range of applications greatly exceeded our expectations, and the presentations will generate great excitement for next year's initiative."

The HPC Analytics program finalists include:

Terascale Music Mining – this project's goal is to create secure, accessible, terascale collections of music materials in a variety of audio, symbolic and metadata forms.

Network Traffic Analysis with Query Driven Visualization -- this entry is used in the field of information management to successfully perform queries, feature identification and characterization of 24 weeks of netflow data consisting of 1.1 billion records of 25 fields each.

SPICE: Simulated Pore Interactive Computing Experiment -- the project aims to understand the vital process of translocation of biomolecules across protein pores by computing the free energy profile of the translocating biomolecule along the vertical axis of the pore.

Bridging the Macro and Micro: a Computing Intensive Earthquake Study using Discovery Net -- results from this group provide valuable scientific insights into earthquake analysis, and also showcase how novel HPC methods and knowledge discovery methods enable such results to be achieved.

Visualization of Large-Scale Unsteady Computational Fluid Dynamics Datasets -- a new scientific visualization tool has been developed for investigators to more effectively interrogate large unsteady datasets.

Real Time Change Detection and Alerts from Highway Traffic Data -- this team has developed a test-bed containing real-time data from over 830 highway traffic sensors in the Chicago region, weather and text data about events that might affect traffic. The goal is to detect interesting changes in traffic conditions in real-time.

The finalists will present their applications on Tuesday, November 15, at 1:30-5:00 p.m during the SC|05 conference. The awards ceremony will commence at 3:30 p.m. on Thursday, November 17. For more information visit its Web site.

###

SuperComputingOnline.com

NCSA to showcase high-definition scientific visualizations at SC05

Monday, Nov 07 @ 12:44 PST
The National Center for Supercomputing Applications (NCSA) will feature striking high-definition stereo visualizations of scientific phenomena ranging from the prediction of destructive weather to the exploration of the mysteries of black holes at Booth 1630 at SC05 (to be held in Seattle Nov. 12-18). The high-definition stereo visualization theater at NCSA's booth incorporates two JVC HD2K projectors with 1920x1080 native resolution, D-ILA technology, and circular polarization for passive stereo. These components are all integrated into a portable enclosure with a 6-foot-wide screen, which was created by Visbox, Inc.

This equipment will be used to display high-definition stereo visualizations created by NCSA's scientific visualization experts, including:
* an interactive three-dimensional visualization of an analysis of a collection of text documents driven by the D2K data mining software developed by the center's Automated Learning Group
* a virtual voyage from our sun through the Milky Way galactic center toward a massive black hole at the center of our galaxy; the scene is excerpted from the upcoming Denver Museum of Nature and Science ultra-high resolution digital dome show "Black Holes: The Other Side of Infinity," which will premiere in 2006.
* a visualization of an F3 tornado within a simulated supercell thunderstorm that was developed in collaboration with atmospheric scientists at the University of Illinois. Using complex computational models, weather data recorded during an F4 tornado that devastated South Dakota with winds in excess of 200 mph, and NCSA high-performance cyber-resources, the team simulated the birth and development of a tornado, generating terabytes of data that the visualization team translated into an artistic, information-rich animation.
* a visualization of a simulation of stratified fluid flow and derivative trajectory data by a scientist at Scripps Institute of Oceanography.

Other displays in the NCSA HD theater will showcase simulations and visualizations created by researchers from the Center for Computation & Technology at Louisiana State University. In a project that hits close to home, LSU scientists Edward Seidel and Shalini Venkataraman will compare computational simulations of the strength and path of Hurricane Katrina generated by storm surge modeling to actual observations of the devastating storm, including storm surge measurements from sensors, GIS data from aerial mapping, and time-varying multi-spectral satellite images. Using flow visualization techniques, the researchers will show time-varying 3D wind fields and storm surges overlayed with large geospatial data from GIS and remote sensing sources.

And turning from the Earth to the stars, Seidel and Werner Benger will show stereoscopic visualizations of curved space, using computer simulation to reproduce the strange visual phenomena which occur near the extremely powerful gravitational fields of black holes.

HPCwire
Thu, November 17, 2005 10:02 am

Borchers Keeps SC Streak Alive

Bob Borchers can recall when the Supercomputing show was nothing more than a grand idea. The calendar had just flipped to 1988 and eight influential figures from the HPC community were gathering around preparing to utter the immortal line that fueled a flurry of Mickey Rooney-Judy Garland movies: "Let's put on a show!"

Borchers remembers it was January and there was urgency that this new show should become a reality later in the same year – fighting conventional wisdom (and reality) that planning and producing such an event would take at least a year, no less a matter of months.

Despite the obvious obstacles, the founding octet remained firm in their conviction that the supercomputing universe needed an annual event around which to huddle and hawk the latest wares and advancements. They pushed on, perhaps not fully realizing the enormity of their task. Finding adequate conference space in an attractive location for an event they hoped would attract 1,000 attendees proved hard.

Finally, they settled on Kissimmee, Fla., at the Hyatt just down the road from Disneyworld. Ever wondered why Supercomputing is held so late in the year? The answer is simple; look back to its first year. "The only date that was available was the week before Thanksgiving," Borchers explained. "Most events are booked out years in advance." The rest, as has been said many times before, is history - and Borchers will be writing a little bit more of it with his 17th consecutive SC appearance this week in Seattle.

He said the rest of the founding members are no longer involved, so he takes special pride in carrying on as the conference's surviving link to the creation of the initial soiree. Also still around since day one are industry organizations ACM and IEEE. SC certainly has rewarded both groups for their continued cooperation.

Every year Borchers seems to find some way to keep his hand - and mind- in the SC picture. He fondly recalls being general chair for the Portland event in 1993 and also remembered San Jose in 1997, where he was network chair. "We've come a long way from the days when we thought it was enough to have six OC-48s running into the show," he commented.

While he has not slowed down his zest for the annual gathering, he admitted one concession: "I've given up trying to get SC to come to Hawaii," Borchers said with a laugh. As the 50th State sustains an unfair rap for being the prime site for perceived boondoggles, he added, "We actually find a way to get a lot of work done there."

At the Maui High Performance Computing Center, an Air Force research lab managed by the University of Hawaii, Borchers stays busy helping provide more than 10 million hours of compute time a year on several supercomputers, including a Cray Inc. XD1, Power3 and Power4 configurations from IBM as well as large Linux clusters. MHPC supports a diverse base of users including the Department of Defense as well as government, research and science users.

MHPCC has a contingent of 18 people at SC05. He expects, aside from booth duty, he and his colleagues will be busy with a slew of vendor briefings as well as investigating some new companies under non-disclosure agreements. "We can't always get companies to come visit us, so we have to take this opportunity to see as many as we can," he explained.

In addition, aside from his participation on this year's InfoStor committee, Borchers said he is working as "an understudy" on the Bandwidth Challenge as he prepares to lead that effort in 2006.

Looking ahead, Borchers said it is also noteworthy that two women are on deck to chair the next two SC extravaganzas. "There has always been a large contingent of women involved in supercomputing," he said. "Next year Barbara [Horner-Miller of the Artic Region Supercomputing Center] and then [in 2007] Becky [Verastegui of Oak Ridge National Laboratory] will carry on that legacy."

Borchers took particular note of SC07 because that conference marks a return to Reno, Nev., the locale of SC89. "Of course, it will be in the new convention center," he added, a triumphant comment on how far the conference has come since its sophomore year.

Carrying on the tradition of high-powered keynotes, as started by Seymour R. Cray, a name synonymous with HPC and the man who created the industry that now bears the name of this conference, Borchers said he looks forward to what Microsoft's Bill Gates has to say to the HPC faithful. "He's a thoughtful guy who's worth listening to," Borchers said. "He will surprise a lot of people with how much he knows about HPC."

###

ComputerWorld

Supercomputing Pushes Toward the Corporate IT Mainstream

News Story by Patrick Thibodeau

Current users cite business gains and call for wider adoption. But some barriers remain -- including a need to convince CIOs to back the technology inside their companies.

NOVEMBER 21, 2005 (COMPUTERWORLD) - SEATTLE -- Loren Miller, director of IT research, development and engineering at The Goodyear Tire & Rubber Co., thinks he can easily make the case for wider corporate use of supercomputing technology.

Simulations made possible by supercomputing have enabled Goodyear to gradually reduce the amount of money it spends on building physical tire prototypes, from 40% of its total research and development budget to just 15%, Miller said last week. The Akron, Ohio-based company is using the money it saves to fund more research work.

"From our standpoint, the results have been dramatic," Miller said at the Supercomputing 2005 conference here. Other companies in the U.S. need to realize that they can gain a competitive advantage from high-performance computing systems, he added.

Efforts are under way to broaden the corporate base of supercomputing users. For example, the Ohio Supercomputer Center (OSC) in Columbus has launched a program called Blue Collar Computing that's designed to provide businesses that lack high-performance computing expertise with tools for testing the technology.

And in a speech at last week's conference, Bill Gates, Microsoft Corp.'s chairman and chief software architect, predicted that one day some supercomputers will cost less than $10,000. He also said that "mass computing" and supercomputing share common technical challenges and could benefit from combined R&D efforts.

William Kramer, head of high-performance computing at the National Energy Research Scientific Computing Center at the Lawrence Berkeley National Laboratory in Berkeley, Calif., said that Gates' appearance was an indication of the growing awareness of supercomputing's importance. "The output of [high-performance computing] activities are no longer hidden behind a curtain, if you will," said Kramer, the conference's general

chairman.

Supercomputing is "being scaled down so more people can make use of these very complicated tools," he added. "And I think that's one of the indications of Microsoft's interest here."

Like Gates, Stanley Ahalt, the OSC's executive director, envisions wide-scale use of high-performance systems by companies looking to run complex simulations and visualizations of products during the design and testing process. The OSC, which is beginning to talk with potential commercial users of its systems, hopes to encourage businesses to adopt the technology by offering help as well as access to its supercomputing resources.

Ahalt said he's convinced that supercomputing is critical to improving the competitiveness of U.S.-based companies. But he thinks that many IT managers still aren't even considering the technology.

"CIOs are focused so acutely on the bottom line, they aren't ready for the next big thing," Ahalt said. IT managers, he added, have to take the message about supercomputing's potential value to corporate executives "and explain that we are about to go through one of these radical shifts on our economic systems." Shortening product-development cycles will become even more important to companies than it already is, Ahalt predicted.

There were about 9,250 attendees at the supercomputing conference -- more than a 10% increase from the prior year's. But many appeared to be from large companies, life sciences firms, national laboratories and academic institutions -- the types of organizations that have already invested in supercomputing.

Better Product Designs

The Procter & Gamble Co. uses high-performance systems to run computerized visualizations of products that it's developing. Thomas Lange, director of modeling and simulation for corporate research and development at P&G, spent time at the OSC's trade show booth last week and spoke highly of the Blue Collar Computing initiative.

Some of the products that Cincinnati-based P&G buys from its suppliers, such as bottle

caps, could benefit from computational design programs that run on supercomputers, Lange said.

"Oftentimes, these suppliers could be making lighter products, stronger products, better products," he noted. "It's the suppliers that supercomputing can make a difference for."

Phoenix-based Ping Inc. isn't a typical supercomputing user. The maker of golf clubs is a midsize company with about 1,000 employees. But earlier this month, Cray Inc. announced that Ping is using one of its supercomputers to simulate golf club designs.

The Cray XD1 system installed at Ping is based on 12 Opteron processors from Advanced Micro Devices Inc. and has 24GB of memory. Eric Morales, a staff engineer at Ping, said the system has enabled the company to drastically reduce product development times. Simulations of product changes that once took a full day to run can now be processed in 20 minutes or less, Morales noted.

"It takes the development [cycle] from weeks down to days, and it helps us get to market faster," he said.

Gaining Ground

Earl Joseph, an analyst at Framingham, Mass.-based IDC, said he expects the worldwide high-performance computing market to reach $7.25 billion this year, a net increase of 49% since 2003.

The technology should continue to gain ground with both technical and commercial users, "primarily due to the attractive pricing of clusters, combined with their growing capabilities," Joseph said.

Price/performance improvements, which are partly the result of increasing use of commodity processors, are helping to make supercomputing more accessible to businesses.

An annual listing of the world's 500 largest supercomputers, which was released last week, showed that more than three-quarters use processors from Intel Corp. or AMD. That's up from 41% in the 2003 listing and just 12% the year before that, according to the researchers who compile the list.

Dave Turek, vice president of IBM's Deep Computing operations, said demand forsupercomputing systems is being driven by forces such as the emergence of new businesses that rely heavily on high-performance systems to support uses such as digital animation and bioinformatics.

In addition, many companies are collecting "vast amounts of data that demand rapid analysis for real-time decision-making," Turek said. In particular, he pointed to the growing use of radio frequency identification devices to track products.

But there are limiting factors as well. Although both the price and performance of supercomputing hardware have improved dramatically, the same isn't true for much of the software used on high-performance systems. Programs such as fluid dynamics applications can be costly because demand for individual products is still relatively limited. And according to a study released by IDC last summer, many software vendors aren't increasing the scalability of their code to take advantage of systems with hundreds or thousands of processors.

Concerns about security are also an issue, particularly for corporate users of shared high-performance systems.

For example, WestGrid, a Calgary, Alberta-based high-performance computing consortium involving seven Canadian universities, has made its systems available to a number of companies for research uses. But it has found that some businesses are reluctant to use the systems for competitive reasons.

"They don't want two companies working on the same problem sharing [computing] resources," said Rob Simmons, a distributed systems architect at WestGrid.

Eric Lai contributed to this story.

###

GRIDToday
Tuesday, November 15

Features:

Calient, MCNC, LSU CCT to Advance Scientific Research

Calient Networks, MCNC, the Louisiana Optical Network Initiative and the Center for Computation & Technology at LSU announced at Supercomputing 2005 a partnership to drive optical networking and Grid computing toward a new realm of advanced scientific applications.

Calient, a carrier-class photonic switching systems and software provider, and MCNC, a network research facility and nonprofit provider of network services to North Carolina universities and state government, have partnered to integrate optical control plane and Grid computing technologies for research purposes. Working with the Center for Computation & Technology (CCT) at LSU in conjunction with the Louisiana Optical Network Initiative (LONI), MCNC and Calient are providing new optical network capabilities to drive extreme-scale scientific applications. At SC'05, CCT and LONI will demonstrate how the control of high-speed networks can be coupled to computing resources, data storage archives, and visualization services for enhancing the study of highly complex and computation intensive scientific problems such as the modeling of black holes and other sources of gravitational waves.

"Network infrastructures are no longer commodity 'plumbing' tools but integral components for 'virtual' research teams that rely on resources, both human and physical, residing in various geographic locations," said John Crites, president and CEO of MCNC. "The combined technological expertise between MCNC and Calient together with other research partners will increase innovation from National Research & Education Networks (NRENs) and foster more scientific collaboration."
"Calient was an early pioneer in delivering optical switching and GMPLS technologies, and has developed great expertise in NRENs demonstrated by their successful involvement in the SuperSINET and JGN-2 programs in Japan," said Mark Johnson, chief technology officer at MCNC. "We are pleased to partner with them, having already seen positive results."

"MCNC operates the nationally-recognized North Carolina Research and Education Network (NCREN) and has a world-class team with a history of pioneering advanced network research," said Charles Corbalis, CEO and president at Calient. "Not only are they leading the Experiment Support Services effort for the National LambdaRail but they are also involved in many other advanced network projects. We are committed to helping them achieve their goals of supporting next-generation network initiatives."

Calient's DiamondWave PXC optical switch has been deployed at MCNC's site in Raleigh, and will be the base of a new optical exchange to spur research activity throughout North Carolina universities and facilitate national and international research.

A working relationship already existed between MCNC and LONI based on shared experiences in operating regional networks dedicated to research and education. "We recognized the value of bringing Calient switches to both LONI and NCREN to enable better provisioning of our optical network resources," said Johnson. LONI has already achieved early success with integrating the Calient DiamondWave PXC product deployed in Baton Rouge with Grid middleware developed by the CCT.

"The ability to dynamically allocate bandwidth, coordinate network intelligence, and vertically integrate Grid computing with network services is a linchpin for worldwide NREN collaboration," added Gigi Karmous-Edwards, principal scientist at MCNC and also the chair of the Control Plane and Grid Integration Middleware Working Group at the Global Lambda Integrated Facility (GLIF).

"Since our research required real-time cooperation and dynamic connectivity between supercomputer sites in Louisiana and with European collaborators we needed new transport protocols and dynamic optical network configurations," said Ed Seidel, CCT director and research scientist. With the network and attached resources, researchers can also share huge data sets of information around the world instantly, such as complex models of storm surge or path predictions.

"More accurate projections can now be realized," said Seidel. "Network provisioning allows for faster, complex real-time problem resolution."

"Despite the impact of recent hurricanes in Louisiana, together with Calient and MCNC resources and researchers, we were able to showcase this new network capability in a relatively short time period. These recent disasters proved the value of collaborative support," said Charlie McMahon, director of LSU Office of Telecommunication. As soon as communications were re-established after the storms, MCNC and Calient personnel were in touch with LSU CCT and doubled their efforts to help LONI get ready for networked computer simulation demonstrations at the October iGRID 2005 and SC'05.

###

GRIDToday
November 15, 2005

Features:

New MATLAB Tools Integrate w/ United Device's Grid MP

United Devices announced that the next generation of MATLAB distributed computing tools, unveiled this week by The MathWorks, will integrate with United Devices' Grid platform.

Because MathWorks provides a generic API in its new products, the Distributed Computing Toolbox 2 and the MATLAB Distributed Computing Engine 2 will easily combine with UD's Grid MP platform to deliver flexible compute power that always satisfies variations in demand from research, analysis and operational applications. MathWorks products let users develop distributed MATLAB applications and execute them in a cluster of computers managed by a third party scheduler to take advantage of an organization's wider IT resources.

"Integrating advanced applications like those from The MathWorks with an organization's overall IT infrastructure is a logical extension of those products' capabilities," said Jikku Venkat, chief technology officer at United Devices. "Scientists and engineers familiar with the highly sophisticated functionality of MATLAB already understand the value of distributed processing and naturally embrace the need to maximize all of an organization's compute power to accelerate and expand research. Combining MATLAB with UD's grid solutions opens new frontiers in this area."

Working together, the Distributed Computing Toolbox 2 and the MATLAB Distributed Computing Engine 2 support high-productivity computing by letting users take advantage of the high-level MATLAB language for algorithm development and then accelerate computation via distributed or parallel execution.

"In the operating mode enabled by Grid MP, execution occurs on a broad and flexible set of resources that can contract or expand based on computational demand," said Lisa Kempler, director of MATLAB product marketing at The MathWorks. "This mode can be used for large data manipulation tasks, which were previously unattainable in a single computer, to take advantage of available resources."

GRIDToday
November 15, 2005

Features:

Mathematica Announces Personal Grid Edition

Mathematica Personal Grid Edition eliminates the barriers to using parallelism as part of your daily workflow -- with no administrative overhead and no contending for shared resources. Users can tackle larger problems and investigate parallel approaches at any stage of the problem-solving process right at their desks.

Mathematica Personal Grid Edition combines the computational capabilities of Mathematica with high-level parallel language extensions to create an optimal computing framework for quad-core machines. Mathematica's high-level language and numeric and symbolic solvers make it ideal for doing exploratory technical computing in a wide range of fields, resulting in increased productivity and pushing the limits further than ever before.

The Mathematica 5 series (5.0, 5.1, 5.2) provided numerous performance and scalability enhancements to the software of choice for demanding technical computing applications. All-platform 64-bit computing has broken the memory barrier, while fast machine-optimized dense and sparse linear algebra gives Mathematica unparalleled numerical speed. These versions also incorporate major new functionality for connectivity -- including built-in universal database link, Web services and numerous file format converters -- as well as hundreds of new numerical and symbolic algorithms.
Grid Mathematica extends this high-level framework to enterprise parallel computing. Now Mathematica Personal Grid Edition is the only solution to bring high-performance supercomputing to the individual desk, continuing the Wolfram Research commitment to cutting-edge innovation.

"Traditionally, parallel computing has been the realm of specialist programmers with limited understanding of different application fields. Researchers had to cross multiple barriers -- including prohibitive cost, administrative overhead, remote locations, and hard-to-use APIs -- to perform parallel computations, often putting this resource beyond their reach," said Roger Germundsson, director of research and development at Wolfram Research. "Now parallel computing can directly be used by individual researchers as an integral problem-solving method. With Mathematica Personal Grid Edition and a quad-core machine, you have an immediately useable environment for exploring parallelism at any level."

Mathematica Personal Grid Edition offers performance improvements of up to 300 percent over standard Mathematica, with only minimal code modifications. For more-intensive parallel applications, programs can run unchanged on Grids or clusters of any size using gridMathematica. This also makes Mathematica Personal Grid Edition a nice prototyping environment for large-scale parallelism, and adds another dimension to

Wolfram Research's ability to conveniently handle all of your supercomputing needs within the same basic framework.

###

GRIDToday
November 15, 2005

Features:

Moab Grid Suite Advances Resource Control

Cluster Resources Inc. announced the release of Moab Grid Suite 4.5, its new Grid management solution that lets organizations consolidate management of multiple clusters to create usable, affordable Grids while maintaining local cluster sovereignty.

Many large-scale Grid sites trust Moab for their Grid creation and management, consolidating local clusters into Grids and connecting to remote resources throughout the world. Moab-based Grids allow them to reduce management overhead, improve job responsiveness and utilize the fine grained control and ease-of-use that Moab provides.

Moab reduces the barriers to enabling Grid computing so organizations can focus on getting the most out of their compute resources. Moab gives organizations the capabilities they need to create a useful, productive Grid across the resources they already own and scales to allow plenty of room to grow. Its flexible and robust policy engine provides solutions to the political problems that come with sharing resources. With its integrated design, Moab makes it simple to go from independent multiple clusters to a high-performance Grid.

While many of the industry's Grid solutions are only compatible with limited resources, Moab can unify multiple heterogeneous software and hardware environments. Moab can work on top of multiple resource managers (LSF, PBS Pro, TORQUE, LoadLeveler, SLURM, BProc, etc.), operating systems (Linux, Unix, Mac, Windows, etc.) and hardware (SMP machines, 32-bit, 64-bit, etc.). This allows organizations with multiple clusters the freedom to move to a Grid while leveraging all of their existing infrastructure.

Moab Grid Suite ensures that even completely heterogeneous systems are highly utilized by intelligently assigning the best-suited hardware to each job. With optimized job and data migration and integrated user and file mapping facilities, users can take advantage of the Grid with virtually no change in submission habits. Moab also tracks both network and storage limitations when making Grid scheduling decisions to better guarantee successful job execution and optimized Grid-level performance.

Moab eliminates the need to change existing cluster policies, underlying resource managers, job scripts and other cluster-specific configurations. Moab's broad compatibility make it a viable option for organizations seeking to consolidate control over disparate cluster resources into a single flexible resource that can continue to apply previous usage controls while adding new, simplified Grid level management.

Challenges in creating a Grid often stem from reluctance of individual cluster **Moab**

management to give up control of resources. Moab Grid Suite makes resource sharing easy between multiple organizations through two methods. First, it allows administrators to maintain sovereignty over their own resources, even in a Grid environment. Second, Grids can be configured in different management styles, including traditional, centralized Grids where there is one central administrator who controls all Grid policies, and a new peer-to-peer Grid model where workload and resource status information can be exchanged and jobs and data can be migrated between trusted peers with any local applied policy. These Moab-unique capabilities let individual clusters maintain independence while still collaborating with the Grid at their desired level.

Moab's advanced policy engine allows organizations to guarantee its users a certain quality of service (QoS) and deliver it. Grid administrators can use Moab to automate policies or set user and group priorities so the highest priority jobs get done faster. Moab also provides organizations with reporting and billing tools so administrators can know exactly how much of the Grid's resources an organization used and, if desired, bill against the resulting information to help share costs fairly.

Moab's Virtual Private Cluster (VPC) feature lets administrators allocate a section of a cluster or Grid or a specific amount of resources to users or groups. These users or groups can then only view, track and allocate resources within their own VPC. The VPC can be used to give political groups the resources they demand while administrators maintain central control of all resources. Using this approach, individual clusters can participate in multiple Grids simultaneously with distinct rules of engagement for each Grid. With the political sharing features Moab provides, Grid organizations can assure local contributing groups that their needs will be met, without compromising the service delivered to other members of the Grid.

"While powerful, one of the highlights of Moab Grid Suite 4.5 is its incredible simplicity," said David Jackson, Cluster Resources' chief technology officer. "We have been doing Grids for a number of years and have found that one of the greatest barriers to Grid adoption is the fear associated with introducing an entirely new layer of software and complexity. With Moab Grid Suite, these issues are effectively eliminated. In many cases, a cluster can be rolled into a Grid in a matter of seconds with no new tools and no new training."

Moab Grid Suite provides features to simplify adoption including a Web access portal that lets users remotely submit and manage jobs through a graphical interface. Moab further simplifies grid computing by translating resource manager commands and scripts, reducing training for end-users and letting them work across multiple resource management environments with scripts knowledge they already have.

Moab makes management easier for administrators with its virtual view of resources that provides information on downed, idle and active nodes, eligible, running, on hold and

blocked job status, as well as resource consumption information by user, project, organization or other groupings. Moab's in-depth reporting tools also allow administrators to view previous months' Grid activity, letting them create reports to see how a policy changed workload or to ensure users are receiving their agreed QoS. Moab's simplistic features help organizations set up grids without lost productivity. It benefits organizations by reducing training time and letting them see what is really happening on their Grid.

###

GRIDToday
November 15, 2005

Features:

Sun's Bjorn Andersson Gives Insight into Big Wins

By Derrick Harris, Editor, GRIDtoday

GRIDtoday editor Derrick Harris spoke with Bjorn Andersson, Sun Microsystems' director of HPC and Grid computing, about what's going on with the company in terms of its Grid and HPC initiatives. Andersson discusses Sun's big customer win for its Sun Grid and the state of the Grid market, and also hits upon some the company's HPC on-goings, such as its largest HPC win to-date with the Tokyo Institute of Technology.

