Second Annual High Performance Computing Users Conference Report

Accelerating Innovation for Prosperity

July 13, 2005

Washington, D.C.

Council on Competitiveness

SPONSORS

The Council on Competitiveness thanks the following co-sponsors for their generous support of the Second Annual High Performance Computing Users Conference:

Defense Advanced Research Projects Agency

U.S. Department Of Energy's National Nuclear Security Administration

U.S. Department Of Energy's Office of Science

National Science Foundation

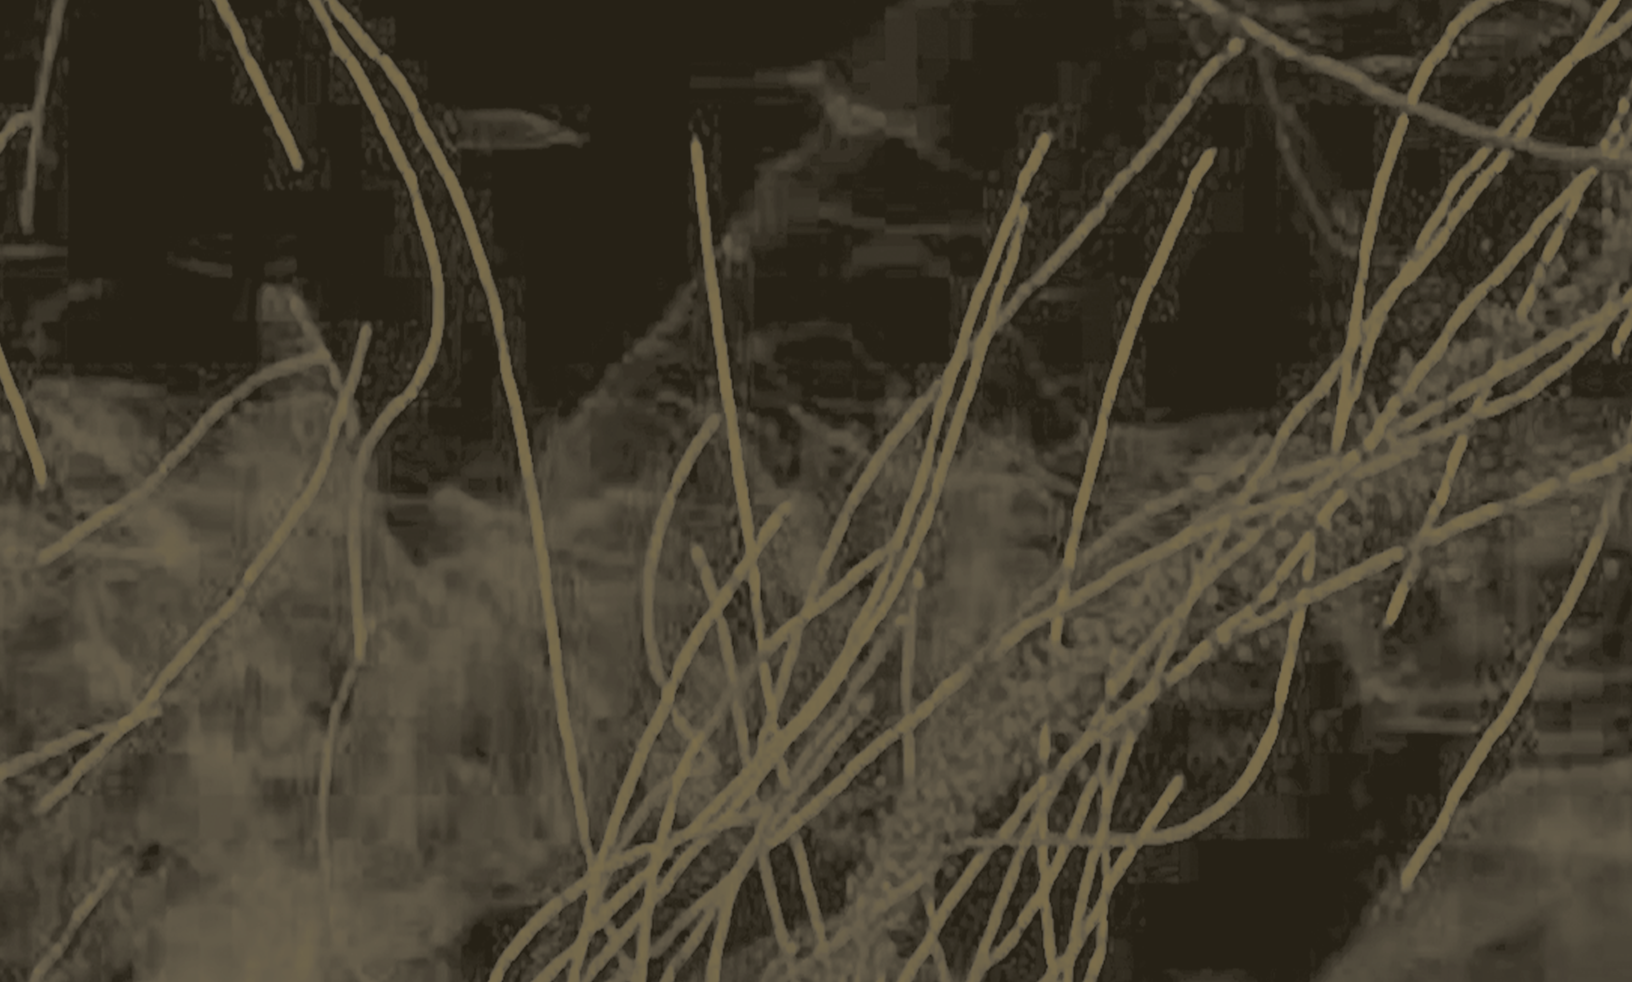

ISBN: 1-889866-28-8

This report can be ordered from the Council on Competitiveness.
To order and/or for more information, please write:

Council on Competitiveness
Publications Office
1500 K Street, NW
Suite 850
Washington, D.C. 20005
Tel: (202) 682-4292
Fax: (202) 682-5150
Email: council@compete.org
Or visit our home page at www.compete.org

The Council on Competitiveness is a nonprofit, 501(c)(3) organization as recognized by the U.S. Internal Revenue Service. The Council's activities are funded by contributions from its members, foundations, and project sponsors.

Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Design: Soulellis Studio

Cover Illustration: Tom White, tomwhite.images-illustration & design

This work was supported by the Defense Advanced Research Projects Agency, under the High Productivity Computing Systems Program, and the U.S. Department of Energy, under grant DE-FG26-04NT42101.

Printed in the United States of America.

HPC simulation of a tornado. Image courtesy of Silicon Graphics, Inc.