GRIDtoday: How are things going with Sun's Grid program?

BJORN ANDERSSON: We have seen significant momentum in our program. Most recently, we announced the biggest purchase yet of compute cycles from the Sun Grid. VCC bought 1 million CPU hours for use in the oil & gas industry. We're very excited to go live with such a big deal as one of the first publicly announced customers.

Gt: Can we expect to see more Sun Grid customers announced in the weeks and months to come?

ANDERSSON: We expect this to be just the beginning of several wins with the Sun Grid. This win may have been one of the larger wins to date however; going forward you will see a mix of larger and smaller customers using the Sun Grid.

Gt: Are companies ready to fully embrace utility computing? If not, when do you think it will reach its peak of ubiquity?

ANDERSSON: I would say it's a matter of how comfortable a company is with the Grid technology and how mature they are in establishing the way of working that Grid. Going forward the obstacles will not be so much on the technology side, but it will be more a matter of corporate culture and how you set up service level agreements between participants in an internal Grid infrastructure versus setting up policies for when you utilize a public Grid. We as an industry will get through this, but it's a learning curve. Many companies have embraced Grid technology in their internal infrastructure today. The volume ramp in terms of usage is well on its way. We're now also starting to see companies use the public Grid infrastructure of the Sun Grid. We expect the volume ramp to follow. I believe we're not anywhere close to a peak in usage of Grid technologies internally at companies and organizations. Utility computing can be seen as building on that foundation, so it's even harder to try to predict a peak for utility computing.

Gt: What about "traditional" Grid computing? Given the current state of the market,

when do you think Grid will become as ubiquitous as the community would like?

ANDERSSON: When you say "traditional" Grid computing, I think of where Grid computing started in terms of what problems it was used for -- traditional scientific and technical computing. We're at Supercomputing 2005 this week, so allow me to look at this from a high-performance computing standpoint. In this area, we have seen a very pronounced trend over the last several years of more and more computing being done on clusters rather than on big monolithic machines. We're also seeing a blurring of the lines between the traditional definitions of Grid and clusters. The whole movement toward cluster computing is also enabling Grids to be more commonplace and used where applicable, while establishing some of the processes and tools needed along the way. What really determines what customers use is the nature of the problem they're trying to solve and what requirements are on latency, bandwidth, memory size, etc. I believe we'll continue to see a diverse environment in HPC that will integrate both clusters and Grids as it applies to the computing requirements.

Gt: What is Sun's take on the current discussions around the need to adopt quality standards before enterprise adoption will really take off?

ANDERSSON: Sun has always worked for standardization of interfaces and continues to do so also in the Grid space. To really reach a level of ubiquity that we want, we need a solid foundation of standards. Then it's up to all vendors to compete on the implementation of these standards.

Gt: Can we expect any big announcements at SC'05?

ANDERSSON: HPC is one of the key focus areas for Sun and we're really investing to accelerate our growth in this market. Just last week we opened our Solution Center for HPC in Hillsboro, Ore. It's a 10 Teraflop capable benchmark and testing facility designed to help customers find their HPC solutions and for us to help remove risk and uncertainties in HPC deployments.

This week, we're also announcing Sun's largest HPC win to-date -- that the Tokyo Institute of Technology (Tokyo Tech) has purchased Sun Fire x64 servers to build Japan's largest supercomputer. Tokyo Tech's project will use Sun Fire x64 servers with 10,480 AMD Opteron processor cores, Sun storage technologies and NEC's integration expertise to build the Tokyo Tech supercomputer. The system will help provide researchers with compute power for a wide range of scientific applications, such as analysis of the complex molecular structure of proteins, simulated blood flow diagnosis in human brains, and clarification of the generation mechanism of Earth and planetary magnetic field We have a big presence at the show and I would recommend anyone to come by our booth (#1416) to take a look at what we're doing in HPC.

Gt: How do shows like SC help to further the cause of Grid computing? Does it help that new Grid technologies and projects are being presented to individuals spanning the entire HPC spectrum?

ANDERSSON: Shows like Supercomputing are great as focal points to bring key individuals together for sharing ideas and provides a unique showcase of what's available in products and technologies. It really helps the whole community to be exposed to this on a regular basis.

Gt: Finally, how do you feel about the SC shows? Is there anything you're especially looking forward to at this year's event?

ANDERSSON: The Supercomputing shows are always full of exciting products, technology demonstrations and events. It's hard to get the time to see it all. Just to give you a couple of examples, in the Sun booth we'll have a demo of OpenIB connected to the InfiniBand fabric at the show. Related to InfiniBand we're also participating in Storcloud with data management products from former StorageTek, which Sun acquired in August of this year. We also will have a remote visualization technology demo in the booth. The list goes on and on, but you can see some trends. InfiniBand seems to have cleared the hurdle to be a key low latency interconnect. Open source and free software continues to be important, we expect to see a lot of interest in OpenSolaris and the fact that the Sun Studio developer tools now are available with a no cost user license. Also, if you're interested in what's around the corner, take the opportunity to learn about the DARPA/HPCS project. Sun, as a key participant in that project, is showcasing some technologies, including a new programming language called Fortress.

Gt: Is there anything else you'd like to add?

ANDERSSON: From Sun we believe that we can make HPC much more attainable and practical. As we do that, it is becoming a tool for many more usages and starts to become a mission critical element for commercial enterprises, whether it's how you do crash simulation on cars, how you do analysis to figure out how to extract more oil out of the ground or finding that next drug that will cure a disease. Once it's being used for the core part of the value a company delivers, it's by definition mission critical for that company. Sun is applying engineering disciplines from high end data center designs and our ability to package solutions together as a way to make it much easier and reliable to build out a HPC infrastructure.

###

GRIDToday
November 15, 2005

Features:

United Devices Founder on Grids, Clusters

By Derrick Harris, Editor, GRIDtoday

GRIDtoday editor Derrick Harris spoke with Ed Hubbard, founder and chief marketing officer of United Devices, about the company's strategy around SC'05. its recent acquisition of France-based GridXpert (a leader in the manufacturing sector), the issues of Grid standards and Grid adoption, and the differences between Grids and clusters -- a topic that arises quite often in the HPC space.

GRIDtoday: How's everything going with United Devices? Will you be announcing any news at SC'05?

ED HUBBARD: Business continues to grow with several of our large customers moving their Grids to a central services model and offering them enterprise-wide.

At SC'05, we'll have a number of announcements, three as a matter of fact. We'll have an announcement with HP, Microsoft and an ISV partner.

Gt: What about the news of the past several months? Can you comment on the GridXpert acquisition? How has it affected the company in terms of customers in new markets and geographies?

HUBBARD: Acquiring GridXpert and its best-of-breed meta-scheduling technology has accelerated a number of our deployments by leveling the complexities of multiple, incompatible schedulers/DRMs. Not surprisingly, we've also opened up a number of conversations within our expanded customer base around the products and services offered by UD. Finally, part of our reasoning for this acquisition was to significantly enhance our position in manufacturing and I believe this has been more successful than we projected. The former GridXpert team (now UD, of course) had a fairly large number of evaluations going on that we've been able to add a lot of value to from a services and technology perspective. I anticipate a very good conversion rate on these evaluations.

Gt: Moving on to the Grid industry as a whole, I'm wondering where you see the market heading in terms of widespread adoption? Would it benefit the cause if the community could agree on some quality standards?

HUBBARD: We continue to architect around open Web services standards, which have served us well with customers. In fact, this is one of the areas where we continue to receive a lot of positive feedback -- namely the ability of our customers to easily create application services and portals around our infrastructure in environments like Java and C#.

On the topic of standards, I don't think they've held the industry back as they've largely been irrelevant to date (at least in our deployments). Our view is that this is changing but not in the way you might expect. We actually see more divergence, which was another problem we wanted to address with the GridXpert acquisition. By embracing the reality of the market (multiple, incompatible schedulers / DRMs) and providing a solution to eliminate that unnecessary complexity for our customers, we believe we have a superior product set to solve what our customers really face in their everyday production environments.

I don't think this is going to change very quickly since there is a huge gap between public/press perception and commercial reality that needs to be closed with some open conversation:

- There is a perception out there that WSRF/OGSA & Globus are standards.
- The reality is that Globus is merely a Grid software toolkit. In other words, it is an implementation geared toward academic environments where it has been somewhat successful -- it is not a standard by any means. OGSA is really WSRF and is currently under review at OASIS and it does not address current enteprise issues around interopability. Most of the requirements addressed in the current document reflect academic needs of running multiple Grid services across different organizations.

We've tried hard to stay out of this debate and just be customer-driven, but this is the core problem of standards for the Grid market that no one seems willing to admit or discuss openly.

Gt: How is it that United Devices is so successful in the Grid space given the somewhat disappointing enterprise adoption rates? Does it have something to do with simply targeting the right customers in the right industries?

HUBBARD: We have 100-plus customers today ranging in size from 20 nodes to 3 million, but this hasn't been an overnight success. We have worked hard and slugged it out with more than a dozen other players that have come and, thankfully, gone. This has left a pretty stable set of solutions that have settled into that age-old competitive triad (UD, DataSynapse and Platform) where there are three quality solutions in the market that all own their respective verticals and compete in the verticals where no player has a clearly dominant position.

I would argue that Grids have already penetrated the majority of the Fortune 500 in some form even if they haven't been adopted as an enterprise-wide service yet. For UD, we chose to focus early-on on life sciences and have been very successful in that vertical market. We've built up a ton of credibility, a high-quality set of references and large, complex deployments that I'd put up against any competitor's reference set. With GridXpert, I hope we've successfully telegraphed that we plan to repeat this same success

in the manufacturing and other environments and, in fact, are already well on our way.

Gt: On to Supercomputing ... what is it about this show that draws so many attendees and exhibitors? For how long has United Devices been exhibiting at the show?

HUBBARD: I think the show draws well because of the natural progress of technology, e.g. what used to be a highly specialized supercomputing application 10 years ago is now running on a Grid and the results can be visualized and manipulated on a single workstation. The democratization of these engineering analyses has continued unabated and the show where you can go to discuss them across the board remains the SC'XY shows.

UD has either exhibited or attended the show (based on our strategy for the given year) since 2000.

Gt: Seeing as how United Devices operates in both the Grid and cluster markets, which do you tend to focus on at SC events, which seem to focus a more on the cluster/HPC market?

HUBBARD: We typically focus on clusters at the SC'XY shows for both practical/logistics reasons and because we just think its a better fit from an application perspective. If you look across the set of applications in use and demo'd at the SC'XY shows, you'll find a large number of MPI-based applications that really require the speed and low interconnect latency of a cluster to run effectively. A Grid is simply a set of resources -- these resources could include clusters. We don't see a Grid market as different from a cluster market. We can effectively manage clusters stand-alone or as part of larger Grid deployments through our Grid MP platform (MPI jobs included).

Gt: Finally, and I'm starting to ask this question a lot because I feel that there is still some confusion out there, can you define, at least from a United Devices perspective, the differences/similarities between clusters and Grids?

HUBBARD: Grids are the superset in our vernacular. A Grid can be made up of any available device that may be useful to the application type(s) the enterprise is trying to accelerate, raise the reliability on, etc. That said, we love starting somewhere on a cluster because we can clearly articulate how this single, little cluster could grow up to be the enterprise's Grid -- we call this a "Grid-ready cluster." Many of our current implementations started out as a cluster and built a Grid of clusters and other resources.

Clusters are, in general, simpler animals. They are typically self contained, built from homogeneous hardware, have dedicated, private-switched networks and storage, are managed carefully by dedicated staff and are often a better path from a price/performance perspective for a number of applications.

Grids, on the other hand, often include a cluster, or multiple clusters. Sometimes running the same underlying scheduler/DRM, but sometimes not. They can also contain servers that are part of "'hot" recovery sites, underutilized machines from, for example, a company's Web farm that do very little at night but which can add significant power to a Grid with high overnight processing demands. Workstations and PCs are also often a part of enterprise Grids and this adds complexity around security requirements, data management and caching, etc., that we've tuned for in our Grid MP platform. Finally, Grids become a much larger, integrated part of the business's core value creation and as such absolutely require a higher level of manageability than a single cluster. This area is one where we have been increasingly focused over the last couple years driven by our customers.

The Grid is fundamentally about removing the current hard, static binding between an application and its execution environment. In the new world, this binding will be soft and very dynamic. The benefits of doing this are very significant for our customers.

Gt: Is there anything else you'd like to add about United Devices, SC'05 or the Grid market in general?

HUBBARD: From both introspection and looking at our competitors, we fully believe that Grids are moving into mainstream adoption. All indications from our current customers and prospects point in this direction.

###

GRIDToday
November 16, 2005

Features:

Altair, Accelrys Integrate for Materials Sciences Community

Altair Engineering Inc. announced a cooperative business arrangement with Accelrys Inc, a provider of computational software and services for scientific researchers. Under the arrangement, Accelrys and Altair have developed an interface that allows Accelrys' Material Studio software to operate with Altair's PBS Professional software.

Accelrys will include a free trial license of PBS Professional with its Materials Studio software on a limited-time basis, and the companies have developed a special promotional pricing program. Materials Studio is a suite of integrated materials simulation software for use in research and development in the materials, chemical, pharmaceutical and biotechnology industries. Both software programs run on Windows, Linux and UNIX platforms. The promotional pricing offer, which represents a significant discount off standard U.S. list pricing, will be in effect for a limited time beginning in December, and may be extended if response is favorable.

"Accelrys is the key ISV [independent software vendor] in the materials sciences sector, and we are pleased to enter into this business arrangement with them," said Michael Humphrey, vice president of Altair Engineering's Enterprise Computing business unit. "I am confident that Accelrys customers will see that this integrated combination provides a compelling value proposition."

PBS Professional is an open workload management solution for Grid computing environments, Linux clusters, and distributed servers and workstations. The software optimizes the utilization of computing resources by intelligently scheduling and managing computational workload in a number of industries, including manufacturing, energy, financial, government research, weather forecasting, life and materials sciences, and digital media. By increasing the efficiency of the hardware and software resources, PBS Professional reduces total cost of ownership and provides true business value to Grid computing customers.

"Integrating Altair's PBS Professional with Materials Studio allows us to provide a more complete solution to our customers," said Steve Levine, senior director of corporate development at Accelrys. "Workload management is an important component of our customers' HPC infrastructure. Providing an integrated, best-of-breed solution will better enable them reduce deployment times as well as administrative costs."

GRIDToday
November 16, 2005

Features:

Platform CEO: If Clusters are Trees, Grid is the Forest

By Derrick Harris, Editor, GRIDtoday

GRIDtoday editor Derrick Harris spoke with Songnian Zhou, co-founder and CEO of Platform Computing, about the company's role in the Grid and HPC markets -- as Platform has a strong presence in both. This is evidenced by the company's recently announced partnerships with Microsoft, IBM and HP, and Zhou goes into detail as to why each of these is important not only to the companies involved, but to the customers.

GRIDtoday: First, I'd like to discuss the three announcements Platform made this week. Can you give a brief recap of the partnerships the company struck with HP, IBMand Microsoft around cluster/HPC resource management?

SONGNIAN ZHOU: Absolutely, I'd be happy to provide a recap. At SC05 Platform made three announcements:

1. We announced an integration with Microsoft on their forthcoming Windows Cluster Compute Server.
2. We announced an expanded relationship with IBM through an integration with their industry leading Blue Gene /L supercomputer.
3. We announced that HP will now be offering Platform Rocks as part of their Linux cluster management solutions.

Gt: What makes these announcements noteworthy?

ZHOU: In the case of Microsoft there are a number of things. Firstly, there's been significant interest in what their plans are for high performance computing (HPC). By working with Platform as their partner of choice for Grid computing, Microsoft has the ability to seamlessly integrate within enterprise IT environments managed by Platform LSF. This integration will allow LSF to seamlessly schedule and pass jobs to and from Windows CCS. The relationship will evolve as we work together on a Web services standard for job scheduling. In addition, Microsoft is going after the lower end of the market, which opens up new opportunities for Platform.

In the case of IBM, we are at the opposite end of the HPC spectrum with our integration to IBM's Blue Gene /L. Powering many of the world's top supercomputers, Blue Gene /L is a powerful and scalable HPC system. Through integrating with LSF, Blue Gene users now have the ability to concurrently process multiple jobs of variable sizes, which in turn improves utilization and decreases turnaround time. In addition, Platform LSF also extends advanced workload management across heterogeneous platforms.

With HP, what we're seeing is an expansion of our relationship to incorporate Platform Rocks -- our industry-leading cluster management solution based on NPACI Rocks from SDSC. HP will now offer this solution on their Xeon-based Cluster Platform 3000.

Gt: What do these partnerships with major IT vendors mean to Platform? What do you gain?

ZHOU: These partnerships come natural to Platform. When we announced Platform Enterprise Grid Orchestrator (EGO) in August, we took a major step in expanding Grid technology adoption into enterprise IT. Mainstream enterprise is much different than HPC, and we realize to be successful we can't go it alone. \ Over the last six months, we have expanded our partner program and broadened relationships with the industry's leading IT providers such as IBM, Dell, HP, Microsoft, Intel, SAS and VMWare, with many others to come. By building and expanding these relationships, we have significantly strengthened our ability to deliver solutions to enterprise IT.

Gt: On the other hand, what do these companies gain from partnering with Platform?

ZHOU: As these companies look to move into adaptive/utility/on-demand computing solutions, Grid is a key enabling technology in the delivery of these solutions. In most cases, it's more efficient to partner with an industry leader in this area than develop a homegrown solution. Platform Computing has 13 years of experience, 1,700 large enterprise customers and a reputation as a leader and pioneer in Grid computing. These factors make us the right choice to work with in this area.

Gt: Moving away from standard HPC and toward Grid computing, I'm wondering what Platform is up to with its Grid business. Can we expect some customer announcements around Platform EGO in the near future?

ZHOU: Absolutely, 2006 is poised to be a banner year for Platform Computing. The customer interest in enterprise Grid is escalating and as we bring to market Platform EGO-based products in early 2006, we expect to make a number of new customer and partnership announcements. Our timing is spot on and Platform EGO is a unique solution that differentiates us in the marketplace.

Gt: Speaking of customers, are you seeing any encouraging, or disconcerting, trends in adoption of Grid computing?

ZHOU: Yes, we're seeing a number of encouraging signs in the adoption of Grid computing. The HPC market remains strong and continues to grow, but we see a new opportunity emerging on the horizon in the enterprise. With many large businesses enjoying healthy ROI from Grid computing, we're now seeing their enterprise IT departments getting in on the action. These initial steps are being driven by a number of

factors in the organization. In some cases, we're seeing server consolidation and the need to improve utilization as a driver. In other cases, virtualization is the jumping off point for enterprise Grid technology. Lastly, were seeing many large applications players integrate Grid technology within their products -- SAS is one such example, embedding Platform EGO technology into their recently announced Grid Manager product. All of these signs are encouraging and point to ways enterprises are beginning to leverage Grid technology within the data center.

Gt: What are your feelings on the importance of widespread standards adoption? Do you feel they are necessary in order to establish widespread adoption of Grid technologies?

ZHOU: The adoption of industry standards is key for enterprises to leverage the benefits of a heterogeneous IT environment. Platform is committed to the support of industry standards and have architected our products to take advantage of them whenever possible. In the case of Platform EGO, we will take it one step further by making a standards-based toolkit, heavily based on Web services, freely available to the developer community so they can easily customize integrations within their environment. We are also are an active participant in several of the key standards bodies for Grid.

Gt: The Platform Web site says that the Platform Globus Toolkit is built on GT3, are there plans to switch over to GT4?

ZHOU: At this point we're not seeing demand from our users to make this change. When this changes, we will look at making the switch to GT4. Platform has a longstanding relationship with the Globus consortium and will continue to support our customers that use the Globus toolkit.

Gt: I've heard on numerous occasions discussions about the differences between cluster and Grid computing. Seeing as how Platform deals in both of these areas, and seeing as how SC really brings these (along with other HPC) technologies together, can you give your two cents on the differences and similarities between clusters and Grids?

ZHOU: That's a great question and one I get regularly. A cluster is usually a group of servers inter-connected to look like a single server environment. A cluster can be large scale and heterogeneous, but is a single management domain. A Grid is when you take a group of clusters and link them together seamlessly for widespread sharing. To draw an analogy, if the CPUs and nodes were leaves, then a cluster would be a tree, whereas a Grid would be the forest.

Gt: And, speaking of SC, are you looking forward to this year's incarnation? What is it about SC that draws so many attendees from pretty much everywhere on the planet?

ZHOU: SC is a unique event in the IT industry. It's the industry's premier event in high

performance computing and really a landmark activity in assessing the state of the industry each year. At no other HPC event do we get such a great collection of customers, partners and industry experts. I am especially looking forward to this year's event. My prediction is that we will start to see more people from traditional enterprise IT in attendance this year. With HPC and Grid beginning to expand into the data center, there's a new audience that's eager to learn about the benefits of this technology and how they can leverage it.

Gt: Is there anything else you'd like to add?

ZHOU: Just that I look forward to seeing everyone at the show and encourage them to stop by the Platform (#2124) booth to say hello.

###

GRIDToday
November 16, 2005

Features:

TACC Showcases New Grid Technologies at SC'05

At SC05 in Seattle, the Texas Advanced Computing Center (TACC) at The University of Texas at Austin showcased TeraGrid technologies, newly released software, and remote and distributed visualization models that enable people to easily utilize diverse, powerful computational technologies that enhance researchers' capabilities for addressing important research problems.

TACC's mission is to enhance knowledge discovery and to improve society and the quality of life through the application of advanced computing technologies. TACC provides world-class terascale computing, visualization, and storage resources to researchers and supports them in using these resources to address the most challenging research problems. Through the National Science Foundation's TeraGrid, these resources and services are made available to the national academic research community.

In addition to providing advanced computing infrastructure and support to enable the research efforts of UT Austin and TeraGrid users, TACC conducts research and development activities to develop new computing techniques and technologies. In doing so, TACC scientists collaborate with many other researchers in these activities, including researchers at The University of Texas at Austin, other Texas universities in the High Performance Computing Across Texas Consortium (HiPCAT), at other TeraGrid institutions, and at other U.S. universities and government laboratories.
This year, TACC presented and demonstrated results from several technology R&D projects, which are listed below.

Remote and Distributed Visualization: This presentation focuses on the key advantages of a distributed visualization model, and TACC's efforts in this area. The hardware and software solutions currently available to TeraGrid users will be described, along with a preview of unique new capabilities yet to come. A live demonstration of a distributed visualization application utilizing resources located at TACC's facility in Austin, Texas, will be shown. Tuesday, 11/15, at 1:00pm. TACC Presentation Area, Booth #2244. Greg S. Johnson, TACC.

TeraGrid User Portal: The TeraGrid User Portal will integrate important user capabilities in a single web interface including allocation and account info, user documentation, consulting and training info, comprehensive RP information services, and simple access to GIG and RP interactive grid usage. In addition, the User Portal will offer users personalization and customization capabilities that allow them to see only information that is most relevant. This talk will discuss current progress and future plans of the TeraGrid User Portal. Tuesday, 11/15, at 10:00am. TeraGrid/ANL Booth #1629. Eric Roberts, TACC.

The GridPort Toolkit 4: The GridPort Toolkit (GridPort) enables the rapid development of highly functional grid portals that simplify the use of underlying grid services for the end-user. This presentation will discuss the motivation, architecture, features, and projects that are using GridPort 4 as well as give a demonstration of a GridPort portal interface. Tuesday, 11/15, at 11:00am; Wednesday, 11/16 at 11:00am; Thursday, 11/17 at 11:00am. TACC Presentation Area, Booth #2244. Eric Roberts, TACC; Maytal Dahan TACC.

Matrix Multiply: This presentation provides an introduction to an important linear algebra operator, the matrix multiply. On many systems, the matrix multiply performs at over 95% of peak and this discussion provides an overview of how an optimized matrix multiply is implemented on a variety of architectures.

Using GISolve and GridShell to Demonstrate the Interoperability between The Open Science Grid and NSF TeraGrid: The presentation will demonstrate a submission pathway for OSG jobs to TeraGrid managed clusters. The demonstration will create a personal virtual cluster composed of OSG/TeraGrid resources using GridShell. Jobs using the GISolve software, developed at The University of Iowa, will then be submitted and managed through the system across OSG and TeraGrid.

Growing & Managing Complex Grid Environments: Numerous organizations across the world have implemented production Grids. In this discussion, United Devices will discuss a number of cases covering a wide range of use cases and organizational environments. Meta-scheduling, cluster management, and capacity management will be prominent topics.

Increasing Linpack Performance: The High Performance Linpack (HPL) program is an important application used to benchmark new architectures and large-scale system installations. Although single CPU performance of the Linpack benchmark can be over 90% of peak, there are still some bottleneck functions. This presentation outlines the problematic functions and presents suggestions for increasing the peak Linpack performance.

GridChem: The Computational Chemistry Grid project (GridChem) aims to help the computational chemistry research community run jobs efficiently on distributed computing/storage resources and increase productivity by providing a homogeneous quantum chemistry desktop environment for building quantum chemistry jobs. GridChem is a three-tiered architecture, consisting of a desktop client, middleware server and a set of compute/data resources. This presentation will include a demonstration of the client and a discussion of the middleware components. Thursday at 2 p.m. TACC Presentation Area, Booth #2244. Chona Guiang, TACC; Rion Dooley, LSU; Sudhakar Pamidighantam, NCSA.

###

GRIDToday
November 16, 2005

Features:

Internet2 Demos Optical Networking Firsts at SC'05

For the first time, three radio telescopes distributed around the world will be connected via dynamically provisioned dedicated optical circuits for an electronic Very-Long-Baseline Interferometry (e-VLBI) observation. Internet2 announced this scientific and networking achievement at the first major demonstration of its nationwide Hybrid Optical and Packet Infrastructure (HOPI) testbed, during the SC.05 conference held in Seattle. The demonstration marks a critical milestone in dynamic or "on demand" optical networking that can support even the most extreme applications used by the global research and education community today.

Leveraging the HOPI infrastructure together with the NSF-funded DRAGON testbed, the telescopes located in Westford, Ma., Greenbelt, Md., and Onsala, Sweden, will be dynamically linked via dedicated low-latency optical circuits to a central data correlator and simultaneously transmit multiple gigabits-per-second of data during a 20-minute observation. The team will also attempt to connect to a fourth telescope in Kashima, Japan, during the demonstration. Historically, radio astronomy data was recorded on magnetic tape or disk at each site and shipped to the central processing location for the analysis.

"VLBI is one of the most powerful techniques available for the high- resolution imaging of distant radio sources in the universe and for making accurate measurements of the motion of the earth in space," said Alan Whitney, principal scientist at the MIT Haystack Observatory in Westford. "These capabilities also allow scientists to measure such things as continental drift and to calibrate the orbits of GPS satellites to enable more accurate position measurements on the surface on the Earth. Advanced optical networks like HOPI and DRAGON, will undoubtedly open new doors for radio-astronomy observations and important science."

As a part of Internet2's mission to design and deliver an advanced network infrastructure to meet the emerging needs of the research and education community, Internet2 has built the HOPI nationwide testbed to investigate next-generation network architectures that combine the best qualities of optical and packet technologies. The testbed is a model for the future of Internet2's high performance Abilene network which serves as a platform for both experimental networking applications as well as stable production IP services.

"The HOPI testbed has far-reaching applications in the scientific, engineering, and medical arenas which have come to require far more sophisticated network and resources than those previously available," said Rick Summerhill, co-chair of the HOPI design team and Internet2 director of network research, architecture and technologies.

"Internet2's HOPI investigation represents a new paradigm in networking that goes well beyond traditional production services of today. In doing so, we hope to catalyze a new era of advanced applications which at this point have only been imagined."
Not only will the demonstration highlight the capability to provision on-demand light paths within an administrative domain, but it also proves for the first time, the ability to provision those optical circuits across multiple network administrative domains for global data transmissions. Utilizing DRAGON-developed inter-domain Generalized Multiprotocol Label Switching (GMPLS) capabilities, which provides control plane capabilities, automated end-to-end circuit provisioning, and management of network resources, the optical routes were seamlessly connected across scientific, HOPI and DRAGON domains. The paths also crossed UKLight, SURFnet, NorthernLight, Nordunet, SUnet, JGN2, StarlLight, GIG-EF and BOSnet.

"We believe the control plane technologies DRAGON has developed and integrated into HOPI pulls together a number of efforts within the R&E community and the international Internet standards bodies to show that these dynamic hybrid network architectures are indeed viable and of great value to the scientific and academic communities," said Jerry Sobieski, lead coordinator of the HOPI Testbed Support Center and project manager for the DRAGON Project. "This demonstration opens the door for both significant advances in radio astronomy and geodesy as well as establishes a foundation on which the global networking community can expand the scope and availability of these capabilities."

Designed to model future optical networking infrastructures, the HOPI testbed utilizes facilities from Internet2's Abilene Network, the National LambdaRail (NLR) infrastructure, The MAN LAN exchange point, and regional optical networks. HOPI nodes, deployed in Seattle, Chicago, New York City, Los Angeles and Washington, D.C., each consists of a 10-Gigabit Ethernet switch provided by Force10 Networks, a fiber cross connect switch provided by Glimmerglass, and HP's servers to measure network performance and manage control plane capabilities.

###

GRIDToday

November 17, 2005

Features:
Catching Up with Walter Stewart

By Derrick Harris, Editor, GRIDtoday

Since Walter Stewart's exit from SGI in September, the company's former business development manager for Grid has incorporated his own cosulting firm and already has two large contracts, including one with CANARIE. *GRIDtoday* editor Derrick Harris spoke with Stewart to find out more about what he's up to and to get his views on the future of Grid computing and "Intelligent Infrastructure."

GRIDtoday: When did you officially leave SGI and what brought about that career change?

WALTER STEWART: SGI began a major restructuring and down-sizing on Sept. 1, 2005. Like many others whose work was less connected with quarterly sales and more connected with future direction, my position was eliminated in an attempt to match costs to revenues. I had eight wonderful years at SGI. I think the company has magnificent technology. I sincerely hope SGI is able to complete the turnaround and return to profitability. I am also grateful for the "push" toward exploring new activities.

Gt: What are you doing now? Can you tell me a little more about your consulting work?

STEWART: I have incorporated "Walter Stewart & Associates." I have two large contracts at present. I am under contract to CANARIE -- Canada's high bandwidth network for research -- to act as senior advisor on Intelligent Infrastructure (II). "Intelligent Infrastructure" is the term in Canada for what in other jurisdictions is referred to as cyberinfrastructure or e-infrastructure. CANARIE is eager to stimulate discussion and planning of the deployment of intelligent infrastructure. CANARIE is focused on using Web services and SOA together with its own developed User Controlled Lightpaths to bring together many existing elements of II into an integrated, highly functional facility to serve research and the wider economy. My role is to engage many in government, universities, funding agencies, not-for-profit associations and the private sector in the discussion with a view to develop a pan-Canadian strategy.

My second contract is to advise a major global corporation, which prefers to have its identity kept confidential at this time, on more effectively working with the research community. The corporation is high-tech, but it is not a computer company.

I am also having conversations with a number of other organizations both within Canada and abroad about other possibilities which generally fall under the headings –

technology, organization and strategic planning.

Both before SGI and while at SGI, I spoke frequently at conferences around the world on aspects of the digital economy. I may well ramp up my work as a speaker.

Gt: Do you have plans to re-enter the vendor space, or will you stick with consulting for the foreseeable future?

STEWART: I am open to possibilities for re-entering "the vendor space," but certainly not wedded to doing so. In a sense, with the second of my two consulting contracts, I am at least partially in "the vendor space." I confess that I am enjoying not having my advice "tainted" with any sense of "he's just here to sell a box."

Gt: You recently received a leadership award from the GGF for your participation in the plenary program committees. What did that mean to you?

STEWART: I was "gob smacked." It meant a great deal to me to have my work recognized by such an organization of highly talented and highly dedicated people. I am a concept person, no one would ever ask me to design a Grid or program some middleware. I so value the work of the gifted people who are capable of doing the technical, hard slog. At the same time, I think we often get overwhelmed in the details -- in the "how" so to speak -- with very little reflection on the "what" and the "why." I consider that reflection my contribution. It was extremely gratifying to have it recognized.

Gt: Your most recent activity with GGF was as chair of the Enterprise Program committee at GridWorld. Overall, how do you think that event went, and how do you think it will evolve in the years to come?

STEWART: I think the program went extremely well. The thing that struck me most about the Grid World Enterprise Program was just how inconceivable it would have been even as recently as one or two years ago. The Enterprise Program this year was almost completely occupied with hearing from people who'd built and now used Grids in the enterprise. Two years ago it was still a dream. This year's Grid World made it clear that it is a dream that is being realized every day as more and more organizations turn to Grid computing to satisfy their demands for infrastructure for data. I think the enterprise program will only strengthen as more and more organizations engage in exploring, adopting and deploying Grid computing. This year we had some significant focus on the challenges of applications on the Grid. I think that is going to remain a major issue for the foreseeable future. The Grid is a disruptive technology that invalidates some existing business models. Developing the new models is a challenge yet to be met.

Gt: Speaking of evolution, how do you think Grid computing has evolved since you first attended a GGF event? How do foresee it evolving in the years to come?

STEWART: I attended my first GGF event in Toronto in 2001. To many at that time, Grid computing was an interesting computer science project. Particularly in North America it was very focused on harvesting unused cycles. I think in five years it has truly expanded to become a vision of infrastructure for a knowledge-based economy. Any economic system requires a distribution system to gather rather materials, process them, and fabricate them into a product. In a knowledge based economy, data are the raw materials we need an infrastructure to ingest, store and manage, process, disseminate, reuse, re-create, etc., data. Grid computing, particularly with its current emphasis on SOA and Web services, becomes a critical element of that infrastructure.

Gt: I've heard a fair amount of talk lately regarding the need for the Grid community to adopt some quality standards, and to do so as soon as possible. Do you share this sense of urgency? What roles do you think the various vendors, standards bodies and consortia have to play in this process?

STEWART: I think the standards work of organizations like GGF is extremely important. We cannot build real, seamless infrastructure for a knowledge-based economy without agreed standards. I don't in anyway gainsay the urgency of the work, but I also would wish to emphasize how important I think it is that these standards are being developed in community in an open process. At times, the process may compromise the ability to meet the urgency, but it is critical that we preserve the open process. We do want proprietary Grids. I believe it is the role of all -- vendors, standards bodies, consortia, the academy -- to bring the greatest possible urgency to the open process and to engage in the process as fully as possible.

Gt: What do you think are some of the biggest misconceptions about Grid computing (and what are the realities), and what can the community do to address them?

STEWART: In some work I did while at SGI with Platform Computing, we agree that Grids are built, not bought. I think in a very consumer-oriented society, we tend to think we should be able to get everything off the shelf. That is simply not possible when you are building real infrastructure. It requires careful planning and integration. It must be seen as iterative. There is a great separation between the vision for Grid computing or Intelligent Infrastructure and what is currently available. Some think the gap is a problem. I rejoice in it -- thanks be to God. I hope we always are deployed in a way that lags behind the vision. We have seen extraordinary progress in well under 10 years. I believe we will see even more progress in the next few years. The "vision lag" is critical to that progress. Someone who had gone to sleep in 2001 in the heyday of cycle scavenging as Grid computing would awaken today in a world near unrecognizable. Even what can be deployed today with SOA and Web services is near miraculous from a 2001 perspective. The promise of what will be achievable is all the greater. We absolutely need to ensure that the vision develops as fast and remains as far ahead of the actually deployed.

Gt: Moving on to SC'05 ... will you be speaking or presenting, or are you just taking in the experience and chatting with colleagues?

STEWART: I am attending SC '05 with my client, CANARIE. CANARIE is part of hugely expanded effort at SC 2005 of Canada's high-performance computing consortia. I am looking forward to not having booth duty and having a chance to really explore developments connected with Intelligent Infrastructure.
Gt: How many of these shows have you been to over the years? How will it be different to attend as a consultant versus attending as a representative for a major vendor?

STEWART: I think this is only my fifth Supercomputing -- as opposed to the 12 Grid Forum's I have attended. I am looking forward to attending as an independent. I never felt overly burdened by being at one level or another "a shill" for technology because I believed in the quality of the technical solutions we had on offer. Nevertheless, I value the chance to walk the floor completely focused on the larger pictures of what are we building for the knowledge-based economy and why are we building it.

Gt: Is there anything else you'd like to add?

STEWART: I regret that other commitments prevent my being at SC 2005 beyond mid-day on Wednesday. I hope that I have the opportunity to see as many people as possible. I feel very privileged to be part of the community who are building 21st century infrastructure.

###

GRIDToday
November 17, 2005

Features:

Grid Bridges 4,800 Miles for Molecular Repositories

In a bid to facilitate collaboration among other biomolecular researchers, the Department of Energy's Pacific Northwest National Laboratory has become the first institution outside the United Kingdom to join the Biological Simulation Grid Consortium of Great Britain.

The BioSimGrid was organized to support research at the universities of Oxford, Southampton, Bristol, Birkbeck, Nottingham and York. Researchers are seeking to learn more about some of the most fundamental building blocks of life. According to Doug Ray, chief research officer at PNNL, collaboration with researchers at these institutions, which rank among the world leaders in biological science, positions PNNL to contribute to breakthroughs in scientific understanding that could have tremendous impact on the environment, energy, security and everyday life.

Molecules may be small, but the data sets resulting from computer simulations of their behavior at atom-level resolution are huge, ranging in size from gigabytes to terabytes. The extreme file sizes mean simulation data generally resides in the home laboratory where the research was conducted and, for practical purposes, remains inaccessible to other research groups.

Yet, sharing simulation data for comparative analysis is a major goal of biomolecular scientists whose work aims to solve problems in such diverse fields as medicine, environment, national security and energy.

"There is an immediate scientific need for the BioSimGrid," affirmed T. P. Straatsma, lead researcher at PNNL. "It will enable molecular scientists to collaborate in a far more sophisticated way than we've been able to ever before. The ramifications of sharing terascale data are huge."

While there have been attempts to launch similar Grid networks in the United States, none are currently active, Straatsma said.

The PNNL node on the BioSimGrid coincided with the lighting of a new high-speed, high-capacity fiber optic connection between PNNL's campus in Richland, Wash., and Seattle. This connection made it possible for the laboratory to accept the BioSimGrid Consortium's invitation to deploy the network infrastructure and software framework that comprise the Grid.

GRIDToday
November 17, 2005

Features:

Gridbus Project Releases Grid Service Broker v2.4

The Grid Service Broker, developed as part of the Gridbus Project at the University of Melbourne, Australia, mediates access to distributed "autonomous" resources by: (a) discovering suitable data sources for a given analysis scenario; (b) suitable computational resources; (c) optimally mapping analysis jobs to resources; (d) deploying and monitoring job execution on selected resources; (e) accessing data from local or remote data source during job execution; and (f) collating and presenting results. The broker provides (a) a declarative and dynamic parametric programming model and (b) a rich set of Java APIs for creating Grid applications. It also allows programmers to create a user-level scheduler plugins that can replace default Computational and Data Grid scheduling algorithms that support service price-based resource allocation.

The Gridbus Broker v2.4 supports/uses the following middleware:

- Globus 2.4.x,Globus 3.2, Globus 4.0.
- Alchemi 1.0.
- Condor 6.6.9.
- OpenPBS 2.3.
- SGE.
- NWS 2.8.
- SRB(Storage Resource Broker) 3.x.
- XGrid Technical Preview 2 & 1.0 (experimental).
- Unicore 4.1.0 (experimental).

The Gridbus Broker v2.4 can utilize a resource via SSH for submitting and executing Grid jobs on: OpenPBS 2.3,SGE and Fork (on Unix-class systems).

What's new:

- Support for Globus 4.0 job submission services using WS_GRAM.
- Improved stability for all middleware.
- Improved persistence in order to support a wider range of databases. Also, introduced in-process and dynamic database creation via HSQLDB.
- New event-based Round-Robin scheduler that provides a simple scheduling algorithm.
- Various bug fixes.

The broker source code,binaries, documentation and manual can be downloaded from www.gridbus.org/broker/.

For further details about the Gridbus broker and/or the Gridbus project, please contact:

- Rajkumar Buyya (raj@csse.unimelb.edu.au).
- Srikumar Venugopal (srikumar@csse.unimelb.edu.au).
- Krishna Nadiminti (kna@csse.unimelb.edu.au).
- Hussein Gibbins (hag@csse.unimelb.edu.au).
- Tianchi Ma (tcma@csse.unimelb.edu.au).
- Marcos Assunção (marcosd@csse.unimelb.edu.au).

The Gridbus broker and the GUI is distributed under the GNU General Public License (GPL) (http://www.gnu.org/licenses/gpl.txt). Other libraries included in the distribution are distributed under their own respective licenses, which are also included.

This product includes software developed by and/or derived from the Globus project (www.globus.org/). This product includes dom4j libraries (www.dom4j.org). Other licenses can be found in the gridbus2.4/licenses directory of the Gridbus Broker v2.4 distribution.

The Gridbus Broker is developed by the Gridbus Project, Grid Computing and Distributed Systems (GRIDS) Lab., Dept. of Computer Science and Software Engineering, the University of Melbourne, Australia. The project is partially supported by Australian Research Council Discovery Project grant, and the University of Melbourne.

###

GRIDToday
November 17, 2005

Features:

Absoft Announces HPC Development Kits for Parallel Apps

Absoft Corp. announced at SC'05 in Seattle new additions to its family of High Performance Software Development Kits. The Absoft High Performance Computing Software Development Kits (HPC SDK) are complete development solutions that allow software developers to compile, run, debug and optimize high performance software applications on clusters, Grids and multiprocessor/multi-core systems built using 64-bit Intel Xeon and IBM POWER processors. The Absoft SDKs are the only HPC software development toolkits available today, and are offered exclusively by Absoft for systems running Linux, OS X, and soon, for Microsoft Windows Compute Cluster Solution (CCS).

Each HPC SDK includes software development tools optimized for 64-bit extensions and is preconfigured for easy installation. Every SDK includes Fortran and C/C++ compilers, debuggers, math libraries, pre-built message passing (MPI) libraries and other development tools. Also included are a comprehensive set of example HPC programs to help new users quickly learn how to take full advantage of the HPC SDK features on their cluster, Grid, multiprocessor or multi-core system. Every HPC SDK comes with 12 months of continuing maintenance for all included components. The HPC SDK from Absoft combines tools from multiple vendors to provide customers with a best-of-class solution.

"We have been using your SDK for just over 6 months now and we have never had a better suite of compilers and tools. The HPC SDK on POWER delivers ease of installation and usage, high performance, flexibility, and reliability that is unsurpassed on the Linux ppc64 platform. Using the HPC SDK Enhanced Edition with its innovative and comprehensive compiler, debugger and development technologies has allowed researchers at the University of Oregon's Neuroinformatics Center to improve their application performance and has streamlined our development process dramatically. We have used the SDK in neuroscience software development, including EEG signal processing and computational head modeling. In addition, we are applying the suite in our parallel performance tools research and development as part of the TAU Performance System project. In our experience, the HPC SDK Enhanced Edition for IBM Linux on POWER has required the lowest maintenance overhead combined with the best feature set for our ppc64 Linux development platforms. The compilers are robust, the parallel libraries are well integrated," said Allen Malony, director of the Neuroinformatics Center at the University of Oregon.

Absoft announced the following new HPC SDK versions along with additional high performance tools:

- HPC SDK Basic Edition 2.0 for 64-bit Intel Xeon.

- HPC SDK Enhanced Edition 1.0 for 64-bit Intel Xeon.
- Intel VTune Performance Analyzer.
- Intel Thread Checker.
- HPC SDK Basic Edition 2.0 for IBM Power.
- HPC SDK Enhanced Edition 1.0 for IBM Power.
- IBM High Performance Computing Toolkit for IBM Power.
- IBM FDPR-Pro post-compilation optimization tool for feedback-directed program restructuring.

"Absoft continues to expand our comprehensive, exclusive line of SDK solutions for high performance software application developers. As seen by our customers' testimonials, the Absoft HPC SDKs are innovative toolkits which provide excellent value," said Jeff Livesay, Absoft chief operating officer. "We are experiencing increasing traction in this exciting growth market and have many new powerful tools now shipping or soon to be available. Absoft is listening to HPC customers' expression of their needs and providing the complete, commercially-supported solution bundles to meet those needs."

The HPC SDK Basic Edition 2.0 for 64-bit Intel Xeon processor-based servers and clusters running is updated with compilers and libraries from Intel including C++ and Fortran 9.0 compilers, Intel MPI, and Intel Cluster Math Kernel Library 8.0. This HPC SDK also includes updated LAM/MPI and new Absoft examples and benchmarks. The HPC SDK Basic Edition 2.0 for 64-bit Intel Xeon processor-based servers and clusters running Linux fully supports the recently announced IBM Grid and Grow hardware and services offering. This Grid computing support makes it easier for companies of all sizes to build and deploy affordable and efficient Grid computing IT operations.

"The IBM Grid and Grow solution based on the 64-bit Intel Xeon processor provides end users a cost-effective foundation for a low-risk quick-start deployment of Grid technology. And, the close collaboration between IBM, Intel Corp. and Absoft resulted in a 64-bit Intel Xeon processor-based HPC SDK, which makes it easy for companies, small and large to use the Grid and Grow foundation to build applications that are solid and high performing," said Robert Fogel, director of worldwide Grid strategy and business development at Intel. "The combination of the Grid and Grow foundation plus the applications developed with the Absoft SDK delivers an innovative and scalable IT solution that can deliver high-impact business value."

The HPC SDK Enhanced Edition 1.0 for 64-bit Intel Xeon processor-based servers and clusters running Linux includes all of the features of the popular HPC SDK for Intel Xeon systems with additional tools from Intel, including, VTune Performance Analyzer to easily streamline code and Thread Checker for OpenMP debugging support. Also included are open source tools such as Marmot, an MPI diagnostic tool. Future versions of the Enhanced Edition of the HPC SDK will feature the FxP parallel MPI debugger

from Absoft.

The updated HPC SDK Basic Edition 2.0 for Linux on IBM Power processor-based servers and clusters running Linux receives many additions, such as IBM's XL C/C++ V8.0, IBM's XL Fortran V10.1 compilers, and the Engineering Scientific Subroutine Library (ESSL). This HPC SDK also includes updated LAM/MPI and MPICH2, and new Absoft examples and benchmarks.

Introduced at the inaugural GridWorld Conference in Boston, this HPC SDK has been specifically designed to work with the IBM Grid and Grow hardware and services offering. IBM Grid and Grow is a powerful solution designed to help organizations get started quickly with affordable Grid computing.

The HPC SDK Enhanced Edition 1.0 for Linux on IBM Power processor-based servers and clusters running Linux includes all of the features of the Basic Edition with the addition of powerful tools from IBM, including the IBM FDPR-Pro tool and the IBM High Performance Computing Toolkit. See the following sections for descriptions of these tools.

The FDPR-Pro tool, developed by the IBM Haifa Research Lab, uses Feedback Directed Program Restructuring to achieve performance gains of anywhere from 5 percent, 10 percent or more for software applications that have already been compiled and linked. FDPR-Pro for Linux on Power is a post-link utility for improving the performance applications that were compiled for the Linux operating system running on Power family platforms.

The tool receives input files in XCOFF format, instruments them, executes them for profiling information, and then optimizes them to achieve:

- Better hit/miss i-cache ratio.
- Reduced number of branches.
- Reduced number of TLB misses.
- Reduced number of page-faults.

The FDPR-Pro tool for Linux on IBM Power handles very large executables (such as an Oracle database application) and shared libraries.

The FDPR-Pro application optimization tool is included in the HPC SDK Enhanced Edition 1.0 for IBM Linux on Power and is also offered as a standalone product. FDPR-Pro will be generally available exclusively from Absoft in the first quarter of 2006.

The IBM High Performance Computing Toolkit (HPCT), developed at the Advanced Computing Technology Center (ACTC), part of IBM's T.J. Watson Research Laboratory

is a suite of performance-related tools and libraries to assist in application tuning. This toolkit is an integrated environment for performance analysis of sequential and parallel applications using the MPI and OpenMP paradigms. It provides a common framework for a number of IBM systems offerings, including IBM eServer pSeries systems. The IBM HPCT includes tools for:

- Hardware Performance.
- -hared Memory Performance.
- Message-Passing Performance.
- Performance Visualization.

"We use the IBM HPC Toolkit at NERSC, and find it to be invaluable in understanding the performance behavior of our applications on our IBM pSeries systems" according to David Skinner, HPC senior applications specialist at NERSC. "We especially like the Toolkit's ability to trace back the performance data to the actual source code statements of our applications in the highly intuitive and visual framework."

###

HPC Wire
11/11/2005

Features:

Kramer Spearheads Impressive SC05 Showing

William T.C. Kramer must be using his HPC prowess to somehow get 28 hours of each day. It's the only way to explain how he could maintain a schedule that includes chairing this year's Supercomputing extravaganza, while serving as deputy division director at the National Energy Research Scientific Computing Center and pursuing his doctorate while also remaining a dutiful and doting husband and father.
It was not surprising that SC05 was first on his mind during a recent conversation from his NERSC office at Lawrence Berkeley Laboratory. After all, this year's conference and trade show is poised to be the biggest in recent years, measuring from many different aspects.

Perhaps most foremost, SC05 has sold out every available inch of allocated exhibit area. According to Kramer, some 220 exhibitors will be occupying "every space the fire marshal would allow us to sell." The booths range in square footage from 10-by-10s to 50-by-50s.

"The industry as a whole is doing well," said Kramer, reflecting on the reasons for the encouraging vendor support. "We're seeing more people and different people." And not just on the vendor side. In particular, he added there has been lots of interest from government initiatives.

This interest has helped Kramer expand the focus of the conference, which should be particularly noticeable to habitual attendees. There is the biggest educational program ever, highlighted by the expected attendance of dozens and dozens of teachers from high schools and colleges. Much of this increase is thanks to additional funding. "There is a renewed energy regarding HPC," said Kramer. "This relates also to how HPC is being applied in a multiple of new and existing areas of study."

Attendance is running ahead of last year's pace. This is particularly interesting because it was only the late 1990s when support of SC was rumored to have peaked and whispers began questioning its relevancy. Hopefully, 2005 will let SC wipe off its Rodney Dangerfield tag altogether. After all, when the gala bounced back with 2002's showing in Baltimore, critics attributed its proximity to the nation's capital. The word went out the next year that there was little interest in Phoenix - and SC03 drew the biggest crowd ever. But the result only led to naysayers predicting a deep drop for Pittsburgh last year. The result: another attendance record.

This kind of pessimism doesn't seem to exist for Seattle's event. The No. 1 reason could

be the opening keynote speaker that Kramer secured: a guy named Sir William H. Gates, Microsoft Corp's chairman and chief software architect.

Kramer admitted his masterstroke of scoring Gates as SC05's opening speaker represents two years of working various channels at Microsoft to get an audience with the man at the top. Once he did, then it still didn't end. Then there was crafting the rational for the HPC faithful as to why the software giant—not even considered a legitimate HPC player by many in the field -- should score the cherished kickoff spot that has traditionally been the roost of non-vendor dignitaries.

After all, holding the event in proximity to Microsoft's Redmond, Wash. headquarters didn't hurt SC's chances. Yet Kramer acknowledged he did not reel in Gates without a Herculean effort from himself and others counted among the SC05 inner circle. "Gates actually canceled a trip to China to be at SC," Kramer said. "He's certainly interested in HPC. He's been getting briefed twice a week to get up to speed on all aspects. Obviously, this is a significant sign."