TABLE OF CONTENTS

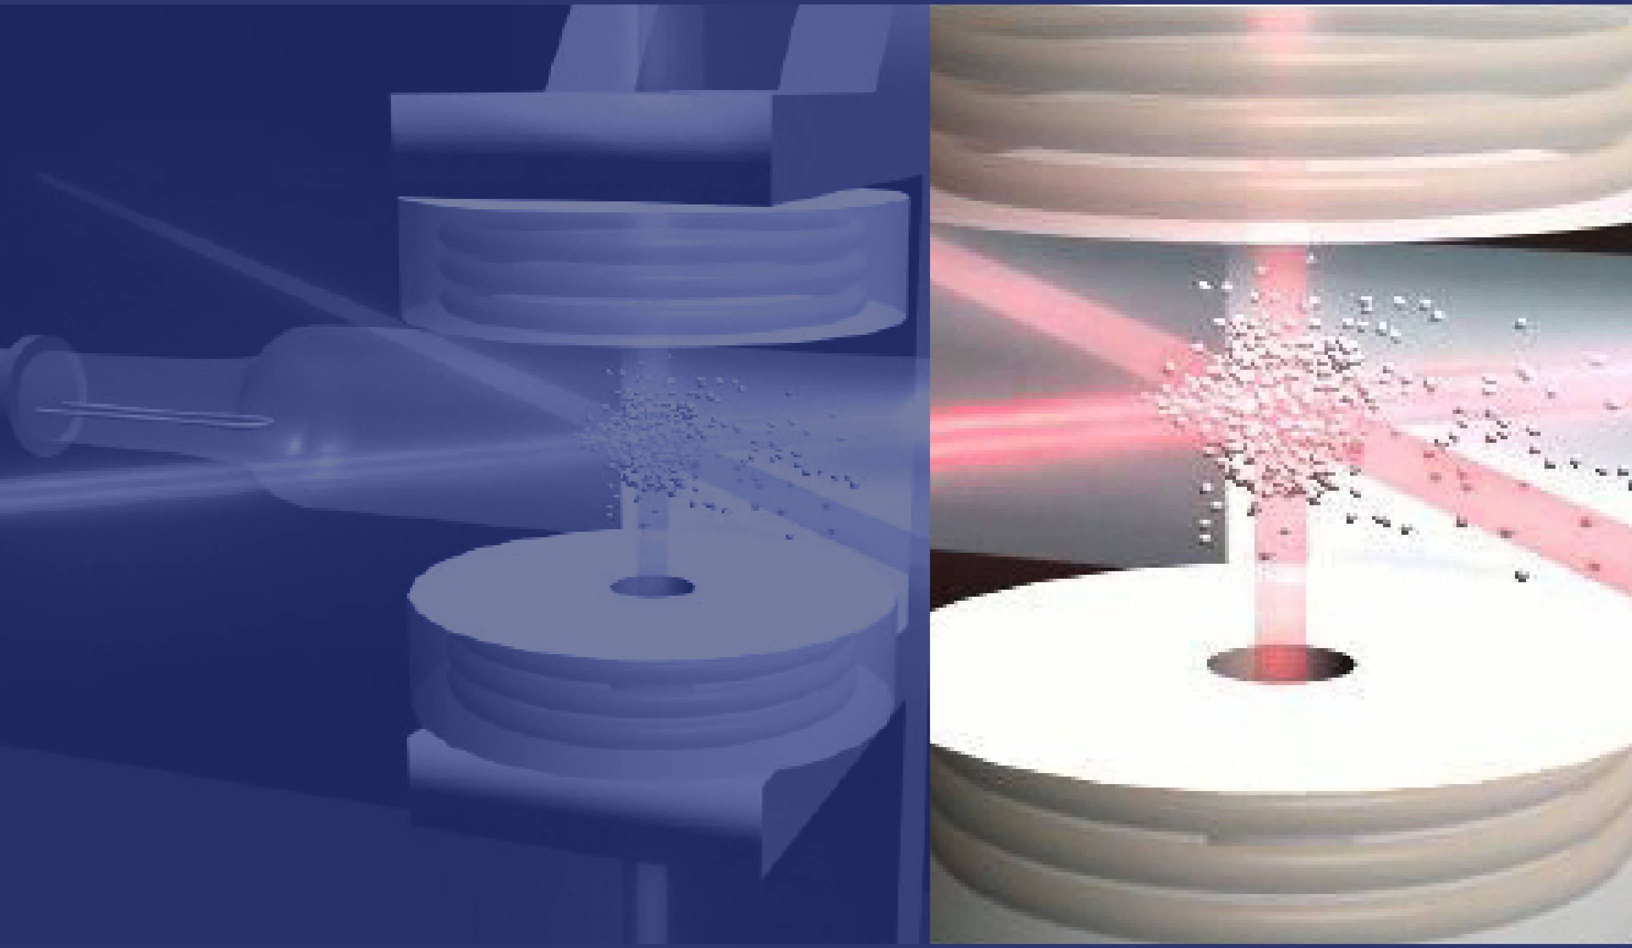

HPC simulation of atom trapping. Image courtesy of the Ohio Supercomputer Center.

PROCEEDINGS

EXECUTIVE SUMMARY

Today, America finds itself at an inflection point, shaped by unprecedented shifts in the nature of global competition, and in the nature of innovation itself. The world is becoming highly interconnected, and competitive and economic interdependencies are growing. At the same time, where, how and why innovation occurs are in flux — across geography and industries, in speed and scope of impact, and even in terms of who is innovating. In many ways, the playing field is leveling, and the barriers to innovation are falling.

In a world where many nations have embraced market economies and can compete on traditional cost and quality terms, it is innovation — the ability to create new value — that will confer a competitive advantage in the 21st century. And there are few areas of technology that hold more promise for stimulating innovation and propelling competitiveness than High Performance Computing (HPC).

HPC has been and will continue to be a key ingredient in America's innovation capacity. It turbo-charges the innovation process by shrinking "time-to-insight" and "time-to-solution" for both discovery and invention. Along with theory and experimentation, modeling and simulation with HPC has become the third leg of science and an important path to competitive advantage.

In order to better understand the potential of HPC as well as why industry is not using HPC as aggressively as it could be, the Council on Competitiveness hosted the Second Annual High Performance Computing Users Conference in Washington, D.C., on July 13, 2005. More than 200 senior government, business and academic HPC users and policy makers came together to explore the potential for this technology, and better understand the lack of application software that industry needs to fully exploit HPC for competitive advantage.

The conference began with the unveiling of "HPC in Everyday Life," an 8-minute video narrated by the penguins from the hit movie *Madagascar.* The video, created by DreamWorks Animation SKG, Inc. in collaboration with the Council, depicts how HPC is behind many of the products and services that we take for granted in everyday life. From

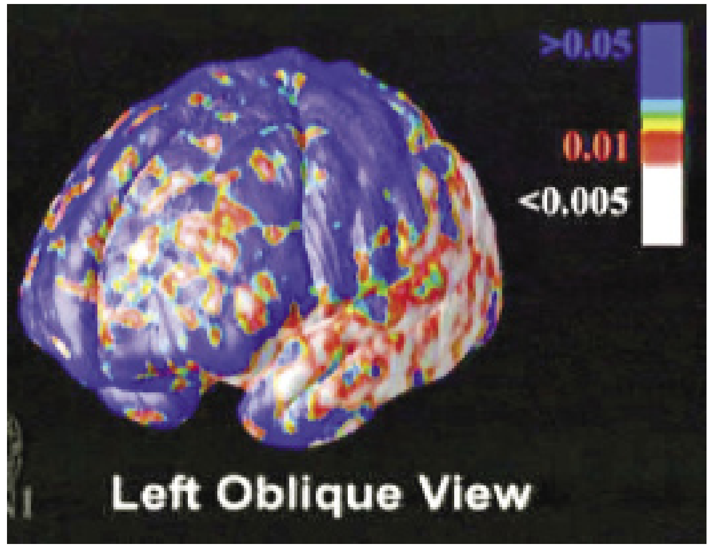

HPC simulation of Alzheimer's disease spreading through the brain of a living patient. Image courtesy of Silicon Graphics, Inc.

medicine to consumer products, and energy security to aerospace, this video illustrates the importance of HPC in breakthrough innovation, and demonstrates the need for continued support.

Following the video, the keynote address explored the crucial role of HPC in the consumer products and computer animation industries, where HPC accelerates creativity, discovery and invention. A call to make "big changes to big things" was reinforced by the first panel, which promoted the economic and societal benefits of solving "Grand Challenges" in fields as diverse as brain research, energy production and management, semiconductor production and video game creation. Their comments underscored findings from the Council's *2004 Survey of Industrial HPC Users*, that "extreme innovation" cannot take place without easy-to-use, production quality HPC application software to take advantage of the enhanced performance of powerful HPC systems.

As a result of that survey, the Council commissioned a *Study of Independent Software Vendors (ISVs) Serving the High Performance Computing Market*, the results of which were released during the conference. The survey explored the reasons behind the lack of production quality application software. Without additional funding, computer resources and a more compelling business case, ISVs are unlikely to invest heavily to develop scalable application software for the HPC market. Without more advanced application software, U.S. businesses will find it difficult, if not impossible, to address their large, complex, and most competitively important problems.

The afternoon panel explored options for bridging the chasm between the software we have and the software we need to drive more aggressive use of HPC across all industries. Panelists included software vendors, national laboratories and industrial HPC users. They offered perspectives on the challenges of using, maintaining, and creating application software, and on the role of universities and national laboratories in helping to accelerate the development of new and/or updated code.

From this discussion, three significant themes emerged:

- Although advanced HPC software is being developed at the national laboratories, it is not readily transferable to commercial users;
- User dependence on legacy systems and their integration into key business processes can slow the development of new application software;
- Some users must develop their own HPC application software in order to gain a strategic, competitive advantage in the market.

Three recommendations were suggested to address these issues:

- More partnerships among HPC stakeholders would accelerate HPC software development;
- ISVs should consider alternate business models to expand the user base and accelerate HPC software development;
- Federal funding for HPC should be balanced, pushing the top end of technology while expanding the usage within the federal government.

Most significantly, there was strong consensus that HPC is tied to U.S. economic competitiveness, and that participants must continue to explore mechanisms that will sustain the long-term health of the HPC market. In order to meet the competitive challenges of the 21st century, our country and our companies must run faster, be increasingly nimble, take risks, and occasionally take a leap into the unknown. Innovation will demand a new threshold of creativity, insight and invention that can be advanced only through HPC. We are just beginning to tap the potential competitiveness benefits of this promising technology.

In this increasingly competitive global environment, we believe that the country that out-computes will be the one that out-competes.

David E. Shaw
Chairman
D.E. Shaw & Co., Inc.

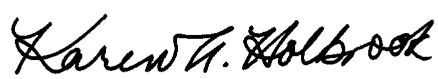

Karen A. Holbrook
President
The Ohio State University

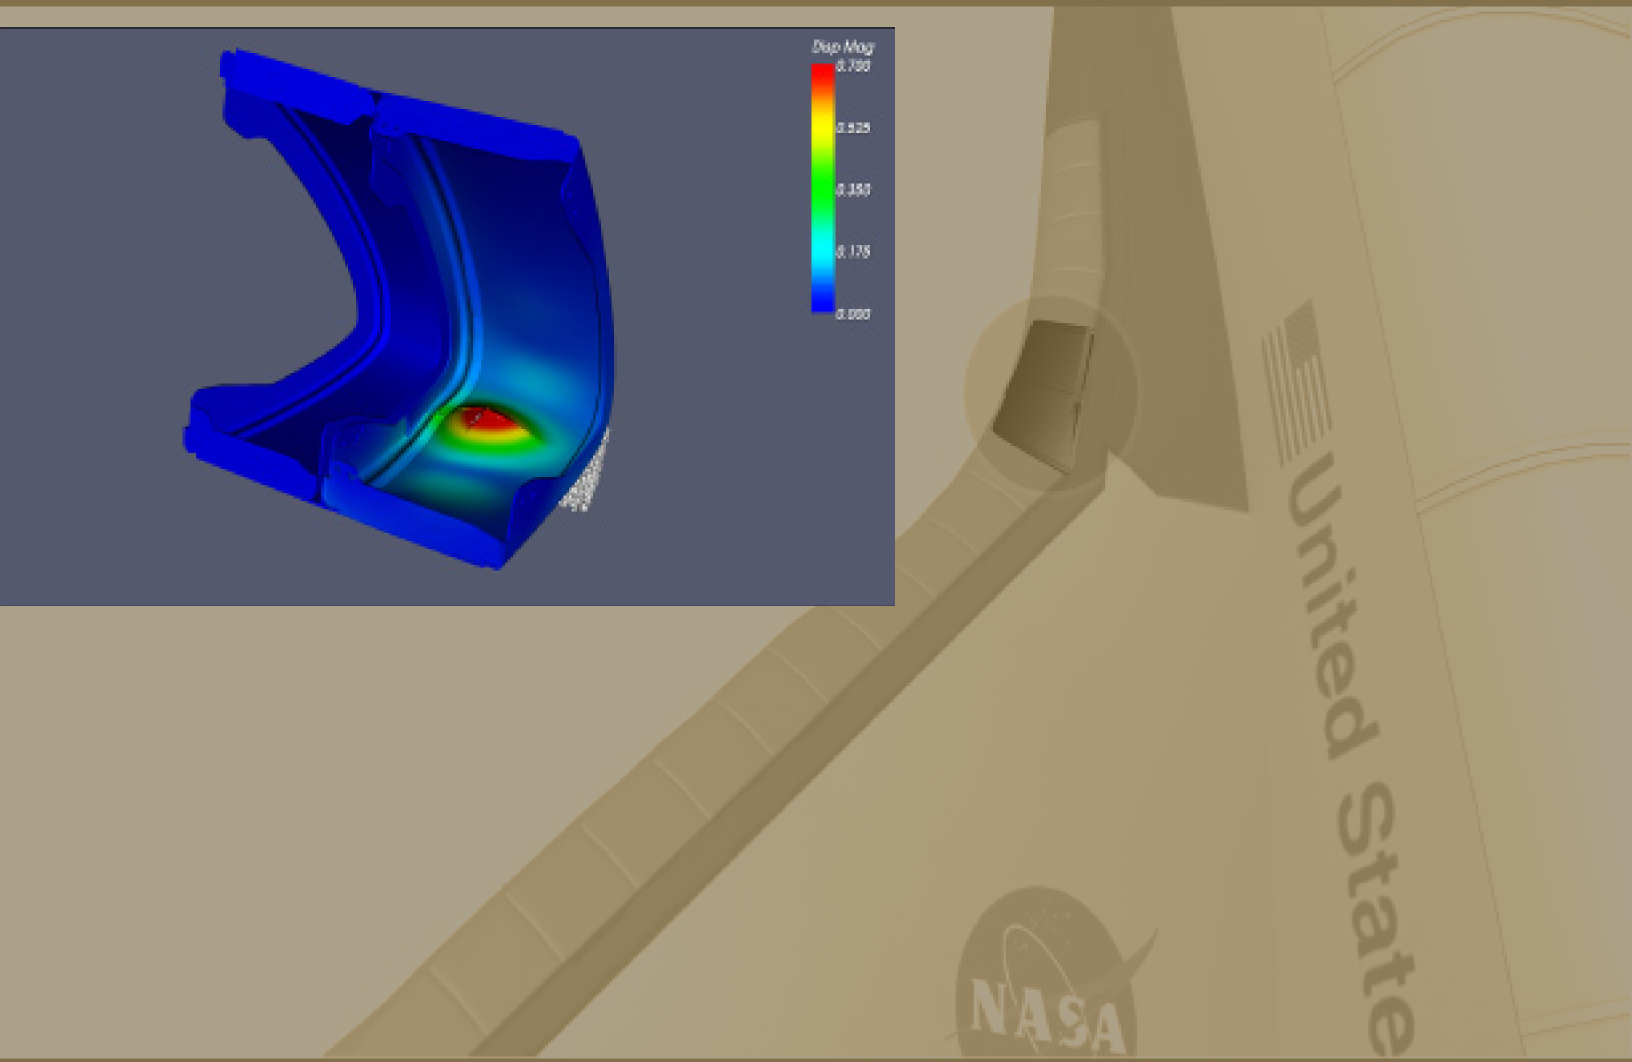

HPC simulation of the leading edge of the wing of the space shuttle Columbia, and the significant panel damage that could occur from foam impacts. Such simulations, which were later confirmed by full-scale impact tests, were critical to verifying the cause of the shuttle accident. Images courtesy of Sandia National Laboratories.

CONFERENCE PROCEEDINGS

High Performance Computing: The Key to U.S. Competitiveness

Council on Competitiveness President Deborah L. Wince-Smith opened the Second Annual HPC Users Conference with a call to action. Reminding those in attendance that HPC will be key to U.S. innovation and global competitiveness, she called on conference participants to identify ways to accelerate the adoption and use of advanced computation across the frontiers of science, in order to spur the creation of cutting-edge products, processes and indeed whole new industries.

The competitive pressure from overseas is mounting. U.S. corporations increasingly find themselves in fierce competition with companies and innovation centers that were not a threat even ten years ago. The global marketplace is becoming a more level playing field, with organizations becoming more interdependent every day. Emerging economies now compete both in low-wage, low-skill markets and in high-wage, high-value-add technology markets. Companies are poised to take advantage of these opportunities in an increasingly interconnected world. For advanced countries like the United States, it is innovation — the ability to create new value — which will be the single most important factor in determining competitive advantage in the 21st century.

HPC is a powerful innovation tool - one that America's competitors are pursuing to strengthen their own innovation capacity. The Japanese automakers, for example, recently have succeeded in using the "Earth Simulator" supercomputer to run a car crash simulation that produced details previously unobtainable. The Earth Simulator is Japan's most powerful supercomputer and one of the most powerful in the world. The Japanese plan to extend the simulation by adding vibration, noise, fuel consumption and other factors critical to car design. Integrating these into one computational model, along with the car crash variables, may require an HPC system even larger than the Earth Simulator[1]. If the Japanese automotive companies are able to run these highly integrated models successfully, they could begin producing automobiles with tolerances so controlled that repairs and other maintenance may become virtually unnecessary. Advances like this, achieved through the use of HPC, could alter dramatically the competitive landscape.

The challenge for the United States, Ms. Wince-Smith declared, "is to create an environment here at home in which innovation can flourish." This is not likely to happen without the robust use of HPC and its integration into the innovation process from concept and design through manufacturing. By leveraging this technology, our nation will realize a range of benefits, from tracking deadly pathogens for emergency preparedness to achieving greater output from oil reservoirs; from better drug design to modeling humans of all shapes and sizes for safety testing. This kind of innovation and outside-the-box thinking also has led to the creation of new industries, such as the computer-generated animation industry, which would not exist without advanced computing technologies.

Unfortunately, although the global competitive pressures on U.S. companies are rising, American HPC remains a niche market. Not only are traditional industrial users failing to apply this technology as aggressively as they could, the number of new users incorporating HPC into their business practices remains small. The conference was about turning this trend around and accelerating innovation through broader use of HPC. "By using HPC to accelerate the innovation process," Ms. Wince-Smith concluded, "we will be able to extend the horizon of applications, and new products and services beyond the range of our imagination."

Discovery and Application Walking Hand-in-Hand

Dr. Arden Bement, Director of the National Science Foundation (NSF), reinforced the link between HPC and competitive advantage by reminding conference participants that the computer and communications industry is still relatively young. He speculated that the research and business communities have only just begun to tap into and harness the power of HPC. He recalled Yogi Berra, saying, "*You've got to be very careful if you don't know where you're going,*

1 The Japan Times: April 28, 2005

because you might not get there." "I take that as a caution about muddled or half-hearted thinking about our future," Dr. Bement said, "and an encouragement to be bold in pursuing our nation's objectives."