One thing's for sure: the HPC faithful have been giving Kramer their opinions on Gates' high-profile selection for SC05. "People love him, people hate him," Kramer chuckled. "But even if they're somewhere in-between, it looks like they will come see him speak." According to Kramer, those attendees that do should know Gates would not be blowing his marketing bugle. "He's going to do a technical talk," Kramer explained. "Sure, he'll bring in Microsoft's interest in HPC, but he [Microsoft] is doing an HPC product announcement in this timeframe." A lot of the keynote also will be based around Microsoft's research efforts in the field of HPC, he added.

SC05 is bolstering its bastions for Gates' speech, having more than 3,000 seats available in addition to two overflow areas. If getting Gates to speak is considered a coup, imagine the extra credit for getting him to answer some unscripted questions in an open question-and-answer forum after the speech. This looks like a reality. "We're very happy with the set-up," Kramer said. "It should make for an interesting discussion."

Kramer also is excited about the debut of what is called SC Desktop, an experiment that will extend some of the conference proceedings to people who could not typically attend the event. "This concept may really change the way the event runs," he said. For example, graduate students who cannot make the trek to Seattle due to budgetary, time or visa restrictions will be able to view some of the happenings broadcast to their computers. SC Desktop takes in Gates' keynote, SC's technical content and one exhibitor forum. Gates' speech will also be available to all interested parties through Microsoft's pressroom on its site.

SC Desktop uses Access Grid technology to connect off-site attendees anywhere in the world to specific parts of the conference programs via an appropriate network

connection. Participants then will have access to speaker plenary sessions, the Cray and Fernbach Award talks and four of the half-dozen parallel technology program tracks. The tracks include two sessions of SC Papers, one Masterworks presentation and one Exhibitor Forum.
Participants can choose to view the SC Global Showcase, which includes a keynote address by Rick Stevens, father of Access Grid, along with presentations on collaborative art, Grid technology, high-resolution visualization tools and collaborative group development.

Aside from watching the goings-on from their desktops, participants can ask questions via email. Kramer again explained SC Desktop is "an experiment for the conference," meaning he "does expect a tremendous amount," but that said, he added, "It's exciting to get more HPC content and information to a broader audience, especially students." Because, at the end of the day -- even long ones like Kramer's -- HPC's future is hinged on getting students and technology leaders interested in the field. "We need to continue growing the awareness," he said. "And for a modest fee, SC Desktop gets the word out." Look for the effort to expand in 2006, he said, with perhaps a virtual exhibit area.

As another way to spread the HPC gospel to the masses, Kramer said SC05 is undertaking the filming of a video. Coming on the heels of the Council on Competitiveness' outstanding DVD, expect SC's version of the HPC story to be more Ken Burns than the penguins from "Madagascar," who starred in the council's movie.
The effort, which will detail why HPC machines are a necessity in our lives, will be distributed at SC. Unlike the council's effort, Kramer, added, this one is available at no cost.

Look for Microsoft to have a huge presence on the show floor, Kramer said. The company specifically targeted the exhibit space on the bridge between the conference's two exhibit halls. According to Kramer, the software giant was treated just like any other exhibitor. In fact, senior HPC stalwarts Cray Inc., IBM and Intel Corp. all passed on the exhibit territory that Microsoft cherished from the start.

Kramer points to SC05's HPC Analytics program as a path for bringing in new users to the event. The increased interest in storage issues, which started last year, also continues to go forward, he added. According to Kramer, computing, networking and storage are the three "stool legs" upon which the conference program is built.

He also credits this year's HPC Analytics Challenge for expanding the legions of supercomputing users who will be on hand. Most notably, Kramer added, HPC heavyweight Boeing will have a bigger presence at SC this year, in part due to its proximity to the event, but other big-name HPC users also have promised stepped-up participation.

The challenge represents an opportunity for industry, academia, government and other organizations to develop and demonstrate applications highlighting analytics techniques for solving complex, real-world problems. Submissions from around the world are under consideration, including ones focused on earthquake research, predictive modeling and traffic analytics.

According to Kramer, the rise in performance and overall computational power in HPC coupled with the rising amount of data collected have resulted in a rising interest in HPC Analytics. At the same time, the field, once the exclusive domain of government and research supercomputing communities, now has generated noteworthy interest from a range of business sectors.

All signs point to surpassing 2004's attendance. One key indicator, hotel reservations, are ahead of last year's tally to the point where additional hotels had to be secured for attendees, he said.

Meanwhile, at NERSC's Berkeley Lab, Kramer's "other" job, the facility's new Linux Networx Inc. supercomputer has been in production since August. According to Kramer, users "like it very much" and it has been kept busy with science production tasks. There is also another new machine currently under scrutiny at NERSC, Kramer said. While it represents the lab's biggest and best-performing machine ever, it has not gained full acceptance yet, he explained.

While he categorized his demanding SC05 workload as "manageable" while the responsibilities at Berkeley Lab seem numerous, he admitted there have not always been enough hours in the day for his Ph.D. studies. He is pursuing his doctorate in computer science at the University of California at Berkeley. "I'm a semester behind now," he said. "But my advisors have been supportive and understanding." Nevertheless, his studies have been relegated to late evening and early morning hours after his young daughter has gone to bed.

While he expects her to be in attendance this month in Seattle, he remarked she isn't expected to carry on the Kramer SC legacy until 2025.

###

HPC Wire
November 11, 2005

Features:

Linux Networx's Ewald: 'Clusters are Here to Stay'

While this time of the year has many people thinking about getting ready to go home for the holidays, Linux Networx Inc. CEO Robert H. "Bo" Ewald already is a jump ahead. He sees next week's SC05 gathering in Seattle as an opportunity to see old friends and members of the supercomputing family weeks before the traditional end-of-year festivities.

"I do feel like I'm going home for the holidays," said Ewald. "SC is a great time to catch up with lots of familiar faces." The trip also permits him to revisit some familiar turf, as he worked at then Minneapolis-based Cray Research Inc. from 1984 to 1996, capped by a two-year stint as president and chief operating officer. Those were some high times at Cray, as the company quadrupled in size to more than $900 million.
And as with any traveler heading home, Ewald will come to SC with a satchel stuffed with goodies and "family" news, more specifically exciting product and company announcements.

Most notably, Linux Networx will use the SC05 spotlight to announce that it has transformed its clusters into a new product series with two distinct families. These families, targeted respectively at high-end and mid-range customers, aim to propel cluster computing out of the perception of being a thrown-together assemblage of random hardware and software into the reality of today's clusters being pre-packaged solutions with tested applications and tight software-hardware integration.

"When you buy a car, you don't start by saying what type of fan belt, tires, etc. you want," Ewald said, who came onboard with the company in June. "Instead you choose the car and then work through the list of options." Linux Networx is following this analogy for its clustered solutions. First customers select the model, then the company works to build out an appropriate list of tightly integrated extras.

According to Ewald, while these two product lines, which will be officially introduced Monday at SC05, should meet the demands of the vast amount of users, the company will continue to build specialized, unique solutions for users - mainly governmental types - that require advanced technology clusters.

The company's new vision received rousing approval when unveiled a few months ago at Linux Networx's first user group meeting. This too must have felt like old times to Ewald, as earlier in his career he was involved in Cray's first meetings of this type when he worked at Los Alamos National Laboratory and then later again as a Cray executive. "I really believe in the user group concept," said Ewald. "I believe in the value of

regularly scheduled meetings to get customers and the company to talk, share ideas, exchange thoughts about what each side needs and set about ways of getting there." Linux Networx's first user meeting, held in September, balanced input on short-term fixes with measured responses to plans about long-term goals. Ewald characterized the meeting's tone as "open and candid."

He took away from the user group one definite theme that he will be expounding at SC05: "Clusters are here to stay."

That said, the dominant user chorus called out for clusters -- or at least what the concept has been perceived of so far -- to move beyond its first generation to what's ahead. According to Ewald, first-generation clusters have been typecast as collections of hardware and software from different places that were pieced together requiring a lot of administrative time. The new cluster reality, he added, must take the form of standardized platforms with integrated hardware and software running applications that have been pre-tested to meet the exceptional performance and price parameters that users expect from Linux-based machines.

"And don't forget: users want the kind of support they are used to receiving from legacy supercomputing vendors," Ewald said. "Users also are clear they want standardization -- across hardware, software and support."

Linux Networx sees its message resounding well in the government market, where the company continues to make serious efforts with national laboratories around the world. Next popular is industry, where Linux Networx has, in particular, targeted aerospace, automotive and heavy industry, Ewald said. A lot of the growth in the next few years is expected to come from defense and intelligence efforts, he added, as well as oil and gas and universities.

A little bit further out, Ewald said he envisions the financial industry becoming a strong user of clustered supercomputing. Linux Networx is in the early stages of working with one large financial concern, Ewald said, and once the vertical market as a whole recognizes cluster vendors can provide necessities such as strong systems support, he predicted more wins will accumulate rapidly.

In the meantime, starting at SC05, Ewald said Linux Networx would continue to emphasize its vision of "turning clusters into systems." The breadth of this possibility will be visible especially at the company's booth on the SC05 exhibit floor, where demonstrations of the two new product families will be displayed prominently. Special focus also will be put on a visualization demonstration of cluster computing in a real-world application, he said.

Overall, look for Linux Networx to be sporting "a bolder, more aggressive look and feel"

at SC05. While tops on that list are the product launches, Ewald said the bigger picture is showing a glimpse of the company's broader strategy for continuing to move quickly after SC05 is just a memory.

"This represents a re-launch of the company," Ewald said. "Our intent is to take the major steps that prove we are the Linux supercomputing company."

###

HPC Wire
November 11, 2005

Features:

Craig Mundie Speaks to HPC's Future

If you are curious about what Bill Gates will say in his keynote address to the supercomputing community next week at SC05 in Seattle, you may be able to get more than a little glimpse of it in "The Next Decade in HPC," an article by Microsoft CTO Craig Mundie, published today in the new issue of CTWatch Quarterly. As this extended excerpt from his article makes clear, Microsoft's thrust into HPC represents a major strategic investment, one designed to enable the company to better anticipate and take the lead in the new directions computing will take.

An excerpt from "The Next Decade in HPC," by Craig Mundie
The global society has an increasing need to solve some very difficult large-scale problems in engineering, science, medicine and in many other fields. Microsoft has a huge research effort that has never been focused on such problems. I believe that it is time that we started to assess some application of our research technology outside of our traditional ways of using it within our own commercial products. We think that by doing so, there is a lot that can be learned about what will be the nature of future computing systems.

Many of the things that we thought of as *de rigueur* in terms of architectural issues and design problems in supercomputers in the late eighties and early nineties have now been shrunk down to a chip. Between 2010 and 2020 many of the things that the HPC community is focusing on today will go through a similar shrinking footprint. We will wake up one day and find that the kind of architectures that we assemble today with blades and clusters are now on a chip and being put into everything. In my work on strategy for Microsoft I have to look at the 10 to 20 year horizon rather than a one to three year horizon. The company's entry into high performance computing is based on the belief that over the next ten years or so, there will be a growing number of people who will want to use these kinds of technologies to solve more and more interesting problems. Another of my motivations is my belief that the problem set, even in that first ten-year period, will expand quite dramatically in terms of the types of problems where people will use these kinds of approaches.

There was a time certainly, when I was in the HPC business, when the people who wrote high performance programs were making them for consumption largely in an engineering environment. Only a few HPC codes were more broadly used in a small number of fields of academic research. Today, it is doubtful whether there is any substantive field of academic research in engineering or science that could really progress without the use of advanced computing technologies. And these technologies are not just the architecture and the megaflops but also the tools and programming environments necessary to address these problems. ...

In parallel with these developments in HPC, we are no longer seeing the kind of heady growth in the number of trained computer scientists produced by the world's universities. In fact, in the United States, this number is actually going down. The numbers are still rising in places like India and China right now, but one can forecast fairly directly that, even if all these people were involved in engineering and science, there will not be enough of them to meet future demand. I think the problem is in fact worse than this because computer science is still a young and maturing discipline.

So another interest I have in seeing Microsoft engage with the scientific community is in helping to bridge the divide between the Computer Science community and the broader world of research and engineering. My personal belief is that what we currently know as computing is going to have to evolve substantially - and what we know as programming is going to have to evolve even more dramatically. Every person who is involved in software development will struggle to deal with the complexity that comes from assembling ever larger and more complicated and interconnected pieces of software.

Microsoft, as a company that aspires to be the world leader in providing software tools and platforms, is thinking deeply about how to solve those problems. One of the features that attracts me toward the world of high performance computing is that it is a world made up of people who have daily problems that need to get done, who live in an engineering environment but who are frequently at the bleeding edge in terms of the tools and techniques. And frankly there is a level of aggressiveness in this community that cannot really exist in basic business IT operations, particularly not at the scale where people are attempting to solve big new problems. So for all these reasons, Bill Gates and I decided that even though technical computing is not going to be the world's largest software market, this is a strategic market in the sense that the HPC community can help us all better understand these challenging problems. We therefore hope that together we can help move the ball forward in some of these very difficult areas. As we look downstream and contemplate some fairly radical changes in the nature of computing itself and the need for software tools to deal with that, we also expect that this community is a place from which technical leaders can emerge. We would like to be a part of that. We think that Microsoft has some assets that could really make a difference for the growing community of people who will need to adopt HPC technologies for their business or their research. Before too long, these people will not only want to solve the problem but will also want to be able to configure and manage these HPC systems for themselves. One thing that Microsoft can do really well is to provide good tools not only for programming but also for administration, management and security.

In the full article, which is available at the CTWatch Quarterly website (http://www.ctwatch.org/quarterly), Mundie goes on to highlight some key challenges that HPC now faces, such as the need for algorithmic innovation and dramatically improved parallelism, where Microsoft research can make a significant contribution.

###

HPC Wire
Tuesday, November 15

Features:

Gates Leads Microsoft into HPC

Today at Supercomputing 2005, an international conference for high-performance computing, networking and storage, Microsoft Corp. Chairman and Chief Software Architect Bill Gates delivered the opening keynote address to more than 7,000 attendees. In the address, titled "The Role of Computing in the Sciences," Gates shared a vision of how the software industry can contribute to accelerating scientific research and engineering innovation, calling for broad collaboration between the computing industry, academia and government to make technical computing easier and more productive. Gates also announced that, as part of the company's steps toward realizing this vision, Microsoft is funding joint research projects at ten academic centers worldwide and has released the beta 2 version of Windows Compute Cluster Server 2003, Microsoft's first product offering for high-performance computing.

"Technical computing is crucial to the many discoveries that impact our quality of life -- from making safer, more efficient cars and airplanes to addressing global health issues and environmental changes," Gates said. "Moreover, most sciences are becoming computational sciences, which is why advanced computing capabilities need to be seamlessly integrated into the end-to-end scientific process. We see many opportunities to collaborate with the scientific community on innovative solutions that will accelerate the pace of insight and discovery."

In addition to computational modeling and clusters -- groups of computers that run as unified systems -- Gates addressed the transformation resulting from the availability of massive amounts of real-world data from low-cost sensors. This powerful combination creates new opportunities, but also new challenges, particularly with how to manage, search, analyze and publish that data and the resulting conclusions. Improving end-to-end workflow will be pivotal in changing the world of work for scientists, engineers and researchers.

Gates urged software-makers and other technology vendors to work with the scientific community to spur the next revolutionary wave of discovery by creating software, tools and standards to help eliminate today's barriers to scientific collaboration.

Windows Compute Cluster Server 2003

Microsoft is working with the computing industry to help facilitate the next wave of discovery and deliver software that addresses some of the barriers scientists, engineers and researchers face. With the delivery of Microsoft Windows Compute Cluster Server 2003 Beta 2 today and general product availability scheduled for the first half of 2006, Microsoft will offer an HPC platform that accelerates the time to insight. Windows

Compute Cluster Server 2003 is designed to be simple to deploy, operate and integrate with existing infrastructure and tools.

To help ensure simple integration of Windows Compute Cluster Server 2003 with a broad set of applications and tools, Microsoft is working closely with software, application and hardware partners across the industry. This collaboration is designed to meet customers' unique needs by enabling them to choose among and run a variety of compatible HPC applications. Microsoft is being joined in its booth at Supercomputing 2005 by 17 software partners and a number of hardware partners, including Dell Inc., HP, IBM Corp., NEC Corp., Orion Multisystems Inc., Tyan Computer Corp. and Verari Systems Inc., to demonstrate 20 industry-specific HPC applications.

"The MathWorks has been providing engineers and scientists with technical computing tools for the last 20 years," said Cleve Moler, chairman and chief scientist at The MathWorks Inc. "The capabilities and direction of Windows Compute Cluster Server 2003 complement The MathWorks' mission to provide software to accelerate the pace of innovation and discovery in engineering and science. The combination of the Windows platform and MATLAB distributed computing tools allows customers to make better and faster progress with high-performance computing tasks and deliver meaningful results that can accelerate the time to insight."

Microsoft is working with software partners across the manufacturing, life sciences and geosciences, and other industries to make a robust set of commercial applications available on Microsoft Windows Compute Cluster Server 2003. This includes integrating applications with the Microsoft Message Passing Interface and the Microsoft job scheduler, and offering performance-tuning technologies to run on Windows Compute Cluster Server 2003. As part of this effort, Microsoft is working with Intel Corporation to provide software vendors with a robust set of resources -- including 64-bit hardware, software and tools from both companies, and joint on-site engineering assistance -- to enable broader application support on Windows Compute Cluster Server 2003 when it is released to the market.

Partnering With and Investing in the HPC Community

Gates also announced an investment in ten institutes for HPC worldwide. This multiyear, multimillion-dollar investment in joint research projects at these institutes will help guide ongoing software research and product innovation at Microsoft to address the most challenging technical computing problems. These institutes are Cornell University (U.S.); Nizhni Novgorod State University (Russia); Shanghai Jiao Tong University (China); Tokyo Institute of Technology (Japan); University of Southampton (England); University of Stuttgart (Germany); University of Tennessee (U.S.); University of Texas at Austin (U.S.); University of Utah (U.S.); and University of Virginia (U.S.).

"Easier access to workgroup and departmental HPC resources, such as that provided by Windows Compute Cluster Server 2003, will play an ever-increasingly important role in furthering scientific and engineering innovation in academia and industry," said Jack Dongarra, Ph.D., distinguished professor and director of the University of Tennessee Innovative Computing Lab. "Our research project will focus on self-adapting high-performance applications and software on the Windows platform, which could help scientists and engineers reduce the time and effort to discovery. I've found that Microsoft's passion for technology innovation around an integrated development and deployment software platform will help expand the potential for many who haven't had access to high-performance computing resources in the past."

In the closing statements of his keynote address, Gates reiterated Microsoft's long-term commitment to working with the HPC community and emphasized his enthusiasm for the future role of computing in scientific discovery.

###

HPC Wire
November 15, 2005

HPC Challenge Keeps Dongarra Excited about SC

HPC stalwart Jack Dongarra admits he still gets excited about going to SC -- even after all these years.

After all, it would be easy for him to be jaded. He's attended the annual industry gathering since the very first in 1988. Yet arguably, in 2005, Dongarra has even more to be excited about -- and especially today.

In the 17th year of SC, he is serving as the co-chair for the first year of the HPC Challenge award competition, the results of which are being revealed today. The goal of the competition, which is sponsored by DARPA's High Productivity Computing Systems Program and HPCwire, is to develop a set of HPC hardware and software capabilities that become de rigeur for the productive use of all HPC systems.

The awards will be presented in two designations. One is based on performance only; while the other spotlights productivity and elegant implementation. The winners will be announced at noon in a Birds of a Feather session.

Victory in the first class is based on the entry (or entries) that delivered the best performance on a base or optimized run as submitted to the HPC Challenge web site. From the suite of choices, the benchmarks judged were: Global HPL, Global RandomAccess, EP STREAM (Triad) per system and Global FFT. The prize is $500 plus a certificate, and up to four winners may be announced.

The results are not subjective; literally, a computer determines the order of finish. So far, more than 80 contributors have dared to measure up to the series of benchmarks, according to Dongarra. Unlike the well-worn Linpack benchmark, which is based solely on speed, the HPC Challenge takes in numerous other attributes to develop an all-around, better-rounded view.

"We're quite happy with the response," he added. "Like Linpack, it was slow at the start but continues to gain momentum." As the deadline goes right up to the last minute, Dongarra said he is maintaining a shroud over the results.

The second award, for most productivity, is based on the most "elegant" implementation of two or more of the HPC Challenge benchmarks with special emphasis being placed on Global HPL, Global RandomAccess, EP STREAM (Triad) per system and Global FFT. This award, which will pay $1,500 and may be split, will be weighted equally between performance and code elegance, clarity and size. This portion of the HPC Challenge awards could be considered the closest the HPC community gets to hosting its own version of "American Idol."

"It's a beauty contest - almost," Dongarra said, with a chuckle. Or perhaps he meant it is a beauty of a contest. Either way, there is likely to be lots for the panel of 10 judges to review. As the deadline for this award was Oct. 15, the judges have been busy reviewing and assessing what turned out to be 10 entrants.

Co-chair Dongarra is one of the judges - and be assured he's no Paula Abdul, Simon Cowell or Randy Jackson. He's been huddling with his associates - co-chair Jeremy Kepner MIT Lincoln Lab; David Bailey, LBNL NERSC; David Koester, MITRE; Bob Lucas, ISI; Rusty Lusk, Argonne National Lab; Piotr Luszczek, University of Tennessee; John McCalpin, IBM Austin; Rolf Rabenseifner, HLRS, Stuttgart; and Daisuke Takahashi, University of Tsukuba. - to determine the three finalists. They will take to the spotlight at the awards session.

While Dongarra said he is "delighted" and "enthusiastic" about the response and the prospects, he also can't help but add that he believes next year's competition will be even better. "We're setting the right tone," he added. "We're expecting an ever-stronger collection of competitors."

HPC Challenge's test suite is designed to provide an overview of the major strengths and weaknesses of HPC systems. With this set of benchmarks, organizations looking to purchase HPC systems will have a better understanding of system capabilities.
"The HPC Challenge benchmark suite is a better indicator than any single test of how an HPC system will perform across a spectrum of real-world applications," said Dongarra. "HPC is much more complicated today then when Linpack was developed 30 years. The HPC Challenge awards are helping set new standards for benchmarking methodology and result-reporting with a control database/repository for both the benchmarks and the results."

The goal is for the suite to take no more than twice as long as Linpack to run, Dongarra said. Future goals are to reduce execution time while expanding the set to include elements such as sparse matrix operations and developing machine signatures.
SC05 represents a homecoming of sorts for Dongarra, one filled with many memories. For example, attendees need to know SC was not a success from the start, he said. While the early conferences were focused on the technical aspects of HPC, it wasn't until the early 1990s that participating vendors started to gain in numbers and prominence. It was then that SC as a whole started to attract broader attention.

"The vendors always put on a nice show," he said. "They're always trying to outdo each other. It's fun to look at the nice, shiny machines. After all, supercomputers are the Ferraris in any stable."

Yet he is quick to add that assuring SC's viable future means there must be continued emphasis on the "other side" of supercomputing, namely software and algorithms.

"Hardware may be attraction, but software and algorithms makes things go and grow," Dongarra added.

For each SC gathering some theme emerges, and Dongarra said he's curious what this year will bring. While in the past Linux and clusters have enjoyed their time in the spotlight, the odds-on money this year is on Microsoft, especially with Bill Gates as the keynoter. "It will be interesting to see how his message resonates," Dongarra said. Dongarra said he hopes Gates' address will clarify Microsoft's interest in HPC, but all that is certain is that it will set the tone for the Seattle gala. "In some sense, he'll be talking to a hostile audience," he added. "The HPC culture is very different than Microsoft's."
One place where perhaps it's not that different is with Craig Mundie, Microsoft's senior vice president and CTO, an HPC cronie of Dongarra's from SC's early days. "I knew Craig at Alliant [Computer Systems Corp., a company that developed massively parallel supercomputers]," Dongarra said. "It's amazing how far things have come since then. We'll see how he and Tony Hey [another high profile HPC veteran who's joined Microsoft] bring about change within the company."

Nevertheless, Dongarra added that some of his favorite parts of SC are seeing the displays and exhibits that the various labs and organizations put forth. He said he is especially encouraged that industry and commercial ventures are working as hard on HPC matters as the labs.

University of Tenneesee has 10 people at SC from Dongarra's group at the Innovative Computing Laboratory. They will be spread out over an number of exhibit booths, including Rice University, ORNL, Los Alamos and, even, Microsoft.

Even when this year's SC05 is a memory, it remains an exciting time in HPC," Dongarra said. "The advent of multi-core systems bares close watching as does the international competition to make petaflop computing a reality sooner than later. Both will put their fingerprint on the future of our industry."

###

HPC Wire
November 15, 2005

Features:

HPC Analytics Returns Bigger, Better and Busier

No less an authority than SC05 chairman Bill Kramer has said that this year's HPC Analytics Challenge should be "one of the most exciting initiatives" at this year's event. Kramer has repeatedly referred to computing, networking and storage as the three "legs" of the "HPC stool," and the HPC Analytics Challenge certainly is an effort that showcases the trio in action together along with sophisticated analysis methods on a mission to solve real-world problems.

The HPC Analytics Challenge is an open competition, drawing participants from academia, education, research, government and commercial interests. The goal of the contest is to provide a forum for researchers, engineers and analysts to showcase their computationally intensive applications that solve real-world problems by way of rigorous and sophisticated methods of data analysis and high-end visualization.

Visual Numerics Inc., a 35-year-old provider of numerical analysis and visualization software, is a prime mover behind this SC05 competition. Tim Leite, the company's director of educational programs, has picked up the baton to lead this year's charge. The program was added to the SC agenda two years ago. Leite gives Phil Fraher, his president and CEO, credit for VNI's involvement. Along the way, HPC Analytics has proven to be a fine vehicle for increasing the company's industry profile and participation.

The idea originally was pitched to Kramer, who responded with interest, but voiced concerns about how the competition could be accomplished. After further consultation, many of the answers were found in assembling a solid HPC Analytics team, including co-chair Donald R. Jones from the Pacific Northwest National Lab, co-chair Paul Fussell and John S. Hurley of Boeing and Michael Schulman of Sun Microsystems Inc., to carry out the mission.