For the National Science Foundation, cyberinfrastructure investment is a top priority. During the past several decades, computing technologies have helped to accelerate scientific discoveries and their associated technological applications at a staggering rate. "Discovery and application now go forward hand-in-hand," he explained, "each nudging the other toward new horizons, at an ever quickening pace, spurring new discoveries, creating new industries and transforming the old." The challenge for the United States in the future will be to make the next innovative leap beyond where other nations are looking. This will not be easy, because other countries are building great economic momentum, which will result in enormous competition for U.S. companies. "We must be increasingly nimble in our thinking, because the path forward is inevitably murky," counseled Dr. Bement. "Innovation," he continued, "is a leap into the unknown — a risky venture that breaks with tradition, stands conventional thinking on its head, plays havoc with established practices, destabilizes the marketplace, and brings big dividends."

He concluded his remarks by focusing on partnerships, pointing out that no sector — industry, academia, labor, or government — can afford to "go it alone." Effective partnerships have long distinguished the U.S. innovation enterprise from others around the globe. "We must all work together to tighten the links in knowledge creation, technological innovation, and business acumen," concluded Bement. "Only with dedicated collaboration will we be able to surmount what may be a tidal wave of competition for this new century."

Big Changes to Big Things

The keynote address, delivered by Mr. Roger Enrico, Chairman of DreamWorks Animation SKG, Inc. and former Chairman and CEO of PepsiCo, Inc. urged conference attendees to transform the United States and its economy by "making big changes to the big things" that will most impact the competitiveness of U.S. companies and spur innovation.

To illustrate his point, Mr. Enrico began with a short animated feature created by DreamWorks Animation SKG, Inc. in collaboration with the Council and narrated by the penguins from the animated DreamWorks movie *Madagascar*. The film showed the audience "big changes" that have affected everyday life as a result of HPC. U.S. industries have relied on HPC to advance Alzheimer's research and develop more effective cancer radiation therapy; model aircraft wings, cabin ventilation systems and engines for quieter, more fuel efficient planes; and develop high-efficiency, low-polluting energy alternatives such as hydrogen fuel systems. HPC, he asserted, is important to nearly every aspect of the U.S. economy. Maintaining and building on America's global technology leadership in HPC will be crucial for the health and wellbeing of every American.

Mr. Enrico admitted that it might seem odd for a "soda and chip maker" or a "cartoon studio" to need HPC. To illustrate its utility, he shared two real-world examples of how advanced computing has enabled both PepsiCo, Inc. and DreamWorks Animation SKG, Inc., two of the companies he has led, to achieve their strategic objectives.

His first example pertained to PepsiCo's Frito-Lay division. In the 1980's PepsiCo, Inc. was faced with weakness in its Frito-Lay division and was struggling to identify the key drivers for market growth. By commandeering the company's computing capacity, they were able to analyze 20 years of complex company and industry marketing data to identify new product releases as the single most important variable in market growth. This revelation drove the company to double research and development spending, boost marketing and focus the entire organization on product development. As a result, during the next decade Frito-Lay

HPC-generated animation courtesy of DreamWorks Animation, SKG, Inc.

saw sales and profit growth rates double and its market share rise from 40 to 60 percent. “Computing power, along with the talent of people who knew how to use it, gave us the confidence we needed to mobilize our entire company and put a large majority of our resources behind a focused strategy,” Mr. Enrico explained. “[It] gave us incredible insight that ultimately transformed our business.”

Mr. Enrico followed this example with a discussion of the importance of advanced computing to DreamWorks Animation SKG, Inc. “Our goal at DreamWorks is to be, both creatively and commercially, one of the very finest companies in the entertainment business.” And in fact, a DreamWorks animated movie generates six times the revenue of the average live action film.

“[But] without the most sophisticated technology, our great stories would never come to life on screen.” DreamWorks Animation SKG, Inc. invests significantly in high-end computing. In fact, without HPC, the computer animation industry would not exist. Audiences expect to see natural character movement and the play of light and shadows on fabrics and backgrounds in an animated film. Rendering these things realistically is a computationally intensive problem that requires HPC. To create *Madagascar*, the company used an HPC cluster of more than 2,500 processors, which in turn crunched highly sophisticated rendering algorithms for more than 12 million CPU hours.

DreamWorks Animation SKG, Inc. is also highly dependent on the talent to use this advanced technology. One quarter of its workforce is made up of scientists and engineers, many with advanced degrees. Mr. Enrico shared his concern about the difficultly of ensuring an adequate talent pool for his and other firms. He pointed out several disturbing trends that, if not addressed, will make this even more challenging in the future.

During the past three decades, the United States has fallen from a ranking of third to 17th worldwide in the relative number of graduating scientists and engineers. This decline is alarming at a time when jobs requiring scientific and engineering skills are growing at a rapid pace, forcing American companies to recruit employees from overseas to fill these positions. In the last ten years, the number of science and engineering jobs held by foreign nationals in the United States has jumped from 14 percent to 21 percent. However, as their home countries ramp up innovation capacity, these scientists and engineers are finding attractive jobs there. What will American companies do when these workers stop coming? Mr. Enrico speculated that they will relocate their critical research and innovation facilities offshore in pursuit of top quality talent. “The tragic fact is: our HPC success combined with a shortage of well-trained Americans may have the opposite effect to the one we seek,” warned Mr. Enrico.

If the United States is to ensure that its innovation system remains world class, and that its companies have the trained talent they need to succeed, it must inspire greater numbers of young people to pursue science and engineering careers. This, he noted, may require a change in the thinking of “hundreds of thousands of youngsters, not to mention, quite possibly governments and the entire education establishment.” Mr. Enrico cautioned that dramatic change such as this “will never come about if we first don’t free ourselves from the tyranny of incrementalism…the belief that somehow dramatic results will come about from un-dramatic actions.”

Mr. Enrico challenged those in the room and at the Council on Competitiveness to come up with that dramatic action…a “Moonshot” which could capture the imagination of this country’s youth, inspire them to complete a science or engineering education, and allow the United States to reclaim its leadership position in these disciplines. “Help make a big change to one of the biggest things I can think of,” he urged… “America’s technology leadership and the young people we need to preserve and grow it. [We need to] take the initiative on this, assume a leadership position, and help secure a more robust future for America.”

Solving Grand Challenges: The Users Perspective

HPC enables groundbreaking innovation by helping to answer the “what-if” questions whose successful solution will have a significant impact on society and propel competitiveness and prosperity. It accelerates the innovation process and provides researchers with added insights and answers throughout the discovery and invention phases of research and development, and makes manufacturing more efficient. It may be possible to capture the imagination of today’s youth by showing them that, by mastering computational science, they too can play a vital role in solving our "grand challenge" problems.

What are some of the grand challenges facing the United States and our industries? How can HPC help to solve them? What is preventing companies from successfully using HPC? What challenges could American companies tackle in the future, if better tools and systems were available? The first conference panel, moderated by Dr. Jeffrey Wadsworth, Director of Oak Ridge National Laboratory, offered four different perspectives from the HPC user community, exploring these issues in greater detail.

Understanding How the Brain Works

Dr. Mark Ellisman, Professor of Neurosciences and Bioengineering at the University of California, San Diego, School of Medicine, introduced the concept of "stretch goals" to the conference attendees, and related it to his research on brain function. Stretch goals, he explained, are 10- and 20-year programs, where scientists are brought together in interdisciplinary teams aimed at tackling grand challenges. A stretch goal in brain research, for example, will enable new understanding of the brain by linking data about macroscopic brain function to the brain's molecular and cellular underpinnings. By gaining a better understanding of what is happening at the atomic and subatomic levels of the brain, researchers may be able to accelerate cures for certain diseases, develop more effective drugs, and design better treatment options.

The human brain is made up of 100 billion nerve cells, called neurons, which gather and transmit electrochemical signals throughout the body. Research in the field of neuroscience demands HPC so that scientists can gather, capture, and analyze data using advanced technologies (such as electron microscopes) and computational tools. Collecting and analyzing brain-related data creates a "tyranny of scale," because the data ranges from the cellular to the atomic level and exists all across the globe at different research centers, making it impossible to manage and manipulate by traditional means. Already, the computational needs in the field of brain research are pushing the limits of current HPC systems.

In order to advance the study of neural pathways and molecular development, researchers are collaborating with colleagues across the United States, sharing data across distributed networks and from remote databases. Dr. Ellisman described the successful implementation of distributed computing and collaboration programs in brain research at the Biomedical Informatics Research Network (BIRN) at the University of California, San Diego. BIRN provides a robust IT infrastructure for researchers to hasten the understanding and treatment of diseases such as Alzheimer's, depression, schizophrenia, multiple sclerosis, attention deficit disorder, brain cancer, and Parkinson's. "We went with this project of linking HPC and data [because] imaging data with humans or animal models is very large," he explained. "Each one of the data sets is pulled from a different location, and they're brought together so that [the researchers] can then work on them." The project currently involves 26 research

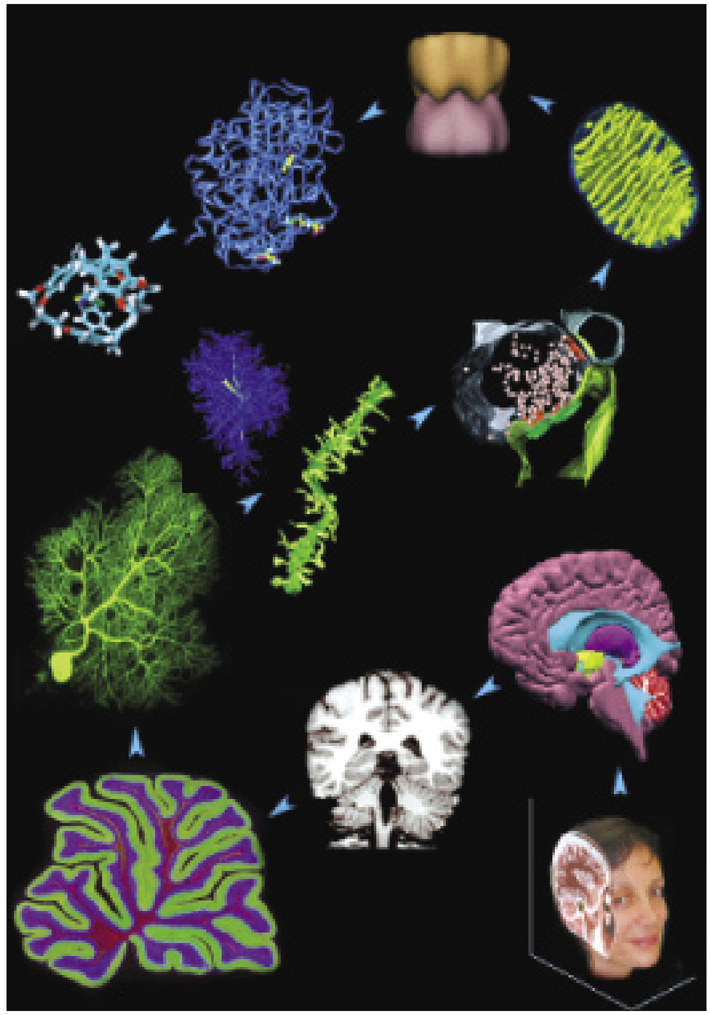

The "tyranny of scale" in brain research results from the massive amounts of data that are generated as researchers delve deeper into the cellular and atomic makeup of the human brain. Image courtesy of Dr. Mark Ellisman and the National Center for Microscopy and Imaging Research, University of California, San Diego.

sites from 19 universities and hospitals. BIRN is creating a federated data-sharing environment that taps into biological data held at geographically separate sites as if it were a single, unified database. Growing this and other networks and continuing to link data, computational resources, scientific instruments, and people via distributed research networks will require a leap forward in computational power and capacity.

Keeping the Lights On: Managing the Energy Grid

In the energy sector, the possibilities for and the advantages to using HPC are enormous. Dr. George Michaels, Associate Laboratory Director of the Computational and Information Sciences Directorate at the Pacific Northwest National Laboratory discussed the need to modernize, and possibly completely reorganize, this country's energy delivery and consumption system for both economic and national security reasons. The current electric grid, he noted, is owned by widely distributed, independently managed enterprises that need to cooperate more effectively to meet America's everyday needs for power, while avoiding expensive and

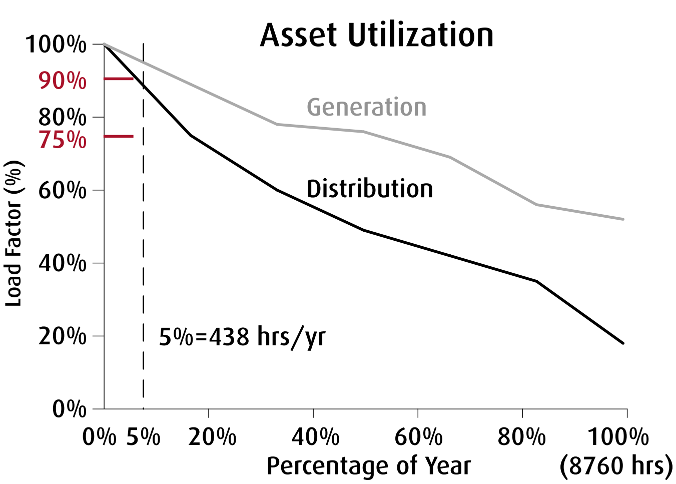

U.S. Power Grid load utilization as a function of full yearly capacity. It is important to note that nationally, the power grid is near peak loading only about 5 percent of the time.

Graphic courtesy of Pacific Northwest National Laboratory.

dangerous blackouts. What is needed is the integrated and holistic management of the energy grid, which is based on the real-time analysis of data that is collected continuously from every part of the grid.

"At every step along the way, there's an opportunity for failure, an opportunity for management, an opportunity for coordination, and there's also a real opportunity for HPC," stated Dr. Michaels. HPC could be the key to effectively managing this country's energy system. The rapid analysis of real-time data from across the nation-wide grid would allow better management of individual generators and transmission lines for greater economy and efficiency. As a bonus, better management of peak-load periods could reduce the need for costly new power plant construction, since peak-load periods occur infrequently during a typical year. Averaged over a year, about 30 percent of available generation and 50 percent of available transmission stands idle waiting for those brief periods of peak demand. Likewise, better grid management based on rapid analysis of real-time data could help to avoid or mitigate disturbances that today may cascade into massive failures, such as the northeastern United States blackout of August 14, 2003.

"The energy grid operates at light speed, and we manage it in minutes, at best," stated Dr. Michaels. "There is a real need to know how to predict what the system is going to do in the next half hour and be able to manage the power capacities that are distributed along the grid." However, the major factor standing in the way of an integrated energy system is the lack of necessary sensors, computers and software to provide sufficient, efficient and reliable management. In fact, many of the sensors, computers and software currently deployed are more than 10 years old. In addition, the systems are not in place to support HPC and real-time decision making, and the underlying simulations and models that are needed to understand the optimal energy solutions do not exist. So while there are plenty of significant "grand challenges" across the energy system that demand HPC, this sector still has a long way to go before it can maximize HPC.

Expanding Market Share

Dr. Michael Zyda, Director of the GamePipe Laboratory at the University of Southern California Viterbi School of Engineering, discussed the growing technical sophistication of electronic video games, and the opportunity to use gaming to significantly enhance education and training. Game production and gaming technologies are now becoming a vital component of the U.S. economy. In 2003, revenue from the gaming industry was about $11 billion, making it

America's Army®

America's Army® is a video game that was originally developed as a recruitment tool for the U.S. Army and has since become the fastest growing online game in the world. However, it quickly became a training tool as well.

The game uses realistic scenarios and proven learning techniques to expose a soldier to the kinds of challenges that he or she might experience, regardless of whether it is in training or in battle. It does the job so well that the drill instructors at Ft. Benning, Georgia now use the game with recruits who fail to perform adequately on the rifle range or the obstacle course, forcing them to play the game until they have mastered the skills. When the recruits complete the appropriate game levels, they are sent back to the range or obstacle course and usually pass their competency tests the next time.

Players hone their skills in rifle range training.
Image courtesy of America's Army®

the fastest growing segment of the $100 billion entertainment industry.

"There is a huge demand for better computer characters and story, and this demand is fueling the requirement for HPC and the capability for its easy deployment in next generation gaming consoles," explained Dr. Zyda. The industry also is exploring artificial intelligence so that games will respond to the emotional reactions of the players. Modeling and simulation of human emotion for the "immersive" environment is the next frontier for networked games and simulations. Games with this complexity simply cannot operate without HPC. In fact, next generation gaming consoles are "portable supercomputers" running at a teraflop or more. A market also is emerging for "massively multiplayer" online games, or games with hundreds or thousands of people playing simultaneously across multiple geographies. This kind of interactive environment will require HPC hardware and software both to run the games and also to simulate them in the development process. In addition, as games become more complex, users are demanding simple and easy to use interfaces.

Games also are evolving from purely entertainment to so-called "serious games" for education, training, simulation and strategic communication. As the next generation of college students and future national leaders comes of age, it seems a natural progression to use gaming as an educational tool, as these people have been immersed in "all things digital" from a very young age. The market for serious games already is emerging in health care, disaster response, public policy and communications, and is extending into America's corporations as a sophisticated training device. Games for these markets will require HPC to model and simulate the diverse human actions and reactions inherent in these multidisciplinary environments.

While electronic games are pushing the development of new technologies and the application of HPC, the gaming industry is struggling to find the talent it needs. Today's games are built primarily by people who are excellent game players, but they are often college dropouts. If the industry is going to meet the market demand for immersive, massively multiplayer games that require HPC to develop and run them, it will need people with undergraduate and master's degrees in math and computer science. Some universities are beginning to respond and are creating "gaming" degrees that prepare students to create these sophisticated games. Interdisciplinary research also is needed on technologies for future interactive games. Dr. Zyda suggested that by creating a science of games, "[we] can absolutely know for sure how to build games that can educate, that can be used for training and strategic messaging, and that can be used for entertaining."