"We wanted to represent vendors, government, industry and academia," Leite explained. "All these segments deal with high levels of forecasting, data mining, optimization and predictive analysis." As with most SC-related functions, the pitch-in, community-style effort is what made it happen.

Leite admits VNI has benefited greatly from its SC involvement. While the company is well known within inner circles of the HPC community, many people still know Visual Numerics' products -- namely, IMSL and PV-Wave -- better than the company that provides them.

The Analytics competition focuses on specific areas, such as algorithms, storage and bandwidth. "Here are HPC solutions and what technology is required for these solutions," Leite explained. "By getting the solutions out there, users get a better understanding of what is possible."

In only its second year, Leite said he expected HPC Analytics would continue to face a steep acceptance and participation curve. "We could have been content with four to six submissions," he said. "But for 18 to express their intent to participate is a success." He credits a broad spectrum of solutions for the unexpected quantity. Nine final papers were submitted and six finalists have been tapped, with judging held in late August in Westminster, Colo. The winners will be announced Tuesday at SC05.

The chosen will receive a certificate and monetary award. They also will present their efforts during the contest session at SC05. Prizes will be given in several categories, most likely including Best Overall, Best Student Entry and Most Original.

As the HPC Analytics Challenge promotes further development of industry benchmarks for analytics techniques while nurturing a forum to advance HPC evaluation techniques, it's a good sign that the competition has received the green light to return in 2006. Leite promised next year's rendition will include "a bigger and better" competition, fueled by the intention to start next year's promotion and recruitment sooner. Who knows, perhaps the competition will even grow into where it gains its own track for technical papers or a tutorial on future SC programs.

###

HPC Wire
November 15, 2005

Features:

D-Wave's Quantum Computer Looks to Revolutionize HPC

By Michael Feldman

Geordie Rose can't wait for the future -- so, apparently, he's not going to. Rose is president and CEO of D-Wave Systems Inc., a company with the ambitious goal of building a quantum computing system for the commercial market. In this exclusive interview, HPCwire editor Michael Feldman spoke with Geordie Rose to discuss his company and what he intends to offer the HPC community.

HPCwire: Could you briefly describe how a quantum computer operates and how it is fundamentally different from that of a conventional computer?

GEORDIE ROSE: One of the most profound realizations of the previous century is that computation and physics are deeply connected -- all of the stuff that makes up our world is constantly "computing", that is, evolving in time, and that when we build computers we are just tapping into that ongoing computation.

When a computer is built out of devices that behave according to the rules of classical physics, like PowerPCs, Pentiums, or balls rolling down hills, some of the computing resources that nature allows are not used. Classical physics can be thought of as a highly restricted version of quantum physics. This restriction removes computational capabilities that are allowed by quantum physics but not classical physics.

When computers are built out of components that can harness the resources made available by quantum mechanics, they gain access to computational resources that are not available to conventional computers. One analogy is to think of computation as painting. In this analogy, completing a computation is like completing a painting. A classical painter only gets to use black and white paints, but a quantum painter gets to use red, yellow, black and white. If the job is to paint a picture of the sun, the classical painter can never get it right, no matter how skilled, because he doesn't have access to the right resources; that is, red and yellow paint don't exist for him, while the quantum painter finds it pretty easy.

HPCwire: Could you present some real-world examples of problems that can be addressed by quantum computers but are impossible to solve with conventional computers?

ROSE: Yes. But I should point out that this capability is not required to make quantum computers worth building. It is likely in the short term that they will be cheaper, easier to use, and just a little better on performance than enterprise-class supercomputers and this

will justify their build-out from a commercial perspective.

The main application where "exponential" improvements occur using quantum computers over classical ones is something called quantum simulation, which is the "in silico" prediction of the properties and/or behavior of an atomic or molecular scale product or process. Even very small quantum computers can outperform the most ambitious supercomputers ever conceived in these tasks because of the inherent advantage of quantum computers for this task.

For many industries, predictive modeling of atomic and molecular scale products and processes would be transformative, including the chemical, biotech, energy and pharmaceutical industries.

HPCwire: There seems to be a general consensus that quantum computing is still in its pioneering stage. What events compelled you to attempt commercialization at this time?

ROSE: It is undoubtedly true that quantum computing was in its infancy when we incorporated D-Wave in 1999. However there were some very good arguments for believing, back in 1999, that the timing was right to structure an organization whose objective was to gain and hold a leadership position in the development and commercialization of quantum computing.

Two of the principle drivers for starting the company when we did were the following: First, while there were literally hundreds of ongoing projects in universities and corporate labs studying aspects of quantum computing, there were only a handful of real efforts trying to build quantum computers. This meant that our effort stood a real chance to become the leader in this field, which we successfully accomplished. Second, even back in 1999, nearly every scientist who had thought carefully about quantum computing technology believed that eventually they would be built -- the question was when, not if. The corollary was that if we started to attack the problem early, using best practices from the semiconductor industry and careful selection of top engineers, management, investors and scientists, we could set up a high-throughput industrial-strength infrastructure to accelerate this inevitability.

HPCwire: Could you describe the overall design of D-Wave's quantum computer?

ROSE: The system we are currently deploying, which we call Trinity, is a capability-class supercomputer specifically designed to provide extremely rapid and accurate approximate answers to arbitrarily large NP-complete problems. This class of computational problem is common in naturally discrete problems with complicated networks of interactions, such as those found in some life sciences, cognition, logistics, electronic design automation and quantitative finance applications.

Trinity has a front-end software interface, implemented in a combination of Java and C, that allows a user to easily state any NP-complete problem of interest.

After such a problem has been stated the problem is compiled down to the machine language of the processors at the heart of the machine. These processors then provide an answer, which is shuttled back to the front end and provided to the user. This capability can of course be called remotely and/or as a subroutine of some other piece of software. The processors themselves are made out of lithographically defined patterns of niobium (a metal at room temperature, a superconductor when cooled) on a standard inert substrate, such as quartz. They are analog, software programmable circuits designed to harness the most robust quantum resources to accelerate computation. In the "computation is like painting" analogy, these processors can paint with yellow, but maybe not red. So you can paint different pictures than the classical black-and-white painter, but you might not be able to paint in red, even though this is an allowed quantum resource.

This is a subtle but important point that many discussions about quantum computing ignore. It is not true that there is a sharp dividing line between what is a classical and what is a quantum computer. Machines can be built that are somewhere in the middle that avoid most of the difficult engineering problems of building entirely quantum computers, but still harness quantum mechanics to massively accelerate computation of high-value, real-world problems.

HPCwire: When do you expect to have a working prototype?

ROSE: The first Trinity prototype will be running our clients' problems by March 2006.

HPCwire: How do you think the arrival of quantum computers will change the landscape of the high-performance computing industry?

ROSE: This is difficult to predict. I suspect that in the short term Trinity will be viewed as a competitor for the current enterprise- and capability-class supercomputer market for the restricted segment of the market Trinity is targeted at -- which includes NP-complete problems, monte carlo simulations and quantum simulation problems.

However Trinity differs significantly from competing supercomputer technologies in several important ways. They are cheap to build; the performance of the machines is expected to scale far beyond current limitations; Trinity generates relatively little heat during its operation because the bulk of the computation is done in superconducting electronics; Trinity takes up very little physical space; the machine is very easy to program. All of these differences will likely lead to further segmentation of the HPC market classifications, drawing market share from the enterprise- and capability-class segments and re-assigning it to the new class of machine that Trinity is. In addition of

course we hope to expand the market size based on the new capabilities of Trinity systems over other capability-class supercomputers.

While at SC05, Geordie Rose will be conducting a Masterworks session on quantum computing. It will take place on Wednesday, November 16, from 4:15 pm to 5:00 pm. For more information visit
http://sc05.supercomputing.org/schedule/event_detail.php?evid=5311

###

HPC Wire
November 15, 2005

Breaking News:
AMD, Sun to Build Largest Supercomputer in Japan

At Supercomputing 2005, AMD and Sun Microsystems Inc. have announced that the Tokyo Institute of Technology (Tokyo Tech), one of the world's leading technical institutes, is creating Japan's largest supercomputer on a foundation of Sun technology. The system is based on Sun Fire x64 (x86, 64-bit) servers with 10,480 AMD Opteron processor cores, totaling more than 50 trillion floating point operations per second, Sun and NEC storage technologies and NEC's integration expertise as well as ClearSpeed's Advance accelerator boards. Using Sun's N1 System Manager and N1 Grid Engine, the system will be provisioned to support the Solaris 10 OS as well as the Linux operating environment. It will be used to help science and engineering researchers dramatically increase their productivity. The Tokyo Tech system marks Sun's largest high performance computing win to date. The grid-based supercomputer plans to expand to more than 100 teraFLOPS by its operation in Spring 2006 and is expected to be one of the five largest supercomputers in the world as measured by the Linpack benchmark.

"Tokyo Tech's system will be leveraged by a wide range of researchers within the university and throughout the world," said Satoshi Matsuoka, professor in charge of Research Infrastructure at Global Scientific Information and Computing Center, Tokyo Institute of Technology. "These researchers are tackling complex problems ranging from analyzing the complex molecular structure of proteins, simulated bloodflow diagnosis in human brains, modeling of the generation mechanism of Earth and planetary magnetic field and their long term effects, to nanoscience simulation of carbon nanotubes -- all tasks that require exceptional computing power and experience working with supercomputers."

"Sun's 'Tools for TeraFLOPS' 64-bit innovations that enable rapid deployment of power- and cost-efficient, terascale compute clusters-are helping us make significant headway in the demanding HPC market," said John Fowler, executive vice president, Network Systems Group at Sun Microsystems. "In conjunction with our partners, Sun is able to provide the critical processing power required for the world's most compute-intense applications while keeping total cost of ownership at a minimum."

Sun partners AMD and NEC are playing pivotal roles in the development of Asia's largest supercomputer. AMD offers customers the x86 processor multi-core technology. NEC, the primary systems integrator, is leading the design of the infrastructure as well as the integration of the various applications which will run on the system, based on their experience in building and managing high-end HPC systems. The Tokyo Tech system plans to be the world's largest and fastest cluster as measured by core CPU count and peak performance, respectively. It will include more than 21 terabytes of memory and 1.1 petabytes of hard disk storage, again exceeding all competing systems within the Asia-

Pacific region. There are planned increases in performance to beyond 100 teraFLOPS with installation of additional ClearSpeed Advance boards (initially from 360 to more than 600) by the time of the system's operation in the Spring 2006.

"Tokyo Tech is among the numerous leading research and academic institutions that rely on AMD64 technology to rapidly and cost-effectively deliver results," said Kevin Knox, vice president, Commercial Business at AMD. "The superb design of the dual-core AMD Opteron processor is optimized for HPC and enterprise environments. Built from the ground-up for 64-bit and true multi-core computing, the AMD Opteron processor delivers leading-edge performance in the same power envelope as single core processors. As evidenced by our growing market share, corporations, academia and government agencies around the world are recognizing the unique combination of industry-leading performance and performance-per-watt AMD brings to x86 server solutions such as the Sun Fire product line."

In addition to working with AMD and NEC, Sun's system also incorporates technology from ClearSpeed Technology Inc., ClusterFS and Voltaire into the Tokyo Tech system. ClearSpeed's Advance board accelerates common HPC algorithms, optimizes the software to achieve maximum performance; ClusterFS's Lustre parallel file system software allows the servers to communicate with the storage in parallel, speeding access to the ever increasing amounts of scientific data being processed; and Voltaire is supplying high speed multi-protocol InfiniBand switches and host card adapters to connect both the servers and storage.

###

HPC Wire
November 15, 2005

Breaking News:
SGI Unveils Native InfiniBand Storage

Silicon Graphics Inc. has announced that it will offer a native InfiniBand interconnect solution that includes both servers and storage integrated in an InfiniBand fabric. This will allow customers to leverage previous investments in InfiniBand by directly connecting servers in a current InfiniBand cluster with native InfiniBand storage systems from SGI using commercially available InfiniBand products.

This week at SC05 in Seattle, SGI, its partners and its customers will demonstrate storage solutions enabled by a new native InfiniBand interface for its SGI InfiniteStorage TP9700 disk array. Supporting the SGI InfiniteStorage vision to integrate seamlessly and scale infinitely, the new TP9700 capability reduces complexity, lowers costs and maximizes throughput for InfiniBand-based Linux OS clusters.

"The most daunting problem HPC customers face today is the mounting challenge of accessing and managing the explosive amount of data they work with. Customers often end up with a patchwork of various networking products," said John Howarth, director, InfiniteStorage Marketing, SGI. "Today, SGI has dramatically simplified this difficult obstacle, giving customers a single complete, seamless, and scalable InfiniBand fabric capable of moving data at up to five times the speed of traditional Fibre Channel. This solution is ideal for an industry thirsty for bandwidth and intolerant of high latencies." The new InfiniBand-based TP9700 system provides four 10 Gb/second InfiniBand host connections which can be plugged directly into the InfiniBand fabric. SGI's TP9700 is based on Engenio's 6498 storage system and interfaces with commercially available InfiniBand components from Voltaire and Mellanox Technologies.

SC05 attendees can see the InfiniBand fabric in action at the SGI Booth (No. 602). The demonstration will incorporate geospatial data residing on a TP9700 array with InfiniBand located at the U.S. Naval Research Laboratory (NRL) in Maryland. The database will be mirrored on a TP9700 array in the SGI booth also with the InfiniBand controller. Users will visualize a data set being shared between two InfiniBand fabrics thousands of miles apart over a WAN connection provided by YottaYotta.

Attendees can also take the SGI InfiniBand Tour at SC05 by visiting SGI's InfiniBand partners, including Engenio Information Technologies Inc. (in the SGI booth), Voltaire (Booth 1630), Yotta Yotta (2118) and Mellanox Technologies (Booth 902).

HPC Wire
November 16, 2005

Features:

Simon Extends Perfect SC Attendance Record

If among the awards given annually at the Supercomputing event was a certificate for perfect attendance, count on Horst D. Simon to be on the elite list of recipients. After all, Simon, the associate laboratory director for computing sciences at Berkeley Lab, has been present at every SC since its debut in Orlando, Fla., in 1988. While admitting he's "gotten tired a couple of times" over the course of the past 17 years, his enthusiasm for the HPC industry's "prom" is solidly rejuvenated for the Seattle gala. As Simon reflects on SC's legacy, he said it's because the conference has again regained much of its early excitement. Credit goes to SC05's expanded program and the zestful willingness of participation from all walks of the supercomputing community.

Simon categorizes SC's "formative years" as its initial phase, from 1988 to 1991. Even after the first two years, he said it still was not clear that SC would emerge as the "must attend" event it is today. "There were still other competing conferences on people's agendas," he explained. "And the supercomputing industry itself was in transition."

The culmination of the first phase came with the gathering in Albuquerque, N.M., in 1991. "That was the big change," Simon recalled. "SC went from being a small technical conference to a big-time event." The upgrade was powered by the first en masse appearance of big-name vendors exhibiting in spacious and well-appointed booths. This earmarked SC as no longer just an ACM- or IEEE-dominated conference, but one that had metamorphosed into a full-fledged industry trade show, he added.

Simon categorizes the period between 1997 and 2001 as a "downtime" for the annual gathering. The "dot com" craze had stolen much of the high tech thunder and the aftermath of the attack on the World Trade Center put everyone in a pall. During that time, according to Simon, the HPC community continued to mature but it didn't seem to be moving forward as it continued to be overshadowed by the hubbub from the flurry of Internet-based start-ups.

So it was only appropriate that a Phoenix-like re-emergence of the conference came when SC ventured to Arizona (Phoenix, to be exact) in 2003. Simon recalled the event was typified by a conversation he had with a well-known high tech reporter from the New York Times. "He'd just left Comdex to come here [to SC]," Simon explained. "While he didn't see much in Las Vegas, at SC he saw lots of excitement."

New trends were being reflected while there seemed to be new and renewed interest from attendees. SC has continued to grow at a healthy rate since then, Simon said. He credits the sustained popularity to influx of people from outside the traditional realm of supercomputing, such as the various supplier technologies, including communications,

networking and storage.

At the same time, supercomputing has seen clusters become big business while big-name latecomers such as Apple, Dell and Microsoft have proved their continued interest in the field. "Two years ago, you would not put 'Dell' and 'HPC' in the same sentence," he added. "We've come quite a way."

The result of this new equation is that SC is "less and less just for HPC purists" and the event has evolved into a more far-reaching attraction beckoning the entire computing industry. This does not mean, however, that the hard work is done.
"There is hardly any European supercomputing industry," Simon said. "It's mostly all in the U.S. Yes, Japan has had some success, notably with the Earth Simulator, but its share is shrinking."

Simon looks for the growth at future SC events coming from a massive surge in commodity systems and parallel computing. But while vendors remain in a bright spotlight at SC, they cannot keep the growth going single-handedly.

"There is a national initiative needed," Simon said. "Washington, D.C. needs to help, to get involved -- and quickly." He cited the petaflop projects underway in China and Japan, and added that the U.S. government must find a way to invest in the future of supercomputing.

"Supercomputing remains one of the key technologies for U.S. leadership," he said. "Yet we do not have the national focus we had back in '91." If SC does not remain "a community event," Simon said it would suffer the fate of many other commercialized industry gatherings, such as most notably of late, Comdex. The once proud, must-attend event for the computing industry has fallen on hard times, being canceled this year and last, basically due to industry ennui. Few will recall that Comdex drew 200,000 attendees in 2000. SC's movers and shakers can learn valuable lessons here.

Yet SC should be able to stave off such an occurrence, Simon said, because the event has skillfully diversified, bringing in and emphasizing new elements of the annual program, such as Education and SC Desktop.

One thing is sure: SC coming to Seattle puts a trio of big-name local companies firmly in the spotlight. Microsoft, for one, will be in everyone's gaze. "It will be interesting to see how Gates' keynote plays out," said Simon. "It will be good for people to pick up on some of what he is thinking."

For sentimental hometown favorite Cray Inc., Simon said the beleaguered supercomputer maker has an opportunity to step up while on familiar ground. "There have been big transitions the past two years," he added. "On its own home turf, Cray can deliver a clear

message of how it is going to go forward."

And let's not forget Boeing, for which Simon once worked -- Boeing Computer Services, to be exact -- in the late 1980s. "After being an industrial HPC leader in those times, there's a resurgence in interest."

###

HPC Wire
November 16, 2005

Breaking News:
HP Advances HPC with Dynamic Visualization Technology

HP introduced visualization technology designed to make high-performance computing (HPC) more affordable and productive for a broader range of customers.

Launched at SC'05, the HP Scalable Visualization Array (SVA) uses industry-standard components and leverages advances in clustering, graphics and networking technology to allow scientists and engineers to visualize and analyze complex sets of data.

HP SVA is a key part of the HP Unified Cluster Portfolio, which is designed to integrate computation, data management and visualization in a single cluster environment, greatly enhancing productivity in HPC environments. Driven by this portfolio, HP continues to lead the HPC market in revenue with more than 30 percent market share, according to second quarter 2005 figures released by IDC.

Customers in aerospace, oil and gas, automotive, manufacturing, defense, medical and scientific research organizations using the HP SVA can match their growing computational capability with powerful visualization tools. SVA technology can produce a vast, high- resolution display wall of 100 million pixels or more at approximately half the cost of competitive proprietary products.

"Our ability to generate large data sets has outstripped our ability to visualize them, both for deriving science and verifying correctness," said Hugh Couchman, scientific director of SHARCNET, a prominent HPC institute in Ontario that conducts leading-edge research in science and engineering. "With HP SVA, integrated with the HP Scalable File Share and XC software, we will now be able to visualize and manipulate huge data sets in real time as an extension of computing and data management in the cluster."

Along with the new array, HP announced additional enhancements to the HP Unified Cluster Portfolio, including:

- Version 3.0 of the HP XC System Software, including support for workstations operated as SVA nodes, large node options and new distributed management and provisioning features.
- Enhancements to HP StorageWorks Scalable File Share, including support for dual-gigabit Ethernet and Advanced Data Guard protection.
- Support for Microsoft Windows Compute Cluster Server 2003.

- Updates to HP ClusterPack for HP-UX, with new tools for easier custom configuration and console management.
- Expanded options for HP Cluster Platforms, including systems with dual-core Intel Xeon processors and HP Cluster Platform Express.
- Novell Validation Suite HPC and Grid Project updates.

The HP SVA consists of a cluster of HP workstations running Linux, industry-standard graphics cards and network adaptors, and an integrated software system. System software includes HP XC System Software for cluster infrastructure and workload management, and HP Scalable Visualization Array software for configuration and job management. HP StorageWorks Scalable File Share and HP Remote Graphics software are optional.

"The HP Scalable Visualization Array reduces the cost of visualization solutions through the use of industry standards and interoperable products," said Winston Prather, senior vice president and general manager of the High Performance Computing Division at HP. "By increasing scalability and flexibility, and integrating computation, data management and visualization with our Unified Cluster Portfolio, HP is meeting the needs of more customers in more application areas and fostering faster scientific discovery around the globe."

In concert with the SVA launch, HP announced a new release of HP XC System Software and updates to HP SFS.

XC version 3.0 supports HP workstations as SVA visualization nodes and large 16-processor node options, providing a single, managed cluster system that can support distributed and SMP applications. New features also include distributed management services, which allocate tasks to multiple nodes to improve system performance and availability, and expanded smart provisioning, which automatically determines cluster topology and the optimal configuration of system services.

HP SFS updates include support for dual-gigabit Ethernet, providing a 50 percent increase in input/output bandwidth for Ethernet connected clusters; high-resiliency Advanced Data Guard protection, which protects against even double-simultaneous disk failures; support for double-density disks, providing up to 1,000TB of storage per file system; and improved resiliency to better detect network failures.

As part of a SC'05 StorCloud demonstration, HP plans to showcase a 64TB, 40 gigabit-per-second HP SFS server connected over a Cisco Systems-powered wide-area-network to a cluster of HP Integrity servers located at Pacific National Labs 200 miles away. HP also announced that it will support Windows Compute Cluster Server 2003, which is scheduled for release in the first half of 2006. With the addition of Windows Compute Cluster Server 2003 to its Unfied Cluster Portfolio, HP will offer customers strong

choices of HP-UX, Windows and Linux to meet their HPC operating environment requirements. Windows Compute Cluster Server 2003 is expected to be available on HP Cluster Platform 3000 and 4000 systems and Cluster Platform 4000 blade systems.

The HP Message Passing Interface (HP-MPI), a leading implementation of the MPI standard and critical for creating and porting parallel applications, will be ported to Windows Compute Cluster Server 2003. In addition, HP and Microsoft are working with software vendors to port and optimize 64-bit applications to Windows Computer Cluster Server 2003.

In support of HP-UX, HP ClusterPack 2.4 offers new tools for easier custom configuration and console management. It also includes many enhancements for easier deployment of HP-UX clusters in production environments where interruptions from adding or deleting compute nodes must be minimized. It is also now integrated with HP System Insight Manager software.

The HP Cluster Platform 3000 now offers dual-core Intel Xeon processors as an option in the HP ProLiant DL380 control nodes. Dual- core AMD Opteron processors are available with both the HP Cluster Platform 4000 and the blade-based HP Cluster Platform 4000BL. HP Cluster Platforms are available in clusters of up to 1,024 nodes.

HP Cluster Platform Express is a faster, easier way to configure and order single-rack cluster solutions compared to traditional customization methods. Customers can select from a menu of popular cluster components, which are then factory assembled to reduce integration costs while shortening the time to deployment.

In other HPC news, HP announced that the infrastructure software for the Novell Validation Suite -- High Performance Computing and Grid Project, announced in March 2005 by HP and Novell, has been validated.

To date, HP and Novell and participating software vendors have validated a range of infrastructure applications, including Altair, PolyServe Matrix Server and Cluster Volume Manager, Scali Connect and Scali Manage, as well as Grid resource managers from Axceleon, Data Synapse, TurboWorx and United Devices. CAE applications from ABAQUS, ANSYS and MSC NASTRAN have also been validated for the project.
The Novell Validation Suite program focuses on defining and certifying integrated multi-application and platform stacks on Novell SUSE Linux to reduce the risk of deploying complex Linux solutions in HPC environments.

###

HPC Wire
November 16, 2005

Breaking News:
AMCC Demonstrates SATA II RAID Solutions

Applied Micro Circuits Corporation (AMCC) will demonstrate high port count, high performance computing (HPC) SATA II RAID solutions based on its recently announced 3ware 9550SX controllers for super computing applications at SC05 November 13 - 15 in booth 6202.

AMCC's new 3ware RAID controllers will also power the SC05 Supercomputing StorCloud, a massive HPC storage network that showcases state of the art heterogeneous devices and technologies that comprise a virtual on-site "storage on request" capability to support researchers and demonstrate high bandwidth applications at the conference. The AMCC demonstration will showcase a benchmark application featuring 24 TB capacity in a 5U enclosure from Pogo Linux powered by four 3ware 9550SX 12 port SATA II RAID controllers. AMCC will also aggregate 2.4GB sustained reads bandwidth across four 3ware 9550SX XXport RAID controller cards.

"Every year, SC reaches scientists, researchers, engineers, CIOs and IT administrators from universities, industry and government fields seeking solutions to improve, accelerate, scale, and streamline operations," said Michael Joyce, Director of Marketing for AMCC Storage. "Our new 9550SX RAID controllers, in conjunction with solutions provided by our onsite partners, provide powerful, manageable, cost-effective RAIDed storage solutions that enable modeling and simulation systems to solve typically unsolvable problems in biotechnology, weather research, astronomy, drug research, wearable sensors and many other critical applications."

"Hitachi drives some of the industry's most powerful supercomputers around the world," said XX of Hitachi XX. "The new XX 500GB SATA drives featuring the latest SATA II specifications are the highest speed SATA drives on the market. When combined with AMCC's new 3ware 9550SX controllers, HPCs are powered by an entirely new class of incredibly high performing, high port count, ultra high capacity storage RAID solutions for next generation virtualization, simulation, and modeling applications."

"Today's HPC customers are seeking performance, scalability, and reliability, combined with cost efficiency in order to build systems that can process terabytes of data without soaring technology and management costs," said Tim Lee, CEO of Pogo Linux. "Pogo Linux is eager to showcase the power of the Linux operating system in such a high performance environment. Our server not only maximizes the inherent power of the operating system, it takes full advantage of AMCC's new SATA II RAID controller. The demo unit is a clear indication that SATA II drives can handle both the speed and complexity of almost unlimited storage access at record speeds."

The new 3ware 9550SX is 200 percent faster than the company's previous generation 9500S, delivering RAID 5 writes at over 380 MB/sec and RAID 5 reads at more than 800 MB/sec and supports the Serial ATA 2.5 specification and a PCI-X host interface.

###

HPC Wire
November 16, 2005

Breaking News:
IBM Wins Big in HPC Challenge Competition

IBM has dominated the first annual HPC Challenge Competition at the Supercomputing 2005 Conference. The DOE/NNSA/LLNL team, using IBM's Blue Gene/L system, swept all four Class 1 benchmark awards. The IBM team tied with the Cray team for the Class 2 "elegant" implemention award.

The following are the winners of the HPC Challenge Class 1 Awards:
G-HPL - IBM Blue Gene/L, DOE/NNSA/LLNL; G-RandomAccess - IBM Blue Gene/L, DOE/NNSA/LLNL; EP-STREAM system - IBM Blue Gene/L, DOE/NNSA/LLNL; G-FFT - IBM Blue Gene/L, DOE/NNSA/LLNL

Class 2: Most Productivity
This award defined as the most "elegant" implementation of two or more of the HPC Challenge benchmarks with special emphasis being placed on: Global HPL, Global RandomAccess, EP STREAM (Triad) per system and Global FFT. The award is weighted 50% on performance and 50% on code elegance, clarity, and size.

The HPC Challenge Class 2 Award was shared by:
Calin Cascaval, C. Barton, G. Almasi, Y. Zheng, M. Farreras, P. Luk, and R. Mak - IBM; Petr Konency, Simon Kahan, and John Feo – Cray

HPC Wire
November 17, 2005
Breaking News:
The MathWorks Upgrades Toolbox

The MathWorks, a company that provides distributed computing tools for technical computing, has introduced the latest version of its high-productivity development tool, the Distributed Computing Toolbox 2. This latest version of the software enables engineers and scientists to simplify the task of developing distributed computing applications. The Distributed Computing Toolbox 2 was featured by Microsoft chairman Bill Gates in his keynote address at the Supercomputing 2005 Conference, on November 15.

"Advancements in high-performance technical computing tools are crucial to engineers and scientists tasked with solving today's toughest computing challenges," said Kyril Faenov, director of high performance computing at Microsoft Corp. "The Distributed Computing Toolbox 2 from The MathWorks, coupled with Microsoft Windows Compute Cluster Server 2003, brings the power of high-performance computing to the departmental and workgroup levels. As a result, the administration and management of clusters can be simplified for those groups who typically had to rely on dedicated IT groups in the past."

In addition to announcing the release of the Distributed Computing Toolbox 2, Company representatives, including MathWorks cofounder and creator of MATLAB, Cleve Moler, shared their expertise throughout the SC05's extensive technical program by participating in tutorials, panels, and presentations at partner booths.

The MathWorks participated in the following SC05 technical programs and demonstrations:

- Loren Dean and Silvina Grad-Freilich of The MathWorks co-presented the tutorial, "Parallel and Distributed Computing with MATLAB," with Jeremy Kepner and Albert Reuther of MIT Lincoln Laboratory. The tutorial used customer stories to illustrate some of the most popular technologies for writing parallel and distributed MATLAB programs.
- Ohio Supercomputing Center's booth featured a Parallel MATLAB mini-symposium.
- Microsoft and Platform Computing featured demonstrations of the Distributed Computing Toolbox 2 in their booths.
- Orion and ClearSpeed partner booths both featured MathWorks tools, while partner Hewlett Packard demonstrated how
- MathWorks tools are integrated with the Enfuzion scheduler from Axceleon.
- The Birds-of-a-Feather session, "2005 HP Challenge Award," included a presentation by Cleve Moler.

Cleve Moler was also a member of the panel "Return of HPC Survivor--Outwit, Outlast, Outcompute." The panel was structured as a series of "rounds," each posing a specific question about system design, philosophy, implementation, or use.

HPC Wire
November 15, 2005

Breaking News:
HPCwire 2005 Readers' and Editors' Choice Awards

HPCwire, the publication of record for high performance computing, has announced the winners of the 2005 Reader's Choice Awards. The annual polling of the HPCwire's global readership produced the winners of The Reader's Choice Awards, and the world class collection of contributing editors selected the winners for The Editor's Choice Awards. The awards were presented to the executives leading these prestigious firms at the Supercomputing 2005 Conference being held in Seattle, WA.

The HPCwire 2005 Readers' and Editors' Choice Awards

Most Significant New HPC Product
Readers' Choice Award Recipient: AMD64 Dual Core Processor
Editors' Choice Award Recipient: IBM

Most Important Overall HPC Technology or Product
Readers' Choice Award Recipient: AMD64 Dual Core Processor
Editors' Choice Award Recipient: IBM

Best Hardware Price / Performance Solution
Readers' Choice Award Recipient: Dell PowerEdge Servers
Editors' Choice Award Recipient: Dell PowerEdge Servers

Most Innovative Storage Technology Vendor
Readers' Choice Award Recipient: DataDirect Networks
Editors' Choice Award Recipient: HP Storage Works Scalable File Share (HP SFS) based on Lustre Technology from Cluster File Systems Inc.

Best Price / Performance Storage Solution
Readers' Choice Award Recipient: DataDirect Networks / RAID Storage Networking System
Editors' Choice Award Recipient: Panasas, ActiveScale File System

Most Significant New HPC Product
Readers' Choice Award Recipient: NPACI Rocks 4.0
Editors' Choice Award Recipient: PathScale, EKO Compiler Suite

Most Important Software Innovation
Readers' Choice Award Recipient: Myricom, Myrinet MX Software
Editors' Choice Award Recipient: NPACI Rocks 4.0

Best Software Price / Performance
Readers' Choice Award Recipient: PathScale EKOPath Compiler Suite
Editors' Choice Award Recipients (Tie): Etnus, TotalView; PathScale, EKOPath Compiler Suite

Most Innovative Visualization Technology or Product
Readers' Choice Award Recipient: SGI
Editors' Choice Award Recipient: SGI

Best Price / Performance In High-end Graphics or Visualization
Readers' Choice Award Recipient: SGI
Editors' Choice Award Recipient: SGI

Most Innovative HPC Networking Solutions
Readers' Choice Award Recipient: Myricom
Editors' Choice Award Recipient: Mellanox

Best HPC Networking Price / Performance Solution
Readers' Choice Award Recipient: Myricom
Editors' Choice Award Recipient: Intel

Most Innovative HPC Cluster Technology
Readers' Choice Award Recipient (Tie): IBM; Linux Networx
Editors' Choice Award Recipient: Linux Networx

Best Price / Performance HPC Cluster
Readers' Choice Award Recipient (Tie): Linux Networx; Sun Microsystems
Editors' Choice Award Recipient: Sun Microsystems

Most Innovative HPC Technology for 2005 / Overall Top HPC Vendor
Readers' Choice Award Recipient: Cray
Editors' Choice Award Recipient: Intel

Most Innovative Implementation of an HPC Application
Readers' Choice Award Recipient(s): Fluent, American Power Conversion, AMD and Ferrari / The Ferrari Data Center

Editors' Choice Award Recipient(s): Fluent, American PowerConversion, AMD and Ferrari / The Ferrari Data Center

Best Collaboration Government & Industry
Readers' Choice Award Recipient: Lawrence Livermore National Laboratory (LLNL) & IBM

Editors' Choice Award Recipient: Cornell Theory Center

HPC Public Awareness Award
Readers' Choice Award Recipient: Intel
Editors' Choice Award Recepient: CERN

HPC Community Recognition Award "For Communicating the Importance of HPC Technology & Raising Public Awareness"
Readers' Choice Award Recipient: Jack Dongarra, University Distinguished Professor at the University of Tennessee and Distinguished Research Staff at Oak Ridge National Laboratory

Editors' Choice Award Recipient: Dr. Thomas Zacharia, Associate Laboratory Director for Computing and Computational Sciences Oak Ridge National Laboratory

###

HPC Wire
November 16, 2005

Features:

Boeing's Shukla Talks Clusters

Suresh Shukla is HPC service manager at the Boeing Co. In this capacity, he is at the forefront of the manufacturer's efforts in high performance computing. He has managed the HPC service for the company for more than a decade, including the first commercial installation of a Cray X1 (2004). His duties also include managing a large Origin, clusters from Linux Networx Inc. as well as IBM gear. HPCwire caught up with Shukla as he prepared for the trip to SC05.

HPCwire: How many SCs have you attended?
SHUKLA: I have attended at least six, as I remember.

HPCwire: How have the annual conferences changed over the years?
SHUKLA: There has been expanding audience, and expanding set of topics. Also, there are more on-site demonstrations of technology today. I would say the conference has grown with the improvements to HPC technology.

HPCwire: Is there anything in particular you are excited about seeing this year?
SHUKLA: I would like to see upcoming technological improvements in the storage and networking solutions for clusters.

HPCwire: Is Boeing planning anything big as this is a "home game" with the event in Seattle?
SHUKLA: Yes. We are "a corporate supporter" of the SC05 reception on Thursday evening at the Museum of Flight in Seattle. The participants should enjoy that experience. We also have a booth. Participants would see some displays of the latest technologies developed at or deployed by Boeing at both the places.
Plus, we have some scheduled talks at the conference. A highlight will be a presentation by Dan Mooney, vice president for product development at Boeing Commercial Airplanes, on how advanced computation enhances Boeing's ability to envision and realize new high technology products.

HPCwire: How will you spend the time?
SHUKLA: I will spend Tuesday listening to speeches, and will start attending sessions of my interest. I will switch between conference and exibit floor on all days.

HPCwire: In conclusion, do you have any comments on the current state of the HPC industry?
SHUKLA: The industry has moved toward cluster solutions due to their cost advantage. I do not know if that helps promote high-end computing research with newer algorithms, faster machines, and improved solutions to upcoming complex problems of the 21st

Century. Industry needs to quantify benefits of high-end computing better, rather than concentrating simply on lowering acquisition costs. They can then participate in furthering high-end computing hardware, software and applications to the next level. This will increase interest from the vendors in delivering even better hardware, software and applications for all involved. This is important for U.S. from competitiveness perspective.

###

HPC Wire
November 16, 2005

Breaking News:

Fernbach, Cray Awards Presented for HPC Contributions

John Bell, a senior mathematician at the U.S. Department of Energy's Lawrence Berkeley National Laboratory, and Steven Scott, the chief architect of the Cray X1 supercomputer, have been named recipients of the 2005 Sidney Fernbach Award and the Seymour Cray Science & Engineering Award, respectively.

The awards are given by the IEEE Computer Society for contributions in high performance computer systems and applications using innovative approaches. The awards were presented at SC'05, the international conference for high performance computing networking and storage. In conjunction with the award, Bell and Scott will gave special lectures today.

Scott, the chief architect of the Cray X1 supercomputer, was named for developing a highly scalable, distributed shared memory multiprocessor, employing custom vector processors. The Cray X1 is a shared memory multiprocessor capable of scaling to kilo-processor count employing a custom vector processor and efficient synchronization between scalar/vector pipelines, single-stream processors (SSPs), as well as multi-stream processors (MSPs).

According to the IEEE Computer Society citation, "There are some in the field who have perhaps one successful system to their credit, but not a sustained career of success. Steven Scott stands almost alone in sustained, commercially successful, high-end computers in the spirit of Seymour Cray."

Bell, department head for the Center for Computational Sciences and Engineering in the Computational Research Division, was named for his outstanding contributions to the development of numerical algorithms, mathematical, and computational tools and the application of those methods to conduct leading-edge scientific investigations in combustion, fluid dynamics and condensed matter.

"John is at the top-most tier of the computational science community," according to the citation. "His sustained record of outstanding technical contributions broadly used and accepted techniques and top quality mathematical software packages clearly qualify him for the Fernbach Award."

HPC Wire
November 17, 2005

Features:

Top Researchers, Projects in HPC Honored at SC'05

Top researchers and their unprecedented accomplishments in high performance computing were recognized at the SC'05 conference this week, where the winners of the Gordon Bell Prize, the HPC Analytics Challenge, the best research papers, Bandwidth Challenge Award and poster winners were announced.

Every year, SC presents awards that recognize the innovative work of conference participants and leaders in the field. The conference gives awards for Best Paper, Best Student Paper, Best Poster, and the Initiatives Challenges Awards. In addition, SC serves as the venue for presenting the Gordon Bell Prizes, which reward practical uses of high-performance computers, including best performance of an application and best achievement in cost-performance.

The following are this year's awards:

- Gordon Bell Prize:

"100+ TFlop Solidification Simulations on BlueGene/L"
Authors: Frederick H. Streitz, James N. Glosli, Mehul V. Patel, Bor Chan, Robert K. Yates, Bronis R. de Supinski (Lawrence Livermore National Laboratory), James Sexton, John A. Gunnels (IBM)

- Best Technical Paper Award (tie):

"High Resolution Aerospace Applications using the NASA Columbia Supercomputer."
Authors: Dimitri J. Mavriplis (University of Wyoming), Michael J. Aftosmis (NASA Ames Research Center), Marsha Berger (Courant Institute),

"Full Electron Calculation Beyond 20,000 Atoms: Ground Electronic State of Photosynthetic Proteins"
Authors: Tsutomu Ikegami, Toyokazu Ishida, Dmitri G. Fedorov, Kazuo Kitaura, Yuichi Inadomi, Hiroaki Umeda, Mitsuo Yokokawa, Satoshi Sekiguchi (National Institute of Advanced Industrial Science and Technology)

- Best Student Paper:

"Programmer Productivity: A Case Study of Novice Parallel Programmers."
Author: Lorin Hochstein (University of Maryland)

- Best Research Poster:

"Ultra-Low Latency Optical Networks for Next Generation Supercomputers."
Authors: Benjamin A. Small, Odile Liboiron-Ladouceur, Assaf Shacham, Keren Bergman (Columbia University); Carl Gray, Cory Hawkins, David C. Keezer, Kevin P. Martin, D. Scott Wills (Georgia Institute of Technology); Gary D. Hughes (Laboratory of

Physical Science)

- Tri-Challenge Award:

Robert L. Grossman, National Center for Data Mining

- HPC Analytics Challenge Award (tie):

"SPICE: Simulated Pore Interactive Computing Experiment"
Authors: Shantenu Jha, Peter Coveney, Matt Harvey (University College London), Stephen Pickles, Robin Pinning (University of Manchester), Peter Clarke (University of Edinburgh), Bruce Boghosian (Tufts University), Charlie Catlett (TeraGrid), Charles Laughton (Nottingham University), Rob Pennington (NCSA/TeraGrid), Sergiu Sanielevici (Pittsburgh Supercomputing Center), Jennifer Schopf (Argonne National Lab), Richard Blake (CCLRC Daresbury)

"Real Time Change Detection and Alerts from Highway Traffic Data"
Authors: Robert L. Grossman, Michal Sabala, Anushka Aanand, Pei Zhang, Jason Leigh, David Hanley, Peter Nelson (University of Illinois at Chicago), John Chaves, Steve Vejcik (Open Data Partners), John Dillenburg, Vince Poor (Princeton University)

- StorCloud Award

"PNNL Computational Chemistry Simulation"
Authors: Kevin Regimbal, Ryan Mooney (Pacific Northwest National Laboratory), Evan Felix.

- Bandwidth Challenge

"Distributed TeraByte Particle Physics Data Sample Analysis"
Authors: Julian James Bunn, Harvey Newman (Caltech), Les Cottrell (SLAC), Don Petravik, Matt Crawford (FNAL)

###

HPC Wire
November 17, 2005

Features:

Muzio Maintains Busy Schedule in HPC Circles

Paul C. Muzio is so diversified in HPC circles that perhaps his business card needs to be two-sided. With his duties as a vice president at Network Computer Services Inc., he keeps quite active as the company is the systems integrator and computing facilities manager for the Army High Performance Computing Research Center, where serves as support infrastructure director.

The AHPCRC, which combines the efforts of government, universities and industry, serves the U. S. Army by helping develop defense technology while providing a forum for researchers with training and research opportunities in science, engineering and the computational sciences. As part of its relationship with AHPCRC, NCSI operates the center's Cray Inc. X1 systems.

Showing another side of him, for the past two years Muzio also has served as the chairman for the HPC User Forum, which held a successful meeting in late September at Oak Ridge National Laboratory.

But there is nothing two-sided about his zest for and appreciation of high performance computing. So it's no wonder that he is geared up for his annual trip to SC. While he's attended most of the gatherings since the 1990s, this year's sojourn will have to be a condensed one as duty -- more specifically Uncle Sam -- has called. "The Army has told me that I must be at a meeting in Baltimore on the Thursday of SC," Muzio said from his Minneapolis office. "So I'll be flying cross-country on the night before."

But before then, look for Muzio patrolling the AHPCRC booth at SC'05. The Army will have on display several unique applications it has developed. One includes three-dimensional projections of lightweight body armor, involving the participation of the National Science Foundation-funded Engaging People in Cyberinfrastructure (EPIC) program. The other, he explained, showcases a high level of fluid dynamics with a dynamically adaptive three-dimensional finite element mesh.

As Muzio explained, while performing fluid dynamics around the static condition of, say, a Boeing 747 airplane is one thing, this presentation shows the concept modeled around the ultra-fast flapping of a hummingbird's wings.

The adaptive movement was solved in three dimensions using parallel computing, Muzio said. Making this a reality translates into the Army being able to use the capability to model the aerodynamic characteristics of small, unmanned vehicles. Aside from

modeling the movement, he said power requirements can also be taken into consideration.

An additional presentation of the application, which runs on a Cray Inc. supercomputer, will be shown at the vendor's booth. The work was written in Unified Parallel C, a newly evolving computer language.

"While the application is still in the preliminary stages, it is far enough along that we wanted to show it is viable, Muzio explained. "It is certainly a unique HPC application."

As for many in the HPC circles, going to SC05 affords an excellent opportunity to see old friends and new products. In particular, Muzio said he is looking forward to seeing some of the new products that was discussed with him under non-disclosure agreements at the recent HPC User Forum gathering. "While we tend to work at the high end, with lots of custom architecture, I am excited to see some of the new, more commercial products that have been talked about." In particular, Muzio said he is curious about some of the custom cluster systems.

While on the showfloor he said he expects the atmosphere to feel like "a farmers' market on the weekend," with lots of bustling about and excitement over what will be on view. While Muzio added that he expected 2005 to be an "up" year for conference, he's remained interested in the event even through some of its down years.

"In the beginning, SC was applications focused," he said. "People were talking about breakthrough computations using supercomputers." While the event has grown, Muzio said the recent SCs have too much of an academic focus.

"There are great sessions on quantum chemistry," he explained. "But I would rather see them focus on the HPC aspects, with more emphasis on the applications."

As with seemingly all SC attendees, Muzio said he is curious to hear what Microsoft Corp.'s Bill Gates will have to say in his opening keynote. The interest is all the more heightened because over the past two years AHPCRC has transitioned to where it is almost exclusively a Mac-based shop, have forsaken Microsoft Windows.

"We started evaluating Macs in November 2003," Muzio said. "After thoroughly testing Macs' support, maintenance, compatibility and reliability, we were convinced, after being initially reluctant, that this was the way to go. So it will be interesting to see what Gates has to say about the future of Microsoft in HPC."

###

HPC Wire
November 17, 2005

Features:

SC05 Deputy Chair Horner-Miller Waits Eagerly for Next Year

You'd never know it was six degrees below zero outside her office by the warm tones of Barbara Horner-Miller's voice. But that is the reality most November mornings at the Arctic Region Supercomputing Center, in Fairbanks, Alaska.

So it's not hard to imagine she's happy to be in the warm -- by any measure of comparison -- surroundings of Seattle this week at SC'05, where she is deputy chair.

And while most everyone in the HPC universe is focused on the current events of SC'05, Horner-Miller actually may have one eye looking ahead. After all, when the hoopla around this week's gathering is a memory, she knows her work will just be beginning.

You see, Horner-Miller, who serves as associate director at ARSC as well as director of user services for the facility, then must step up and start things in motion for SC'06, of which she is the chair.

Without putting SC'05's horse before SC'06's cart, Horner-Miller said her main goal for next year's event in Tampa Bay, Fl., is to make sure "it's a little bit different" from what everyone is experiencing in Seattle. First off, logistically, the move from the Pacific Northwest to the far southeast should assure at least some of differentiation. But after that, she added there needs to be some changes and additions to the conference program.

While attendees will see the HPC Analytics Challenge return in 2006, they must bid adieu to the StorCloud initiative after Seattle, Horner-Miller said. "We need to find some way to keep data in our program," she added. While just what will be a hot topic when formative meetings begin for SC'06, right now Horner-Miller is most enthusiastic about next year's debut of the Exotic Technologies initiative.

"We needed a way to bring in new, unfolding technologies and small companies that otherwise might not have a place in the conference," she explained. The idea actually has its origins in two separate factions that came directed to her at the same time. "We need to put the spotlight on things that are unfolding but are not yet necessarily part of even the leading edge systems," she added. "At the same time, we need to bring in small companies that may not even have prototypes, but one-offs, that could bring value to the HPC community over time."

The deciding factor for the new initiative, according to Horner-Miller, is how willing these fledging companies will be to showcase their wares before the clamor of the SC crowd. Nevertheless, she is encouraged by the fact she convinced none other than Bill Kramer -- this year's SC chair -- to head this new initiative. "It is an honor to get Bill to

do this," Horner-Miller said. "Typically, last year's chair would want the year off to recuperate, but Bill was excited to lead this new effort."

One fact regarding SC'06 that should be a challenge is that the exhibit space available in Tampa is smaller than Seattle's, she said. "We'll need to come up with some innovative ways to use the space," Horner-Miller added. While the deletion of StorCloud will open some room, clever thinking will be needed to maximize the area.

Horner-Miller also knows she will be under scrutiny to at least meet the booming demand for the event's technical program. "More registrations have been sold for SC'05 than SC'04," she reported. "This is wonderful news. We probably broke the 3,000 mark and we expect to continue this trend in Tampa."

As for this week in Seattle, Horner-Miller said she is taking lots of notes and observing the progress of the event -- and learning from it. "It's the only way not to get swamped for next year," she added.

As for this year's SC, Horner-Miller is leading a 17-person ARSC contingent at the show. While ARSC, a Department of Defense center that also is tied to the University of Alaska at Fairbanks, relies on Cray Inc. and IBM supercomputers as well as Sun Microsystems Inc. and StorageTek storage, she said ARSC will be looking at the entire gamut of products on display at SC'05. While Horner-Miller has been a strong participant in the direction of Cray Users Group since the early 1990s, she stressed the importance of taking in all the variety SC'05 has to offer.

ARSC's 20-by-20 foot booth at SC05 also offered variety. The space was built on four pillars that highlight the 50-person center's various activities. One pillar features general information, such as the current projects underway and the diversity of ARSC users; the next highlights ARSC's systems, such as "Klondike," its 128-processor Cray X1, "Iceberg," an 800-processor IBM system and "Iceflyer," a 32-processor IBM Regatta. The final two pillars were devoted to the StorCloud initiative, including ARSC's focus on its tsunami model and portal project. Various ARSC-penned papers on the matter also were available to attendees via DVDs at the booth.

###

HPC Wire
November 17, 2005

Breaking News:
Virginia Tech Extends HPC Challenge to College Students

A new high-end computing research challenge for students was announced this week at SC05. Patterned after the well-known Gordon Bell prizes presented annually at the SC conference, the goal of the Virginia Tech High-End Computing Challenge (VTHECC) is to encourage participation by graduate and undergraduate students from the US and throughout the world in high-performance computing.

General Chair Kirk Cameron (Virginia Tech) invites students to submit technical papers describing original, innovative approaches that increase efficiency and reduce time to solution of real parallel applications on real high-end computing systems. Topics of interest may include -- but are not limited to -- novel parallel or distributed algorithms, applications, architectures, systems, or middleware that improve state-of-the-art HEC applications and systems. Papers may focus on -- but are not limited to -- peak performance, price/performance, energy/performance, etc.

Submission deadline is July 10, 2006. A distinguished awards committee will select four winners who will be awarded cash prizes and be invited to participate in a research symposium, to be held at Virginia Tech in Fall 2006. For complete details, visit http://vthecc.cs.vt.edu.

HPC Wire
November 18, 2005

Breaking News:
Super Network Supported Supercomputing 2005

For the first time ever in a real-world environment, Pacific Northwest Gigapop (PNWGP) and its strategic partners have brought together more than one-half terabit per second (500 gigabits per second) of bandwidth in deploying SCinet, the very high performance network built to support Supercomputing 2005 in Seattle. The network is provisioned through multiple dark fiber strands brought by the University of Washington from the convention center to major telecommunications facilities in the city.

DWDM gear from Ciena, Cisco and Nortel were used to provision more than 50 10Gbps circuits and a native 40Gbps circuit. These circuits were then interconnected to numerous high-bandwidth national backbones, including National LambdaRail, CANARIE, Internet2's Abilene Network and UltraScience Net. International networks worked with these various North American facilities to reach the Seattle venue. In particular, Pacific Rim networks in Japan, Korea, Taiwan and Australia were able to utilize the Pacific Wave distributed peering exchange facility, a joint project between PNWGP and CENIC.

"As a direct result of many strategic investments by the University of Washington and the Pacific Northwest Gigapop, Seattle is one of the few places in the world where SC05 could benefit from an abundance of first-rate networking resources including metropolitan fiber, carrier-grade telecommunications facilities, a world-class engineering team, and an ever growing concentration of national and international networks," said Steve Corbato, director of network initiatives for Internet2.

"This staggering amount of bandwidth," he continued, "was deployed seamlessly and provides a truly impressive demonstration of the rapidly evolving suite of network capabilities in support of leading-edge computational science."