Sustaining Market Competitiveness

Mr. Guru Bhatia, General Manager and Director of the Engineering Computing Information Technology Group for the Intel Corporation, shared how Intel uses HPC to design its processors and enhance its competitive advantage.

As the number of transistors on a chip grows, so do the complexity and the cost of chip design. Just as Moore's Law predicts doubling of transistors on a chip every 18 months, it has required a doubling of computer power every year since 1997 to support the design of these faster chips.

In 1998, Intel projected that by 2002 it would cost more than $300 million a year for the dedicated workstations and servers needed to do their chip design. Intel bet that an advanced HPC architecture could help control these costs and provide better performance. The bet paid off. Currently, 98 percent of Intel's chip design is carried out using this HPC architecture in a global design environment that supports 20,000 design engineers and 300 massively parallel applications from 45 design sites around the world. According to Mr. Bhatia, "in the last four years, the company has spent $1 billion less on computing."

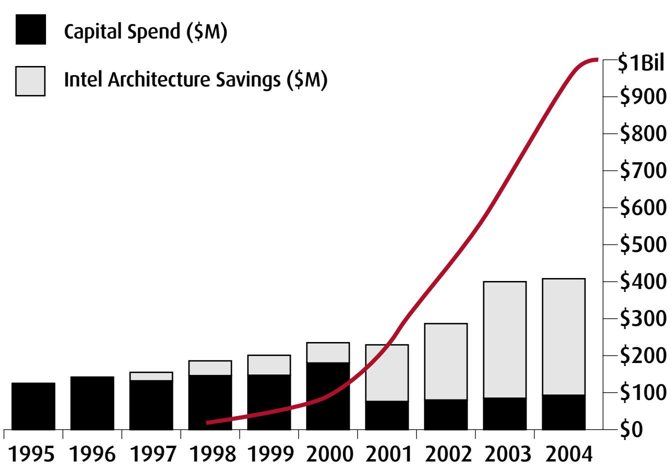

Accumulated savings as a result of Intel architecture migration. Graphic courtesy of Intel Corporation.

Mr. Bhatia explained that the next grand challenge for the semiconductor industry is to solve the optical distortion problem that arises as the feature size on a chip becomes smaller than the wavelength of the light used for lithography on the chip. In 2004, fourteen to fifteen thousand computers were used to compute the resolution enhancing techniques (RET) that control this distortion. However, by 2012, Mr. Bhatia speculated, solving this same problem will require 700,000 computers.

Independent Software Vendors Survey: The Need For Better HPC Application Software

The Council's *2004 Study of U.S. Industrial HPC Users*[2] identified application software limitations as a significant barrier preventing more aggressive use of HPC as an innovation driver across the private sector. Put simply, there is a lack of "production quality" HPC application software with the capacity to scale to hundreds or thousands of processors, and this is preventing companies from capitalizing on HPC. While the government is investing heavily in advanced computers that will reach quadrillions of operations per second (petascale), the software that would make this capability accessible to commercial users is lagging.

To understand the lack of advanced, commercial HPC application software, the Council, in partnership with the Defense Advanced Research Projects Agency, commissioned a first-of-its-kind survey to map the landscape and market dynamics surrounding the independent software vendors on whom many companies depend. Dr. Earl Joseph, Research Vice President for the High Performance Systems Program at International Data Corporation, explained the results. *The Council on Competitiveness Study of ISVs Serving the High Performance Computing Market*[3] reveals that the niche status of the HPC market discourages commercial development of HPC application software. More than a third of the software vendors serving this market qualify as small businesses, earning less than $5 million a year. The niche HPC market does not generate sufficient revenue to support investment in the research and development required to produce more advanced HPC application software. As a result, most independent software vendors must pursue the larger, commercial computing market, and the needs of HPC

2 See the Council's ***Study of U.S. Industrial HPC Users***, available at http://www.compete.org/hpc
3 See Appendix for Executive Summary of ISV Study. Complete Study available at http://www.compete.org/hpc

users, while important, are often a secondary concern. It also is unlikely that the market will solve this problem by itself. Unique public-private partnerships will be required to ensure that the U.S. businesses that rely on HPC for competitive advantage have the application software they need to solve their most complex, and competitively important problems.

Dr. Joseph highlighted the seven major findings from the survey:

- **The business model for HPC-specific application software has all but evaporated in the last decade.**
  As for-profit companies, ISV organizations are now targeting their software development so that it focuses on the broader computing markets (workstations, PCs, Macs) where they can realize significant growth. Today, the high-end HPC market typically represents less than five percent of an ISV's revenue, and in many cases, it is less than one percent.

- **ISV applications are important for improving and maintaining U.S. business competitiveness, but they can exploit only a fraction of available problem-solving power of HPC.**
  Today, high performance computers can be equipped with hundreds or thousands of powerful processors, and yet almost no ISV application takes advantage of more than 128 processors. In fact, today 82 percent of the codes can only scale to 32-processors or less, and 25 percent are strictly single CPU applications.

- **For many applications, the ISVs know how to improve scalability but have no plans to do so.**
  The adoption of clusters has allowed most ISVs to grow their revenue significantly with only normal feature enhancements and "technology updates." As a result, the ISVs are not forced to scale their software. Furthermore, when asked about how they might spend an infusion of research and development money, a majority of the ISV respondents indicated that it was unlikely that they would spend it on increasing scalability for such a small part of the total market.

> **When asked specifically about their plans for:**
>
> - Scaling to hundreds of processors: NO PLANS for 37 percent of codes
> - Scaling to thousands of processors: NO PLANS for 44 percent of codes
> - Scaling to tens of thousands of processors: NO PLANS for 60 percent of codes

- **The open source community is not now, nor has it been a significant source of new application software for HPC.**
  The survey also found that while the open source community has made a tremendous contribution to operating systems and middleware, it has not been a major contributor to ISV application development.

- **There is a lack of readiness for petascale systems within the ISV community.**
  Very few codes currently scale to thousands of processors or are even being considered for that level of scalability. ISV applications that are able to scale to that level do so because the underlying problems are relatively easy to parallelize. However, many of the most complex and consequential industrial problems are far more difficult to scale and therefore have not progressed.

- **Market forces alone will not address this problem and need to be supplemented with external funding and expertise.**
  The survey found that neither the ISVs nor the HPC hardware vendors possess the funding to make major research and development investments to provide fundamental rewrites of their codes. Absent a strong business case, ISVs will require external expertise and support to help them improve the scalability of their codes.

- **Most ISV Organizations Would Be Willing To Partner With Outside Parties To Accelerate Progress.**
  The vast majority, 83 percent, of the survey respondents stated that they would be open to developing partnerships with other organizations. When asked who their preferred partners would be, the top three were: other

code developers (25 percent); government labs (25 percent); and universities (22 percent).

The sobering findings from the ISV survey can be reduced to the following: Without additional funding and computer resources and a more compelling business case, ISVs are unlikely to rewrite their codes to address current scaling limitations, much less take advantage of petascale systems when they are available. The limited scalability of today's application software will likely serve as a major competitive barrier for U.S. businesses because their large, complex, and most important problems will not be solved within reasonable timeframes, or possibly at all.

In Their Own Words:

A Selection of Respondent Remarks from the Survey of ISVs Serving the High Performance Computing Market

"We already have enough creativity. What we need to do this is more time and human resources."

"We have made some significant strides in modifying our application for HPC, but we can't justify investing more."

"Show me the business case."

"To keep up with HPC hardware, there need to be better software developer tools."

"It would be great to have a stable five-year funding horizon to meet these HPC requirements."

"It's about time and money. To scale up for HPC, we'd need to reduce the risks."

"We need long term access to large systems with 10,000 plus processors, and we can't afford them."

"We need technical expertise and access to more experts in our field."

Challenges in HPC Application Software Development

The afternoon panel, moderated by Dr. Graham Spanier, President of The Pennsylvania State University, provided a discussion of the challenges of creating and maintaining application software suitable for a competitive, corporate "production" environment, and the role of universities and national laboratories to help accelerate development of new and/or updated code. It offered perspectives from three different communities — software vendors, national laboratories, and industrial users of HPC. Panelists revealed a "chicken-and-egg" challenge: On the one hand, the niche HPC market largely relies on independent software vendors to provide HPC application software. On the other hand, most software vendors have difficulty meeting this demand because the market is too small to support their software development costs. Panelist comments focused on three aspects of this dilemma.

Some users must develop their own HPC application software in order to gain a strategic, competitive advantage in the market.

Although most companies depend on commercially available HPC application software, some larger firms also are making significant investments in internal software development to ensure they can address unique, strategic needs. In the oil and gas industry, for example, firms like Chevron Corporation develop their own seismic processing and reservoir simulation software to gain a competitive edge. "If you can image exploration prospects before your competitors can, then you can get to that lease position first." explained Dr. Donald Paul, Chevron Vice President and Chief Technology Officer. Chevron recently completed the next research and development phase of their next-generation reservoir simulator from scratch, investing more than 100 man-years in development. Dr. Paul stated that it would take an order of magnitude more investment to create a supportable product and distribute it globally to over 5,000 engineers.

Mr. Loren Miller, Director of IT for Research, Development and Engineering at The Goodyear Tire & Rubber Company, offered a similar explanation for his company's investment in internal software development. "I'm not trying to distance myself or the company from the ISVs... In fact, I don't think there's a software package that either [Fluent or MSC.Software] offers that we don't have at least one and probably many copies of... We're doing what we're doing for competitive advantage ... short and sweet."

He went on to explain that computational analysis is taking the place of designing and building "test" tires in the development process. As a result, Goodyear has cut its cycle time substantially, as well as its costs. "We've gone from spending 40 percent of our development budget on building and testing tires to 15 (percent). That's a lot of time as well as money." Goodyear accomplished this through better computational analysis, simulation, and predictive performance testing. "Reducing our development cycle time is a big driver," admitted Mr. Miller, " but the only rationale we have for developing our own computational analysis software, and we try to be quite rigorous in doing this, is that it gives us a competitive advantage." As the only major U.S. tire company, Goodyear's competition is worldwide, coming from Japan, France, China and Korea. "Our tires have to compete worldwide. In order to solve our computational analysis problems, we are required to develop our own software."

Although advanced HPC software is being developed at the national laboratories, it is not readily transferable to commercial users.

Even though the federal government makes substantial investments in HPC software development, panelists noted the difficulty in leveraging that investment to help meet industry's HPC software needs. Dr. William Camp, Director for Computation, Computers, Information and Mathematics, at Sandia National Laboratories and Dona Crawford, Associate Director for Computation at Lawrence Livermore National Laboratory (LLNL) both indicated that most of their HPC application software is developed internally and usually does not meet the needs of industrial users. During the last decade, Sandia has invested more than $800 million in internal software development. LLNL is spending about $230 million a year in software development that includes verification and validation, model development and software tools. Despite this enormous investment, both acknowledged that the specific applications, the scale to which the software must be written, and the complexity of the code make it difficult for industry to use and for the ISVs to commercialize.

The complex research software required to support the national security mission of these laboratories often does not have a ready application in the private sector. "Sandia is doing problems that other people wouldn't need to do at this point or see a need to do, " said Dr. Camp. LLNL is driven by "big science challenges for our national security mission," said Ms. Crawford. "We're trying to understand what's happening inside a nuclear weapon when it's exploding, without testing that. Some of our computational applications are beyond even experiments." Additionally, the laboratories are writing codes that must scale to 10,000 or more processors. Most companies simply cannot afford this infrastructure.

The national laboratories also operate on time scales that are not feasible for companies facing time-to-market pressures. It can take six months to prepare a problem to run on a 10,000-processor system, far longer than industry can afford. Ms. Crawford pointed out that when the Department of Energy launched the Advanced Simulation and Computing program in 1995, it established a nine-year goal to be able to compute a certain resolution and a certain amount of physics in three dimensions on a particular weapon prototype system. Few companies take this kind of long-term strategic approach to using computation to transform their entire business.

Software written by the research and development community at the national laboratories often is difficult to use, not only by industry but also by the design communities within the laboratories. The researchers writing the code have master's degrees and doctorates and a lot of experience within their technical fields. The software they create requires a great deal of expertise to use. But industry and the design communities within the laboratories want fast desktop tools that don't require a Ph.D. to use them. The world of technical computing has evolved from one that was managed by a few highly skilled and educated engineers to one in which computing resources are spread throughout an enterprise to a wide variety of people. "It's a real disconnect that our software is largely developed by eggheads for eggheads...to use a pejorative term that I don't really mean," said Dr. Camp. "But it describes the fact

HPC simulation of an SUV stabilization system. Image courtesy of Ford Motor Company.

that even though the software we develop compared to 15 years ago is extremely well software-engineered, it's not particularly user-friendly." And so there is a "gap" getting this software out to a broader user base.

Finally, both Dr. Camp and Ms. Crawford expressed concern about another "gap": There is not an obvious way for the national laboratories to work with ISVs to identify those laboratory codes that might have commercial appeal and better understand what the laboratories could do to make those codes easier for ISVs to adopt. Dr. Camp speculated that "more likely than not, the ISVs would like to adopt some of our techniques into their codes rather than grab our code unchanged."

User dependence on legacy systems and their integration into key business processes can slow the development of new application software.

Panelists noted that long-time HPC users often are dependent on legacy systems, software and models, and are therefore reluctant to change, even when more advanced software is available. Introducing new or updated software could change the way a problem is run, which might change the results. For many long-time HPC users, consistency in results is often more important than improved performance.

Dr. Reza Sadeghi, Vice President, Enterprise Computing at MSC.Software Corporation explained some of the challenges his firm has in supporting customers of MSC's widely used structural dynamics code NASTRAN. MSC is one of the largest independent software vendors, and has made substantial investment in adapting NASTRAN to run on more processors. "'Crossing the chasm' by the end user is a challenge for us," he noted. "We have customers using 6-8 processors every day with NASTRAN and using codes like NASTRAN. ...(But) NASTRAN today scales up to 64 processors. Nobody knows, and nobody wants to invest in the infrastructure to actually test it... Even for crash solutions where the codes easily scale up to 128, I don't see more than 16 processors utilized on an average."

Users of legacy software want the ISVs to continue to support this older code. This often discourages ISVs from developing new HPC application software. Supporting legacy software can deplete an ISV's limited research and development funds. For example, in the case of MSC, "Anything that either rolls or flies was either designed at some level with MSC's NASTRAN or is certified with MSC NASTRAN," explained Dr. Sadeghi. "This means that there are thousands and thousands of existing models (sets of data) that need to be supported." Given the limited research and development resources of most ISVs and the requirement that publicly traded companies like MSC.Software Corporation and Fluent Inc. must report quarterly earnings, ISVs invest first in projects customers will pay for today, such as supporting legacy models. Only then can they consider higher risk, more costly efforts such as developing new solutions from scratch or undertaking the "hundred man-year investment" that is needed to correctly re-configure a complex code like NASTRAN for future petascale systems.

Industrial and National Laboratory Software Environments

	Industry	National Laboratories
People	• Bench designers and engineers mostly trained to B.S. or M.S. level with few computer scientists • Limited experience developing algorithms or writing code; mostly run ISV-supplied code • Seldom modify codes; intolerant of difficult-to-use codes or poor interfaces; need reliability and simplicity in codes more than performance	• Computer scientists and discipline scientists, many with Ph.D.s and substantial research experience • Accustomed to developing algorithms and writing own codes • Frequently modify codes for new problems; tolerant of difficult-to-use codes with poor interfaces; seeking revolutionary performance to make scientific discoveries
Resources	• Constrained by need to make profit; limited computer resources — hard to justify large computers unless direct contribution to sales or profit • Short-term projects associated with engineering and design • Very sensitive to hardware and software costs	• Less tied to "bottom line," especially for national security missions; ready access to large systems if needed by mission • Long-term projects associated with basic research. Projects may last years before results become available • Less sensitive to hardware and software costs
Problems and Codes	• Need answer yesterday to incremental design and engineering problems with short developmental schedule; expect intuitive user interfaces to codes written by others • Codes are stable — reliable, seldom modified, support legacy problems • Seldom attempt revolutionary performance advances; need support for legacy problems	• Less time pressure because most projects are long-term — can afford to take six months to set up a problem if necessary • Codes often unstable with poor interfaces because constantly being modified as research advances. Users tolerate this because often they wrote the codes • Regularly attempt revolutionary performance advances for new research discoveries; seldom consider legacy problems

Panelists also noted that industrial HPC users often are reluctant to make changes to hardware or critical software when it is integrated into key business processes. The business process re-engineering required to make the switch to newer, more efficient and effective applications is frequently too time consuming and risky when companies are under market pressures to deliver products to customers on time and within known and/or previously approved design and quality parameters. Dr. Paul explained that competitive advantage often comes from "how you integrate all of the applications that you use into major workflows and how well those are integrated into and drive the way the business works. The business process structures that you can build around the technology are as important as the technology, and in fact, they're dependent on the technology." While changing software or systems might be acceptable in a research environment, Dr. Paul explained that it is often intolerable in a business production environment where a single application may be integrated in a complex stream of business processes. "When you get down to production line people, they'll say: 'No thanks. I'm not interested unless tomorrow everything works just as well as it does today, only better'."

Moving Forward

The day's discussion exposed some of the challenges in ensuring the availability of application software suitable for a competitive, corporate "production" environment. Three possible action items emerged.

Accelerate HPC software development through partnerships among HPC stakeholders.