Among the many events relying on this bandwidth were massive storage- and dataretrieval tools, the Internet2 Land Speed Record attempts (IPv4 and IPv6), data grids, multipoint real-time, high-definition video from points around the world, super highdefinition video, and massive 3D imaging.

Professor Larry Smarr, director of the California Institute for Telecommunications and Information Technology (Calit2), a partnership of the University of California at San Diego and UC Irvine, and principal investigator of the National Science Foundation's OptIPuter project, offered this observation: "The Terabit Era has arrived! This unprecedented achievement of PNWGP and SC05 demonstrates that the United States needs to broaden its strategic technology leadership agenda from a focus on faster individual supercomputers to supernetwork-connected resources on a global scale."

BIO-IT World

SC05 Recap: Microsoft, Top500, Parties, and More

By Salvatore Salamone

November 22, 2005 | The SC05: Gateway to Discovery conference, held last week in Seattle, featured Microsoft's official entry into the high-performance computing (HPC) market, the release of a new Top500 supercomputer list confirming IBM's dominance in the field, over 250 exhibitors, about 10,000 attendees, excellent vendor parties, and a drink—the Super Computini—named in honor of the conference.

For anyone attending the conference, the most noticeable thing was the size of this event. Exhibitors filled the entire main hall and most of an upper level hall of the Washington State Convention and Trade Center. All told, there were 256 exhibitors. The mix included many of the academic supercomputer centers, most of the government research labs, and the entire gamut of HPC vendors.

The founders of this conference could never have imagined how big this event would become. Over time, its focus has shifted from the sole purview of the government labs and university supercomputing centers to include a wider cut of the scientific computing community. In fact, over the last few years, the show has expanded its scope from just supercomputing to a much broader high-performance computing emphasis, including HPC systems, networking, and high-performance storage. And this year, high-performance analytics was added as a new dedicated topic for the show.

Microsoft's Big Push For the past year, **Microsoft** has been talking about its take on HPC (see: Make Room for Microsoft).

But this year's conference was Microsoft's official push into the HPC market. Microsoft Chairman and Chief Software Architect Bill Gates delivered the opening morning keynote, the company released the beta 2 version of its HPC cluster software, and numerous software partners touted their support of Microsoft's efforts in a pavilion on the show floor.

The keynote had a bit of a rock concert atmosphere. A standing room-only crowd packed three connected ballrooms while concert sized monitors broadcast a short video on the history of scientific computing before the talk. (Talk about playing to an audience and whipping the crowd into a frenzy!)

In his address, titled "The Future of Computing in the Sciences," Gates shared his thoughts on how innovations in high-performance computing will improve collaboration and speed research developments by connecting scientists, computing resources, and data.

Gates noted that several industry megatrends are coming together to alter the way

research is conducted. These trends include increased processing power of HPC systems, service-oriented architectures that make it easier to share and access data, new computer form factors (Tablet PCs, PDAs, cell phones, etc.) that help collect data. and the digitalization of simulation and product design (i.e., molecules and drugs). "With computer modeling, we're figuring out what is happening in the real world," said Gates. Specifically, he noted that as scientists have more data and more-complex data to work with, researchers "will need new software to figure out what resources are needed to work with the data," said Gates. One area where he thinks a lot more work needs to be done is on visualization to provide insight into data and to get more value out of data.

Gates also spoke about the work Microsoft has been doing in the HPC arena (more on the Gates keynote.) The results of its efforts were also announced at this show.

Specifically, Microsoft announced the availability of the beta 2 version of Windows Compute Cluster 2003. The software comes on two CDs—one that contains Windows Server 2003 Compute Cluster Edition and a second CD with the Microsoft Compute Cluster Pack. The first CD includes the basic operating system; the second offers a job scheduler, MPI features, and other tools to run jobs on an HPC cluster. Microsoft broke the offering up in this way so that organizations that use third-party schedulers or MPI technology may continue to do so.

Like many other Microsoft efforts, the company will rely on partners and independent software vendors (ISVs) to reach particular markets like the life sciences market. To that end, partners in Microsoft's booth included many vendors familiar to the life science community, including **Accelrys**, the **BioTeam**, **MathWorks**, **Platform**, **Wolfram Research**. Microsoft also included the parallel molecular modeling dynamics code **NAMD** in the partners' exhibits.

Trends Galore On the show floor and in the conference sessions, some of the more interesting themes included discussions about personal supercomputing systems, acceleration technology, and the Top500 list.

Gates spoke a bit about the implications of having supercomputing power available from the desktop to the high-end cluster. Among the vendors targeting the personal supercomputing (or at least the desk-side supercomputing) market included **Orion MultiSystems** and **Penguin Computing**.

Orion was showing its DS-96, a 96-node cluster, that earlier this year was a finalist in the Bio-IT World Best of Show Awards program (for details see: Announcing the 2005 Best of Show Winners).

Penguin Computing unveiled its Penguin Application-Ready Cluster Portfolio, a line of workgroup or personal clusters that combines hardware and software. Essentially, the systems are HPC turnkey clusters designed for researchers with no IT support.

On the accelerator front, attendees packed a room for exhibitor presentations on new approaches to using dedicated hardware to accelerate HPC applications. Among the companies presenting included **CPU Technology**, **Mitrionics**, and **Synective Labs**. And on the show floor, systems vendors including **Cray** and **SGI** were showcasing their latest technologies to tap the processing power of field-programmable gate arrays (FPGAs).

And when it came to the session on the Top500 supercomputers, much of the drama was squelched as it became clear that IBM dominated the top end of the list. IBM systems accounted for the top three and five of the top ten systems on the list. In fact, the top two systems were both Blue Gene systems coming in with benchmarked performances, respectively, of 280.6 and 91.29 TeraFLOPS (trillion operations per second).

However, at the birds-of-a-feather session where the list was officially announced, the **Top500.org** committee discussed some ideas about getting a truer measure of performance than using the single Linpack benchmark.

"With Linpack you get the good, the bad, and the ugly," said Jack Dongarra, university distinguished professor in the department of computer science at the **University of Tennessee**. Dongarra is one member of the committee that puts together the Top500 list twice a year.

He noted that the good from using Linpack is that it is a single number and it allows historical comparisons (the Top500 list has been produced twice a year since 1993). The bad is that Linpack "emphasizes peak CPU speed and the number of CPUs, it doesn't [reflect or stress] local bandwidth or the network," said Dongarra. And the ugly is "benchmarketeering hype."

To overcome the limitations of Linpack, Dongarra and others are exploring benchmarks that might be used to complement it. For example, he noted the **HPC Challenge (HPPC) Benchmark**, which includes Linpack and six other benchmarks that stress non-CPU aspects of an HPC system.

Outside the Conference SC05 reached a stratum that IT veterans will appreciate. As was the case with the mega-conferences of the past like Networld, Interop, Networld+Interop, Comdex, and CeBIT, SC05 included some blockbuster vendor parties.

Two of the biggest events were a Microsoft and **Intel** party at the **Experience Music Project**, a museum of musical history founded by Microsoft co-founder Paul Allen. And the **Linux Networx** party at **Safeco Field**, the home of the Seattle Mariners.

The Microsoft and Intel event featured a live performance by Sheryl Crow. The Linux Networx party featured an on-field pitching setup complete with a radar gun to measure your fastest pitch and an under-the-stands batting cage. (Linux Networx also used the show to announce the LS Series, which is a new family of Linux Supersystems.)

And to show appreciation for the large number of guests in the local hotels, the master mixologists at the Westin Hotel Seattle created a special drink called the Super Computini just for the show. The ingredients of this martini-like drink were kept secret (actually no one bothered to ask for them), but word on the exhibit hall floor was that a couple of these drinks increased networking speed and improved collaboration.

(*Note: A roundup of products introduced at the show will be in a separate online story to appear later this week.*)

###

BIO-IT World

SC05 Keynote: Gates on Reducing Time to Insight and Discovery

By Salvatore Salamone, Bio-IT World

November 15, 2005 | Microsoft Chairman and Chief Software Architect Bill Gates kicked off this year's SC05 conference by sharing his thoughts on how innovations in high-performance computing (HPC) will improve collaboration and speed research developments by connecting scientists, computing resources, and data.

The talk was before a standing-room-only crowd that filled three connected ballrooms in the Washington State Convention and Trade Center in Seattle. And while Gates used the talk to announce the beta 2 release of Windows Compute Cluster Server 2003, the focus of his keynote was on the changes he expected in how researchers work, as new ways to use computing resources become available.

Gates noted that several industry megatrends are coming together to alter the way research is conducted. These include increased processing power of HPC systems, service-oriented architectures that make it easier to share and access data, new computer form factors (Tablet PCs, PDAs, cell phones, etc.) that help collect data, and the digitalization of product design and simulation. "With computer modeling, we're figuring out what is happening in the real world," said Gates.

As these trends give scientists more data and more-complex data to work with, researchers "will need new software to figure out what resources are needed to work with the data," said Gates. One area where he thinks a lot more work must be done is visualization - to provide insight into and get more value out of data. And he believes leveraging the megatrends is a way to reach these goals.

Workflows Rule At a conceptual level, Gates sees computational workflows that leverage the benefits of megatrends and integrate distinct tasks as the key to speeding insight and discovery.

This workflow includes computational modeling of the physical system being researched; acquisition of experimental data; searches and incorporation of persistent distributed data (such as libraries or public databases); and a software core that includes job scheduling, data mining, and analysis algorithms. All these elements produce information that would be interpreted by a researcher to get new insight into a problem.

While this approach is used in many organizations today, Gates thinks there are several opportunities to use newer technologies to improve the process. For instance, he believes personal supercomputers (systems with supercomputer power that cost less than $10,000, and that sit next to a researcher's desk) give researchers the ability to quickly develop and test new models.

Additionally he sees new resource management tools helping to provide a transparent way for organizations to selectively use computing resources without the scientists having to be IT managers. For example, when a scientist submits an analysis job, the job could be assigned to run on a personal supercomputer or a larger cluster based on pre-defined workload criteria.

Focus on Life Sciences Gates' talk focused on issues of importance to anyone working with large amounts of data, computational models, and research and development. But in many parts of the talk, Gates related how life scientists were using the methods he described.

For example, he called out Kyril Faenov, director of HPC at Microsoft, to demonstrate how to use these ideas in a research environment. Faenov's demo showed how a cancer research project was using an iterative data analysis technique to identify biomarkers and proteomics features in mass spectrometry data. In that demo, different analysis jobs were run on a local Linux cluster and a remote Windows HPC cluster.

Gates also spoke of an AIDS vaccine design project effort between Microsoft Research, the University of Washington, and the Fred Hutchinson Cancer Research Center. He noted the project was "the intersection of the advances in science and HPC."

He concluded noting that "the next 10 years will be really exciting, not just in [computing], but in science."

###

Bio-IT World
InfoWorld

SC05: IBM Holds on to Top500 Supercomputer Lead

By China Martens, IDG News Service

November 14, 2005 | IBM Corp. retained its lead of the Top500 list of supercomputers with its BlueGene/L System installed at the U.S. Department of Energy's (DOE's) Lawrence Livermore National Laboratories in Livermore, California. The system topped the twice-yearly list of the fastest computers in the world for the third consecutive time and is likely to remain number one for some time since its size doubled earlier this year.

The list, the twenty-sixth to be issued, was due to be announced Monday at the Supercomputing conference (SC05) taking place in Seattle through Friday.

There was some shakeup among the global top ten supercomputers with new entrants displacing some incumbents on June's list.

Cray Inc. notched up one new system and one revamped system, while IBM and Dell Inc. had one new system apiece in the top ten. Two IBM eServer Blue Gene systems on June's list dropped off the top ten -- the Blue Protein supercomputer at the Computational Biology Research Center in Japan and a Blue Gene machine at the Ecole Polytechnique Fédérale de Lausanne in Switzerland. Also exiting the top ten was the Thunder supercomputer at the Lawrence Livermore laboratory based on Intel Corp.'s Itanium 2 processors.

IBM's Blue Gene/L was measured with a Linpack benchmark performance of 280.6 teraflops. A teraflop is one trillion mathematical calculations per second.

In second position to the BlueGene/L was IBM's Watson Blue Gene (WBG) eServer system which the company installed at its Thomas J. Watson Research Center in June with a performance of 91.3 teraflops. New in at number three was the ASCI Purple system built by IBM and based on its pSeries 575 server, which is also installed at the Lawrence Livermore lab. It was measured at 63.4 teraflops.

Silicon Graphics Inc.'s Columbia system at the NASA (National Aeronautics and Space Administration) Ames Research Center in Moffet Field, California, slipped from June's third position to fourth place with 51.9 teraflops.

The DOE's Sandia National Laboratories had two newly installed systems at number five and six on the list, Thunderbird, a Dell Inc. PowerEdge-based system just edging ahead of Red Storm, a revamped Cray machine with maximum performances of 38.3 teraflops and 36.2 teraflops respectively.

NEC Corp.'s Earth Simulator in Yokohama, Japan, which topped the Top500 list for five consecutive times until IBM displaced it a year ago, was in seventh position with a performance of 35.9 teraflops, a slip from June's number four position.

In eighth position was another IBM machine, the fastest computer in Europe, the MareNostrum at the Barcelona Supercomputer Center in Spain, with a performance of 27.9 teraflops. The MareNostrum was number five on June's Top500 list. After that came another IBM eServer BlueGene machine, owned by Astron and run at the University of Groningen in the Netherlands with a performance of 27.4 teraflops. In tenth position was another new entry, a Cray machine at the DOE's Oak Ridge National Laboratory in the U.S. with a performance of 20.5 teraflops.

IBM had the most supercomputers on the list with 43.8 percent, followed by Hewlett-Packard Co., with 33.8 percent of all systems, though the latter had no showing in the top ten list.

Two-thirds of the Top500 machines are powered by Intel's chips, with 81 of the 333 systems using the chip giant's EM64T-based processors. IBM's Power chips appeared in 73 systems. Intel's rival Advanced Micro Devices Inc. (AMD) had processors in 55 systems, more than doubling its standing on the last list put out in June. Geographically speaking, the U.S. dominates the list, accounting for 305 of the Top500 supercomputers, trailed by Europe with 100 systems and Asia with 66 machines. Germany, which had been the leading European supercomputer country with 40 systems on June's Top500 list, only had 24 systems on the new list. with the U.K. becoming the number-one European player with 41 systems up from June's 32 systems.

The Top500 list was compiled by Erich Strohmaier and Horst Simon of the National Energy Research Scientific Computing Center (NERSC)/Lawrence Berkeley National Laboratory, Hans Meuer of the University of Mannheim, Germany, and Jack Dongarra of the University of Tennessee, Knoxville.

###

Jupiter Networks
November 15, 2005 06:09 AM US Eastern Timezone

Juniper Networks Routing Platforms Enable Core of SCinet at SC05; T-series Participating in SCinet High Performance Core Network for Fifth Consecutive Year

Supercomputing 2005
Booth #530

--(BUSINESS WIRE)--Juniper Networks, Inc. (NASDAQ:JNPR):
When: November 12 - November 18, 2005

Where: SC05, Washington State Convention and Trade Center,
Seattle, Wash.

For the fifth straight year, Juniper Networks, Inc. (NASDAQ:JNPR) T-series routing platforms will participate in the core of SCinet, a collection of high-speed networks built to support the annual International Conference for High Performance Computing, Networking and Storage (SC05). SCinet provides the infrastructure for both an extremely advanced high-capacity experimental network, and the high-performance, production-quality network supporting exhibitor's networking needs.

As part of the core of SCinet, Juniper Networks T-series routing platforms will help support all SC05 network traffic, including the supercomputer applications being run by researchers participating in the conference.

The T-series delivers the unique combination of advanced routing features, best-in-class IP/MPLS capabilities, and unmatched reliability and scale that makes it the ideal platform for both production networks and advanced research and experimental networks. Shipping since April 2002, over 1,000 T-series units are deployed in networks worldwide. With the multi-chassis TX Matrix technology, the T-series provides a seamless migration path to multi-terabit speeds.

For more information, please visit the Juniper Networks booth (#530) at SC05. In the booth, Juniper Networks will also be highlighting a range of core routing, security and application acceleration solutions.

About SC05

SC05, the premier international conference on high performance computing, networking and storage, will convene November 12-18, 2005 in Seattle. Under the theme, "Gateway to Discovery," SC05 will showcase how high performance computing, networking, storage and analytics lead to advances in research, education and commerce. Exciting technical and education programs, workshops, tutorials, an expanded exhibit area, demonstrations and many other activities await attendees. SC05 is the one place that attendees can see tomorrow's technology being used to solve world-class challenge problems. SC05 is sponsored by ACM and IEEE.

Lightwave

National LambdaRail to provide WAN connections for SC05 conference

November 10, 2005 Cypress, CA -- National LambdaRail's (NLR) network infrastructure will be used to provide the majority of bandwidth for next week's SC05 conference in Seattle, Washington. More than 240 Gbit/sec of networking capacity over 24 individual wavelengths of the NLR infrastructure will be used for the conference, including a 40-Gbit/sec capable wide-area link, provided in partnership with Cisco, specifically established for SC05.

NLR is a major initiative of U.S. research universities and private sector technology companies. SC05, the premier international conference on high performance computing, networking, and storage, will showcase how these technologies lead to advances in research, education, and commerce.

Network connectivity for the conference includes ten wavelengths being provided by NLR and its partner Cisco, as well as 14 NLR wavelengths already acquired and in use by CENIC and PNWGP (5); Internet2 (2); Oak Ridge National Laboratory (4); the U.S. Department of Energy's ESnet (1); and the University of California, San Diego and the University of Illinois, Chicago (2).

NLR says that because it owns the underlying optical fiber and networking equipment, it is able to flexibly provision connections for SC05 with more than 150,000x the capacity of typical home broadband, to support demonstrations in areas such as grid computing, visualization, and collaboration, as well as next generation networking development.

"We are pleased that the NLR infrastructure is the wide area networking choice for some of the most demanding applications being demonstrated at the SC conference this year," says Tom West, CEO of NLR. "NLR is committed to providing the research and education community the power and flexibility of next-generation IP and optical networking technologies, and the SC conference is an ideal opportunity to see these capabilities in action.

BusinessWire

November 09, 2005 09:50 AM US Eastern Timezone

SC05 Presents Special Video on Impact of Supercomputing to Society; Video to Premiere before Tuesday Keynote by Microsoft's Bill Gates

Supercomputing 2005

SEATTLE--(BUSINESS WIRE)--Nov. 9, 2005--SC05, the premier international conference on high performance computing, networking, storage and analysis, will make public a special video that explores the vital influence of supercomputing in solving the grand challenges of social and commercial problems affecting the quality of life for people throughout the world.

This informative video kicks off SC05's opening session on Tuesday, November 15, 8:30 a.m., at the Washington State Convention and Trade Center. A grant from ACM and IEEE Computer Society, cosponsors of the SC05 conference, allows supercomputing centers, research labs and higher education to freely use the video as one of their key educational pieces for years to come.

"Many people outside of high performance computing, networking, and storage do not realize the long process that is necessary from initial concept and prototype to when an application impacts our lives down the road," said Bill Kramer, SC05 General Chair. "Supercomputers play a critical role in jumpstarting many major applications that we have come to rely upon."

The video will explore the influence of supercomputing technology and research applications that ultimately arrive to the general consumer market as low-cost technologies.

"As a key component of information technology, supercomputing offers advances in natural disasters, transportation, environmental monitoring, and health and medicine," said David Patterson, ACM President. "We want this video to be used by educational and research institutions to broaden awareness and understanding of the essential role of supercomputing for computer science professionals, public policy officials and the public that benefits from these exciting advances."

The video tracks how the components of supercomputing - hardware, software, and algorithms - combine to produce world-altering discoveries. Extensive financial, labor, and time investment from government and industry produce supercomputing applications that "trickle-down" to commercial products and services. These applications result in major improvements in the production of automobiles, airplanes, pharmaceuticals, financial markets, and even animated movies.

Examples of initial supercomputing investments resulting in industrial and other major innovations include:

-- Flow analysis programs in the 1970s were created at national laboratories; by the early 1990s, these programs were being used for cell phone plastic molds

-- Vibration analysis (NASTRAN) was developed for government projects and is now available to all mechanical engineering firms to design airplane wings, auto bodies, and the Tacoma Narrows Bridge

-- A Cray-2 was loaned by the National Aeronautics and Space Administration (NASA) to the National Institutes of Health (NIH) to do the first analysis of X-ray defraction crystallography data of the AIDS Virus

-- By the early 70's, graphics-based finite element analysis software, such as LS-DYNA, was limited to expensive mainframe computers generally owned by the aeronautics, automotive, defense, and nuclear industries. Now, FEA is used for safer and cleaner cars (crash analysis), lighter milk cartons, and running shoe materials

Supercomputing algorithms have resulted in human genome mapping that detects birth defects and Alzheimer's disease. Even the design of sharper and longer-lasting disposable razor blades was achieved by using computational algorithms. With the aid of computer infrastructure models, researchers are able to accurately predict the effects of the recent hurricanes and plan for future natural disasters.

This video demonstrates how supercomputers have closed a gap between scientific theory and its application. Computer science professionals and their use of supercomputers have changed our world for the better.

About SC05

SC05, sponsored by ACM and IEEE Computer Society, is the premier international conference on high performance computing, networking and storage, will showcase how high performance computing, networking, storage and analytics lead to advances in research, education and commerce. The conference includes technical and education programs, workshops, tutorials, an exhibit area, demonstrations and hands-on learning. Visit us at http://sc05.supercomputing.org.

###

ComputerPartner

SC05 - Linux Networx shows off LS Series

14 nov 2005

Hardware and clustering company Linux Networx Inc. unveiled its LS Series of Linux machines Monday. The Series contains two families, the midrange LS-1 Supersystem and the high-end LS/X supercomputer.

The company is using "supersystem" to describe the midrange machine, because when many people hear the word "supercomputer," they typically associate it with a device with a million dollar price tag, according to Ben Passarelli, vice president of product marketing at Linux Networx. The LS-1 is designed so that much of the integration work involved in setting up cluster computing is done at the factory so users don't have to spend as much time on the task, he said in a phone interview Monday.

Linux Networx made the announcement at the Supercomputing conference (SC05) taking place in Seattle through Friday.

The LS-1 and LS/X are powered by Advanced Micro Devices Inc.'s (AMD's) Opteron processors. Linux Networx does plan to have an Intel Corp. version of the machine once the chip giant comes out next year with its Dempsey processor for two-socket servers, according to Passarelli.

The LS-1 is available in a desk-side model and a rack version as well as with a variety of different interconnects, he said. A desk-side LS-1 with 16 CPUs (central processing units) will cost around US$40,000, while an LS-1 with 64 processors will be priced about the $130,000 mark, Passarelli said. The system can scale up to 128 nodes, he added, and will be generally available in the first quarter of next year.

Linux Networx is making pretested subsystem modules for the LS-1 available in a number of areas including visualization, data storage and application acceleration, according to Passarelli. The company is currently working on Memphis, co-accelerators for a variety of hardware functions, including a floating-point accelerator (FPA), that are due to appear sometime in the first half of 2006, he said.

As for the LS/X, which has been under development at Linux Networx for the past two years, "We're aiming for extreme performance," Passarelli said. The supercomputer can scale up to 6,144 nodes and uses InfiniPath input/output technology from PathScale Inc. Linux Networx has developed an integrated switch infrastructure that improves the system's resilience and substantially cuts down on its cabling requirements by as much as 70 percent, he added. The LS/X will also become generally available in the first quarter of 2006.

Linux Networx was very pleased with its ranking on the Top500 twice-yearly list of the fastest supercomputers in the world, which was announced at the conference Monday,

according to Passarelli. The company has 16 systems in the current list or 3.2 percent market share as compared with only five systems or a one percent share on June's list.

"We're right there with Cray and SGI," Passarelli said, in terms of market share. Both Cray Inc. and Silicon Graphics Inc. (SGI) had 18 systems listed on the November Top500 list, equivalent to them each having a 3.6 percent market share.

###

IT and Informatics Weekly New Product Roundup

By Salvatore Salamone, Bio-IT World

November 23, 2005 | This week's roundup focuses on product news from the SC05 conference, held in Seattle last week. Most announcements are for high performance computing systems and software to run applications on HPC systems.

Microsoft announced the availability of the beta 2 version of Windows Compute Cluster 2003. The software comes on two CDs -- one that contains Windows Server 2003 Compute Cluster Edition and a second CD with the Microsoft Compute Cluster Pack. The first CD includes the basic operating system, the second offers a job scheduler, MPI features, and other tools to run jobs on an HPC cluster. Microsoft broke the offering up in this way so that organizations that use third party schedulers or MPI technology may continue to do so.

Microsoft will rely on partners and independent software vendors (ISVs) to bring Windows Compute Cluster 2003 to particular markets such as the life sciences market. To that end, partners in the Microsoft pavilion at SC05 included several vendors familiar to the life science community including **Accelrys**, the **BioTeam**, **MathWorks**, **Platform**, **Wolfram Research**. Microsoft also included the parallel molecular modeling dynamics code **NAMD** in its partners pavilion at SC05.

Penguin Computing unveiled Penguin Application-Ready Cluster Portfolio, a line of workgroup or personal clusters that combine hardware and pre-installed software. Essentially, the systems are turn-key HPC clusters designed for researchers with little or no IT support. The entire portfolio of products is available with single- and dual-core **AMD** Operteron processors or **Intel** EM64T processors

The Application-Ready Cluster Portfolio includes three product lines. The Penguin Portable Clusters include 6 to 24 CPUs and delivers up to 200 GigaFLOPS (200 billion operations per second) of processing power. The Penguin Performance Clusters can scale from 32 to 256 nodes and delivers up to three-quarters of a TeraFLOPS (trillion operations per second) of processing power. And the Penguin High Density Clusters support up to 240 CPUs and delivers up to 2 TeraFLOPS of processing power.

Linux Networx announced the LS Series, a family of Linux Supersystems. The series includes two product lines: the LS-1 and LS/X. Linux Networx says all the systems in the LS Series combine optimized performance and reduced total

cost of ownership.

The LS-1 Supersystems are pre-configured Linux clusters aimed at a wide range of scientific computational and simulation applications and they can be scaled from modest to high-end performance. The LS/X Supersystems are aimed at the high-end, which Linux Networx calls extreme performance. Linux Networx offers LS/X systems that deliver up to 100 TeraFLOPS of processing power.

Wolfram Research announced Mathematica Personal Grid Edition, software that combines the computational features of its flagship data analysis and visualization software Mathematica with high-level parallel language extensions that allow the application to be ported to run on four computing cores.

Specifically, Personal Grid Edition combines four Mathematica kernels, a Mathematica front-end, and a parallel computing toolkit. The software allows a researcher to quickly modify applications to run on four systems or cores producing up to 300 percent increase in performance. Wolfram says the applications can also be run unchanged on grids or clusters using gridMathematica.

SGI announced the SGI Altix 4000, a new platform in its Altix computing line. The Altix 4000 uses blade components, Intel Itanium2 processors, and SGI's NUMAflex shared memory architecture. The first model in the line, the Altix 4700, supports up to 512 processors and as much as 128 TB of shared memory. The system comes with a choice of **Novell** SUSE Linux Enterprise Server 9 or **Red Hat** Enterprise Linux Advanced Server 4 operating systems.

MathWorks announced Distributed Computing Toolbox 2, a new version of its programming tool for developing distributed computing applications. Enhancements in this release include support for third-party compute job schedulers and new inter-process communication features for distributing and executing parallel algorithms in a cluster of computers using MATLAB.

Interactive Supercomputing (ISC) announced ISC Star-P, software that allows scientists to write algorithms and create models on their desktop and then run them instantly and interactively on SGI Altix servers. Specifically, scientists can work with familiar mathematical software packages such as Mathworks MATLAB to develop data modeling and analysis routines as they normally would. Once done, the Star-P software takes the desktop application and automatically parallelizes the code so that it can run on any multi-processor system (such as the Altix) or parallel cluster. Star-P will ship by year's end.

Rackable Systems announced the C5100-SMP Server and the C2002 High Efficiency Series product line. The C5100-SMP is a 5U, 8 Intel Xeon processor server with 32 GB of memory. Rackable Systems touts it as an x86-based alternative to RISC-based SMP servers. The C5100-SMP is available now with a list price of $69,000. Dual-core and 16-processor systems are planned. The C2002 uses DC power and a number of design features to offer low power consumption and a thermally efficient system. The C2002 systems use dual-core AMD and Intel processors and are intended for use in high-density data centers.

Briefly…

Liquid Computing, a developer of scalable HPC servers, announced that Europe's Council for the Central Laboratory of Research Councils will participate in the beta testing of Liquid's Interconnect Driven Server (dubbed LiquidIQ)…

Appro showed its new AMD-based XtremeBlade systems, which were announced the week before the show.

(*Note: Next week's product roundup will be a holiday special featuring gifts and gadgets for scientists and IT professionals and their staffs.*)

###

New York Times

CNET News

Writing the fastest code, by hand, for fun

By John Markoff
November 27, 2005

SEATTLE--There was a time long ago when the word "computer" was a job description referring to the humans who performed the tedious mathematical calculations for huge military and engineering projects.

It is in the same sense that Kazushige Goto's business card says simply "high performance computing."

Goto, who is 37, might even be called the John Henry of the information age. But instead of competing against a steam drill, Goto, a research associate at the Texas Advanced Computing Center at the University of Texas at Austin, has bested the work of a powerful automated system and entire teams of software developers in producing programs that run the world's fastest supercomputers.

He has done it alone at his keyboard the old-fashioned way--by writing code that reorders, one at a time, the instructions given to microprocessor chips.

At one point recently, Goto's software--collections of programs called subroutines--dominated the rarefied machines competing for the title of the world's fastest supercomputer. In 2003 his handmade code was used by seven of the 10 fastest supercomputers. (The Japanese Earth Simulator, which was then the world's fastest machine, however, did not use his software.)

In the most recent ranking of supercomputers, IBM machines overtook a number of supercomputers using Goto's software to capture the top three spots in the fastest computer rankings. Still, the Goto Basic Linear Algebra Subroutines, or BLAS, as his programs are known, were used by four of the world's 11 fastest computers.

Goto has become a legend in the supercomputing community because of his solitary crusade. And he shows no signs of flagging in the contest to wring every ounce of computing speed from the world's fastest microprocessor chips.

But for all the acclaim he has received, Goto is a relative newcomer to the supercomputing field, having made his breakthrough about a decade ago.

"At first I didn't know anything," he said in an interview at the annual supercomputing conference held in Seattle in mid-November. "This was all trial and error, but now I have experience."

The value of his work goes far beyond setting speed records. Because his programs can more efficiently solve complex linear equations, they can offer better solutions to virtually every computational science and engineering problem. For example, the subroutines are used in simulation programs to model the flow of air over the surface of a plane or a car more precisely.

Chip versus the hand

One of Goto's principal rivals is a software project known as Atlas, created by a group of researchers working with Jack Dongarra, a computer scientist at the University of Tennessee. Atlas is an automated effort to find the most efficient way to solve linear algebra functions for specific microprocessors--a task that Goto does meticulously by hand.

Like chess-playing software, the Atlas project tries to overcome the shortcomings of different kinds of computer designs by systematically testing thousands of solutions for each chip to find the most efficient one for each type of microprocessor.

By contrast, Goto uses only a program called a software debugger that allows him to track how data moves among different components of a microprocessor.

He then reorganizes the individual software instructions so that his subroutines perform crucial algebraic functions more quickly to gain small amounts of processing speed from a specific type of computer chip.

Typically these are highly repetitive operations that can consume vast amounts of computing capacity. For example, one challenging type of calculation requires the microprocessor to multiply numbers from two tables stored in memory together.

Dongarra acknowledges that Goto's hand-tuned programs are more efficient and can still outperform Atlas.

"I tell them that if they want the fastest they should still turn to Goto," said Dongarra, who is one of the researchers who maintains the Top500 listing of the world's fastest-performing computers from a computing speed race held twice a year.

Writing the fastest code, by hand, for fun

Goto came to his passion for supercomputing almost by accident. Educated in power engineering at Waseda University in Tokyo, he worked as an employee of the Japanese Patent Office, doing research on early inventions like video recorders.

To help in his work, Goto purchased a Digital Equipment workstation based on the Alpha microprocessor in 1994 to perform a simulation.

But when it arrived he could not understand why it was performing so slowly. So he explored the Alpha's design to see where the performance bottlenecks were.

He later purchased a second Alpha-based computer and by rewriting the crucial subroutines was able to improve its performance to 78 percent of its theoretical peak calculating speed, up from 44 percent.

No formal training

Although he was not formally trained in computer or software design, he perfected his craft by learning from programmers on an Internet mailing list focusing on the Linux operating system for the Alpha chip. His curiosity quickly became a passion that he pursued in his free time and during his twice daily two-hour train commute between his job in Tokyo and his home in Kanagawa Prefecture.

"I would frequently work on these problems until midnight," he said. "I did it to relax."

As a teenager, Goto developed a passion for electronic design, building his own stereo equipment from the most basic components.

His current interest, he says, is not in the pure mathematics of the linear equations, but rather in finding clever ways to overcome the shortcomings of the architecture and internal organization of microprocessors that are used in every kind of computer, from hand-held devices to supercomputers.

Modern computers are organized to offer the programmer a hierarchical series of data storage areas that range from the computer's disk drive DRAM memory, as well as relatively small temporary memory areas called caches. Typically, the fastest memories are also the smallest.

One of the simplest ways to speed a program is to keep the calculation in the memory unit, which is closest to the microprocessor's calculating engine.

Every time the calculation engine is required to stop what it is doing to get new data from a more distant memory area, processing speed slows. But in some cases, keeping data in

Writing the fastest code, by hand, for fun

the closest memory cache may not be as efficient as keeping it in a larger cache that is farther away.

Robert A. van de Geijin, a computer scientist who works with Goto at the Texas Center, said that Goto's special skill was in the step-by-step reordering of software instructions to take the greatest advantage of the performance trade-offs offered by each type of chip.

"He combines both scientific insight and engineering skills," van de Geijin said.

They met in 2002 when Goto took a sabbatical from his job at the patent office to spend a year at the Texas center. (He has since resigned from the patent office.)

Once Goto arrived in Texas, he turned his attention to optimizing the speed of the Pentium 4 microprocessor. When computer scientists at the University at Buffalo added Goto BLAS to their Pentium-based supercomputer, the calculating power of the system jumped from 1.5 trillion to 2 trillion mathematical operations per second out of a theoretical limit of 3 trillion.

The increase was so astounding that the record keepers for supercomputing Top500 called the researchers in Buffalo because they did not think such a speed was credible.

"I teased them and suggested that the speed of light was faster in Buffalo than it was in Tennessee," van de Geijin recalled.

Recently there has been a quiet controversy around the Goto BLAS because Goto has been slow to offer his work as open-source software, the free model of software distribution.

Some programmers have suggested that Goto has not joined the open-source movement because he wants to protect his secrets and strategies from competitors.

That is not so, he said recently, noting that the Goto BLAS software is freely available for noncommercial use. And he said he was preparing an open-source version.

He said his next big challenge was to expose chip designers to his ideas to help speed their processors.

"Computer architects are stubborn," he observed. "They have their own ideas." His ideas on computing efficiency, he said, speak for themselves.

###

InfoWorld

Microsoft shows prominently at Supercomputing '05

Microsoft enters HPC fray, with a splashy booth and beta release of Windows Compute Cluster Server 2003

By Paul Venezia
November 21, 2005

Microsoft's (Profile, Products, Articles) presence at Supercomputing '05 in Seattle last week was unmistakable. The conference, which spotlights innovations in HPC (high-performance computing), featured more than 200 exhibitors, including HPC stalwarts such as Hewlett-Packard, **IBM** (Profile, Products, Articles), **SGI** (Profile, Products, Articles), and **Sun Microsystems** (Profile, Products, Articles). Despite the fact that Microsoft is not really associated with supercomputing, the company seemed bent on keeping a high profile at SC05: Microsoft was a major sponsor of the show, and Bill Gates was a keynote speaker. Also, Microsoft's booth was easily the largest on the floor, holding down the prime real estate at the middle of the show, requiring all attendees to pass through it as they wended their way between exhibit halls.

The main thrust of Microsoft's efforts at SC05 was for Windows Compute Cluster Server 2003, which is still in beta. Show attendees could grab a free time-bomb demo copy of Windows CCS Beta 2. Aside from that, the booth consisted of other vendors showing off various products loosely tied to clustering, running on some Windows variant.

Microsoft didn't seem so interested in talking to the press, and all that was readily available was a handout showcasing Microsoft's one day of workshops to be held at the show -- a handout that somehow got the name of the show wrong, dubbing it "Supercomputer '05."

Microsoft's CCS is aimed at those who feel they need the raw horsepower of clustering but want to stay on the Windows platform. Historically Windows has been a very tough sell in the HPC market, so it's safe to say that the established HPC software vendors don't have too much to worry about from Microsoft -- yet.

InfoWorld

StorCloud at SC05

By Paul Venecia
Novermber 14, 2005

The show hasn't even officially opened yet, and I just got a brief but enlightening tour of the StorCloud infrastructure. 650TB of storage from a dozen vendors, linked through SCInet via 10G, 1G, Infiniband, and FC to various clusters elsewhere on the show floor. Mostly driven by a mess of Cisco's biggest iron, it's quite impressive.

Some vendors are mapping luns directly via FC or iSCSI, some are leveraging NFS, but all that storage is live and available where it's needed. Some of the challenge events are using it, as are some of the other demos on the floor.

Very, very cool.

Posted by Paul Venezia on November 14, 2005 12:08 PM | TrackBack (0)

InfoWorld

SC05 and the low end of interconnects

By Paul Venecia
Novermber 13, 2005

I'm spending the weekend and Monday at SC05 in Seattle, taking in the latest in supercomputing concepts. Today, I've been focusing mostly on interconnects. The cluster interconnect talk by DK Panda proved quite interesting, focusing on Infiniband usage in HPC.

One very interesting item of note was an off-the-cuff discussion I had with a conference attendee. We were talking about the usefulness of Infiniband in smaller clusters given the rise of affordable gigE interconnects. Since it's a given that commodity servers already have gigE NICs, it makes more sense to move in that direction rather than Infiniband or Myrinet for some smaller HPC clusters.

While discussing this, he noted that his lab had tested 40 commodity gigabit switches for their performance in HPC applications, and the standout was the NetGear GSM7248 -- at least I believe it was that model. Apparently, in an effort to reduce cost within the 7248, NetGear designed the switch with only two 24-port gigE ASICs, rather than the 8- or 4-port ASICs found in most other switches. Given the lack of interconnect requirements between ASICs on the switch backplane for each 24-port block, the HPC latency and performance numbers were apparently outstanding as long as you didn't cross the ASICs and only used the first or second 24 ports. The inter-ASIC communication introduced significant latency increases. For roughly $1,300, it became the switch they used to build a good-size HPC cluster on gigE.

Posted by Paul Venezia on November 13, 2005 10:49 AM | TrackBack (0)

InformationWeek

Supersized Challenge

By Aaron Ricadela, InformationWeek Nov. 21, 2005 URL:
http://www.informationweek.com/story/showArticle.jhtml?articleID=174400167

Microsoft is an afterthought in supercomputing. Changing that will take overcoming Linux--and recruiting a new breed of employee.

The most powerful lessons in life sometimes draw from long-ago experiences. So it may have been when Bill Gates reminisced last week about bypassing Harvard's computer-science classes 30 years ago. "I was taking physiological psychology and economics," the school's most famous nongraduate recalled. "If you look at my course sign-up, you wouldn't think I was a software person at all."

Now, like its founder back then, Microsoft is taking a more-interdisciplinary approach. Broad changes in computing, science, and education are compelling the company to look beyond personal and business computing for its next advances. One of the most important is the rarified stratum of supercomputing, where its Windows operating system is a nonfactor today. Gates predicts the kind of supercomputing server clusters used by a small but growing number of companies will become a foundation of business computing, helping drive the next breakthroughs in science, medicine, product design, and finance.

To even get its foot in the door, however, Microsoft needs better products to unseat a well-entrenched Linux operating system. And perhaps more difficult, it needs new talent. The explosion of fields like genomics and nanotechnology is interwoven with leaps in software, which means many of today's most-computer-savvy graduates major in biology, physics, or chemistry--fields where Microsoft isn't seen as a prestige employer. Those realities will put Gates' interdisciplinary approach to the test.

In a keynote at the SC05 supercomputing conference and a subsequent interview in Seattle last week, Gates laid out his plan for entering the scientific-computing market. Microsoft will release a special version of Windows; launch a multimillion-dollar outreach program to harness the high-performance computing IQ of professors at 10 universities in the United States, Germany, China, and elsewhere; and increase the degree to which its 600-plus researchers collaborate with scientists in fields such as AIDS research, astronomy, and oceanography. Given the expanded use of scientific computing in industries such as aerospace, auto engineering, pharmaceuticals, medical imaging--even consumer-product design--the effects could be widespread.

"When we say science, think about people designing cars, think about people designing planes, think about people thinking through the design of a Web site," Gates said. "It's not just new medicines, although that alone would justify all this work. It's not just modeling

the environment, although that alone is a supercritical thing that we absolutely need to do."

Microsoft last week issued a second beta-test release of a new version of Windows designed for scientific workloads due in the first half of next year. The 64-bit operating system, Windows Compute Cluster Server 2003, is designed to distribute computing across clusters of as many as 100 or 200 machines--relatively small potatoes in the supercomputing market but far more than the company's software has been capable of. Microsoft says it will be priced lower than the business edition of Windows Server 2003, which runs $800 in volume.

The company also is designing new capabilities for its graphical development tools that could help scientists program those clusters. "Scientists have all this data but can't really bring it together and get insights into it," Gates said. "They spend a lot of their time not thinking deep scientific thoughts but rather re-entering data and writing code that they shouldn't have to."

The Linux Challenge With clustering technology, users can chain together dozens or hundreds of inexpensive PC servers with special cabling to run software that distributes work among the processors. The approach lets universities, supercomputing research centers, and corporate IT departments access supercomputing power at much lower costs than specialized architectures have allowed. So far, though, cluster installations have been dominated by Linux, since many of supercomputing's practitioners cut their teeth in the shareware world of Unix. It's a group that's not particularly inclined to use Windows in the same way. "There aren't a lot of people who like Linux and like Windows," says Pete Bradley, an associate fellow for high-intensity computing at Pratt & Whitney, a United Technologies Corp. division.

The maker of commercial and military aircraft engines, a longtime supercomputing user, has added to its arsenal a network of 2,050 engineering workstations running Windows for message-passing applications when the machines sit idle at night. It also has a 64-node Windows cluster at its East Hartford, Conn., headquarters and a small Windows cluster in Seattle. But much of its staff expertise in high-performance computing is in Linux. "Microsoft sees this rapidly growing segment that's almost all Linux," Bradley says. "As cluster computing is becoming mainstream, this is clearly an area Microsoft is going to want to play in."

'Trickle-down' computing Sales of technical computing systems grew more than 30% last year, to $7.25 billion--faster than revenue growth in personal and business computing, according to research firm IDC. Sales of clusters are growing even faster, up 96% last year. And technical systems account for about 16% of all servers sold worldwide.

Supersized Challenge

Microsoft is betting that growth will ratchet up even more as businesses' research-and-development arms take advantage of high-performance computing techniques borrowed from university labs. Procter & Gamble Co. uses supercomputers to design packages for Folgers coffee, Pringles potato chips, and Pampers diapers. Rolls-Royce plc uses high-performance computers to track data from sensors installed on aircraft engines, then analyzes it to predict if faults might arise. And new streams of data are coming from satellites, environmental sensors, industrial radio-frequency identification tags, and medical devices, all of which could feed a market for software that can house and analyze it.

"There's this new era of hybrid computing," says Steve Wallach, a longtime supercomputer designer and adviser to CenterPoint Ventures, a venture-capital firm. The importance of data mining, visualization, and pattern recognition for business applications means commercial and technical computing markets are using each other's tricks. "You need thousands of processors," he says. "You have to access terabytes if not petabytes of data. And the computing isn't initiated by hitting carriage return. It's happening in the background."

By tapping into the world's high-performance computing expertise, Microsoft is positioning itself and its customers for an era in which the computing power resident in today's high-end systems becomes commonplace in workgroups or even on the desk. It's a problem as well as an opportunity. As chip designers run up against physical limitations to build ever-faster microprocessors, engineers at Advanced Micro Devices, IBM, Intel, and Sun Microsystems are packaging two or four processors on each chip to boost performance. But that makes it harder to program software that takes advantage of those chips' design. The problem will only get worse as the industry moves toward perhaps 64 processing cores on a chip by 2012, and hundreds by 2015. To prepare its products for those scenarios, Microsoft will need the supercomputing sector's know-how; those researchers have been distributing algorithms across scores of chips for decades.

"There's a trickle-down effect," says Jack Dongarra, a computer-science professor at the University of Tennessee and Oak Ridge National Laboratory. "What we would consider supercomputing today eventually filters down into commercial and home machines. Think of the high-performance computing market for Microsoft as an incubator for ideas that will eventually become mainstream."

New Breed Of Talent To speed up the incubation process, Microsoft is funding 10 Institutes for High-Performance Computing, at a cost of several hundred thousand dollars each, at the universities of Tennessee, Utah, Virginia, Washington, and Cornell; Stuttgart University in Germany; and Shanghai Jiao Tong University in China, among others. In exchange for the grants, professors will supply Microsoft with design help and feedback on its products. Getting respected industry figures such as Dongarra on board could boost Microsoft's credibility among a skeptical audience of supercomputing users.

Supersized Challenge

The company also is expanding its on-staff talent with supercomputing ties. In June, Microsoft hired Tony Hey, a former director of the United Kingdom's scientific computing initiative and a fellow of Britain's Royal Academy of Engineering, as corporate VP for technical computing. This month, Fabrizio Gagliardi, a former top technical manager at the European particle physics institute CERN, joined the company. Microsoft veterans with experience in high-end computing include Craig Mundie, senior VP and chief technical officer for advanced strategies and policy, who worked at Data General in the '70s and founded supercomputer company Alliant Computer Systems in '82; Gordon Bell, a former executive at Digital Equipment Corp. instrumental in developing the company's PDP and VAX machines; and researcher Jim Gray, a winner of computing's prestigious Turing Award.

But the star power also reveals a weak link: Microsoft attracts computer people. It needs to work more closely with scientists in disciplines outside of computing, as Mundie and Gates apply more of the company's brainpower to broad societal problems in science, medicine, and engineering. "The only raw material you feed into Microsoft is smart people," Mundie says. "They come from all walks of life, not just computer science. In the R&D part of the company, we're still overweighted toward people who come out of computer-science programs."

Microsoft traditionally has funded university research through grants to computer-science professors. That helps attract graduate students to projects and potential employees to Microsoft. Now Mundie wants to put company resources where they can benefit projects in broader fields. Gray, an authority on high-performance databases, has worked with astronomers for several years to virtually connect their databases of telescope observations using Web services into an online storehouse that can search for 100 million celestial objects. Another Microsoft researcher, physician David Heckerman, recently applied machine learning and data-visualization technology developed at Microsoft to AIDS research at the University of Washington.

Top grads in physics, chemistry, and biology are expected to be adept computer programmers, Mundie says, but Microsoft is short on connections to attract them. New types of collaboration could net a broader swath of the "intellectual elite," he says. "We have to look more broadly at people who are knowledgeable users."

Supercomputing Impact Microsoft isn't the first company to pursue low-cost supercomputing. Intel in the early '80s introduced a "personal supercomputer" that failed in the market. And the rise of symmetric multiprocessing led to a number of "mini-supercomputer" companies, including Mundie's Alliant, which never took off.

Today, the world's fastest machine is IBM's Blue Gene/L, installed at Lawrence Livermore National Laboratory to help maintain the country's nuclear weapons stockpile. Blue Gene/L reinforced its lead with a top speed of 280.6 trillion calculations per second,

according to a new list of the world's 500 fastest supercomputers released by Tennessee's Dongarra and other professors last week. Blue Gene runs a version of Linux on cells of special embedded, low-powered processors and uses ultrafast memory connections to achieve its speed.

To Dave Turek, VP of deep computing at IBM, one virtue of Blue Gene--there are 19 of the systems in the Top 500--is its ability to grow or shrink in size as customers need. "Blue Gene/L is a pretty effective price/ performance machine," he says. "From a design perspective, we don't present customers with discontinuities as they scale up or down. You don't suddenly need to put in new software or change your application."

Still, the high-performance segment is largely characterized by customers' ability to choose operating systems, middleware, and applications from different vendors to achieve the best performance. Microsoft's entry "has the potential to change the dynamic of the market in some ways," Turek says.

Smaller Installations Among the most powerful systems, Windows is nearly absent. According to the latest Top 500 list, only a system at Cornell University, which Microsoft heavily funds to use its products, runs Windows. Changing that could take years. Granted, the Top 500 is a market segment more powerful than what Microsoft is targeting with Windows Compute Cluster, but more entries could lend prestige.

For now, Microsoft is aiming for much smaller installations--clusters of fewer than 200 machines. That's the fastest-growing part of the market and the most profitable. "Microsoft is being very strategic and going after the sector of the market that pays margin," says Scott Studham, CIO at Oak Ridge National Lab.

Part of the appeal of Windows for technical clusters would be the ability to rapidly prototype mathematical models on a desktop PC or notebook, then run a simulation on a powerful cluster without porting the code. Merck & Co. is testing Windows Compute Cluster on 20 machines running The MathWorks Inc.'s widely used Matlab program. Eric Schadt, a senior scientific director for genetics, says the drug-design industry's frantic growth can strain software-development cycles. "The biology field is exploding," Schadt says. "We're constantly developing new algorithms. When we find something that works, we don't want to wait another six months to get that prototype running on Linux. We'd like to go directly from prototype to running it in the high-performance computing environment."

Microsoft is raising the ease-of-use quotient by building into its new version of Windows point-and-click management tools that would let a biologist configure a cluster, and a version of the popular MPI middleware for clusters the company developed with Argonne National Laboratory. Also in the works are visual-development environments for scientists and extensions to Office apps for scientific note-taking or outsourcing Excel

calculations to remote systems. "We can make this technology available to more people," Mundie contends. "Today, there's too high a degree of wizardry."

At the Seattle conference, Microsoft demonstrated a prototype "personal cluster" of four machines running at 25 gigaflops--equivalent to the power of about 10 PCs--that it says could fit under a desk and sell for less than $10,000. Microsoft showed the workgroup system running a rough series of genomics computations, then offloading them to a larger system in another city for more detailed analysis. That kind of wide-area supercomputing could have business appeal. "If you can let somebody submit a job onto the Internet and find the cheapest place to run it, that's not just interesting for scientific cluster computing, that's interesting for business computing," Gates said. Rare types of analyses or jobs that run remotely in a disaster-recovery scenario could benefit from the resource-sharing work Microsoft is developing. "We'll get plenty of benefits from these advances in the business realm," he added.

Where will these trends leave the industry? Probably with unimaginable computing power at users' disposal within a decade--or maybe two. The IBM chip inside Sony Corp.'s PlayStation video game uses vector-processing techniques that were a hallmark of Cray Research supercomputers in the '80s, and the work done by today's clusters could run on a single chip in 15 years. That could usher in simpler ways to interact with technology. Exponential increases in computing power could yield computer vision and speech interfaces that work in daily life or more intelligent filters for what information crosses our desktops.

Microsoft won't pull off any of this unless it can feed its voracious need for new talent--and expand the pool of people eager to work there. "We're always on the lookout for somebody who loves software but knows it so well they're seeing how it can be applied in different ways," Gates said. But declining enrollments of computer-science grads could hinder Microsoft down the road. So could difficulties recruiting from the hard sciences. Yet if Gates and Mundie are right, and the future of everything from medicine to finance hinges on hyper-talented people working with better software, integrating these fields could help us all. Said Gates, "The fact that we need software understanding to advance the sciences means the shortage is all the more acute."

###