Panelists discussed the need for, and potential benefits from, partnerships among the research community, HPC users and ISVs in order to accelerate the development of both proprietary and commercial HPC application software.

Partnerships between industrial users and the national laboratories can be very beneficial to both sides. Mr. Miller and Dr. Camp discussed their collaboration, which began in 1993 with a small amount of seed money from the federal government to help Goodyear offset the initial risk in transferring technology from Sandia. Mr. Miller explained that the relationship has progressed into a genuine joint research and development program that provides Goodyear with valuable knowledge and software. He also pointed out that now "there isn't a penny of taxpayer money other than our own that goes into our joint R&D." "This is not a government subsidy," explained Mr. Miller. "Goodyear pays for all of the time and effort that Sandia puts into working on jointly developed technology."

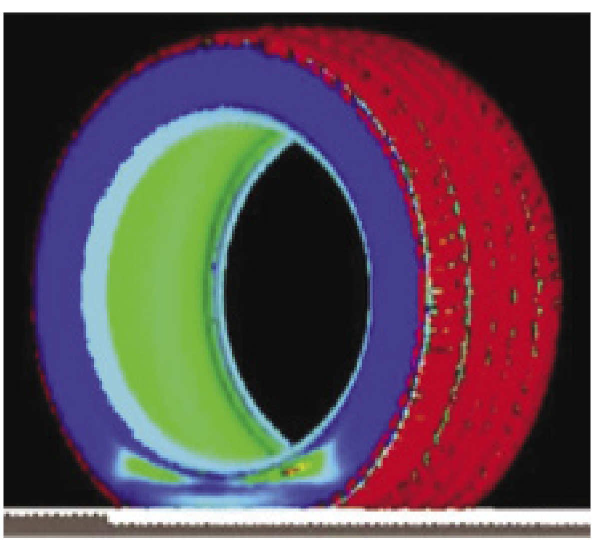

HPC simulation of tire stress points. Image courtesy of The Goodyear Tire & Rubber Company.

Dr. Camp confirmed Mr. Miller's comment, stating that: "We're forbidden by the government from doing work for free... and we don't do work for free." He emphasized that Sandia has benefited as much from the partnership as Goodyear, but in different ways. He explained that the national laboratories must guarantee every year to the United States that in a worse case accident with a nuclear weapon, there is less than one chance in a million that the explosive force (equal to the force of a half million tons of TNT) will be stronger than a half pound of TNT. "We're really good at high-end things. We're learning from Goodyear how to move things out of the high-end R&D community right on to the designers' desktops and making a difference at the start of the design process, rather than later on verifying that you got it right. This is incredible payback."

Partnerships between the national laboratories and the ISVs would help to supplement the manpower the ISVs need to both support current products as well as develop new solutions. "Yes, we do have a shortage [of manpower] and yes, we do need collaboration with the laboratories," confirmed Dr. Sadeghi. He suggested that such collaborations could significantly shorten the time to introduce new HPC software from 10 years to three-to-five years.

Panelists agreed that research partnerships among industrial users, national laboratories and universities are easier and therefore more common than those linked to commercialization. The national laboratories are prohibited from spending program dollars for anything other than their programs, so research in those areas that is also of interest to industry becomes the best means for collaboration. Ms. Crawford noted that a new "national mandate" might be required so that the national laboratories could be funded to assist the ISVs and industry in making its research software more useable in an industrial production environment. "When industry wants to take that and commercialize it, that's not our business, and we're not going to do that unless someone to makes it our business. If we had funding that said economic competitiveness is a national security competitiveness issue, it might actually change the way we do business."

Within the private sector, there is an increasing willingness among competitors and suppliers to partner on research projects. However, panelists noted the cost and difficulty of commercializing that research into a product that can be supported outside of the research and development community across a broad user base, and that it is probably unreasonable to expect the ISV community to take this on alone. They suggested instead that users and ISVs could collaborate, and might even be able to form non-profit organizations. The users would fund the development and would get free perpetual use of the software, paying only for support. "A not-for-profit structure doesn't mean it's not a viable business model," pointed out Dr. Paul. "Not-for-profits can be a very viable technology development mechanism if supported correctly." Dr. Sadeghi indicated that MSC.Software is beginning to see "first-tier companies" form consortia and ask his firm to help them "move to the next level." "I have a feeling," he said, "that when it comes to parallelization and true scalability and cluster computing, there will be more consortia formed to address these issues."

Panelists agreed that as "HPC for everyone" becomes more of a reality, there will be new opportunities for different kinds of relationships and partnerships across the HPC community.

ISVs should consider alternate business models to expand the user base and accelerate HPC software development.

Panelists discussed the critical need for new models that would allow for greater access to HPC to grow the user base and reduce the risk for users to experiment with this technology. They expressed particular concern for small- and medium-sized companies that still "build and test" in the face of increasing global competition, when in fact it may be faster to apply HPC computational methods in research and product development. While Fortune 100 companies can "do what they need to do," small and mid-sized firms cannot afford the HPC hardware and software needed to experiment with potential solutions to prove innovation or competitive advantage. "The barrier to entry is simply too high," said Mr. Paul Bemis, Vice President of Marketing at Fluent Inc. "We haven't delivered [HPC] to them in a way that allows them to use it without having to be a Ph.D., without having to spend a tremendous amount of money on hardware themselves or to build the infrastructure themselves and manage it."

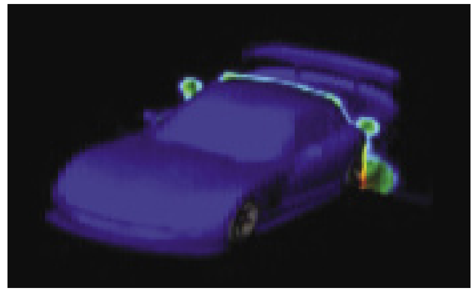

HPC simulation of NASCAR stock car aerodynamics. Image courtesy of DaimlerChrysler.

Dr. Sadeghi voiced a similar concern: "We have to change the time scale of evolution. We must bring the small and medium sized companies into the market so we can make faster and bigger steps in the world of computational science."

Mr. Bemis indicated that "some of the blame" for the high cost of entry rests with the ISVs and their pricing practice of providing unlimited use annual licenses. "Parallel processing is periodic heavy usage, and it doesn't fit an unlimited-use, annual model." He noted that this pricing model is evolving, so that users will be able to purchase software for high intensity periods of limited duration. This should attract small and medium sized companies who might not otherwise consider using HPC.

Dr. Paul suggested that a service model like the one that has evolved in the oil and gas industry might help grow the ISV revenue base needed to support HPC software research and development, and engage more new users. The oil and gas industry in the United States is comprised of hundreds of independent oil and gas companies, some with as few as three people, some with as many as a thousand. Because most of these firms are too small to own their own hardware and software, service companies have developed to do the processing for them. These companies bundle HPC hardware, software and specialized services "and for all practical purposes just deliver the pixels out the back end," explained Dr. Paul. "This is a very efficient system. All I need as a customer is the right business interface and somebody who knows what they're doing. And then I get the output, and I don't own any of the infrastructure or any of the software." These service companies have enhanced the value of their software products by wrapping higher value services around them.

The oil and gas industry model prompted several panelists to express the need for, and the potential benefits of, the emerging "utility computing" model, and its similarity to the old "service bureaus" run by Boeing Computer Services and United Computing in the 1980's. The panelists said the difference today is the delivery mechanism; the Internet, the browser and the power of the PC to handle the human interactive response. "What's needed," said Mr. Bemis, "and not yet available, is a graphical user interface that's very easy to use on a powerful PC that accesses the Internet as the medium to take advantage of large scale computing."

The challenge for the utility providers is to guarantee deliverability — what the user needs when it is needed. "I need to be able to purchase on-the-grid utility computing at competitive rates, that can be harnessed for delivery to the consumer as a value add to the solution," explained Mr. Bemis. "So I need to purchase it like I would electricity... at a dollar per CPU hour or at some rate so I can deliver it back to the market." The infrastructure to do that today doesn't exist.

Dr. Paul noted that while the service provider in Fluent's example has to consider these complexities, the customer really doesn't care about them. "And that puts more pressure on something that is really needed in HPC: the interface. In the service model, the interface is absolutely everything. If I as the customer can't get to the pixels and do what I want to do with them, I'll find someone else who can help me do that."

Panelists also noted that serious cultural changes would be required in the user community for utility computing to succeed. "Engineers are used to doing things the way they've done them for 20-30 years," noted Dr. Sadeghi. "Now we're giving them a browser and a PDF file or some HTML file and telling them to be content... this is the verification of your design and this is the data you want. That cultural change is moving very slowly." Mr. Miller similarly observed that a paradigm shift has to occur with design engineers. "Many fail to put multiple CPUs on a problem," he explained. "The design community is not used to conceptualizing and formulating design problems in a way that takes advantage of parallel solutions and then actually doing it. Parallel computing is not a trivial paradigm shift."

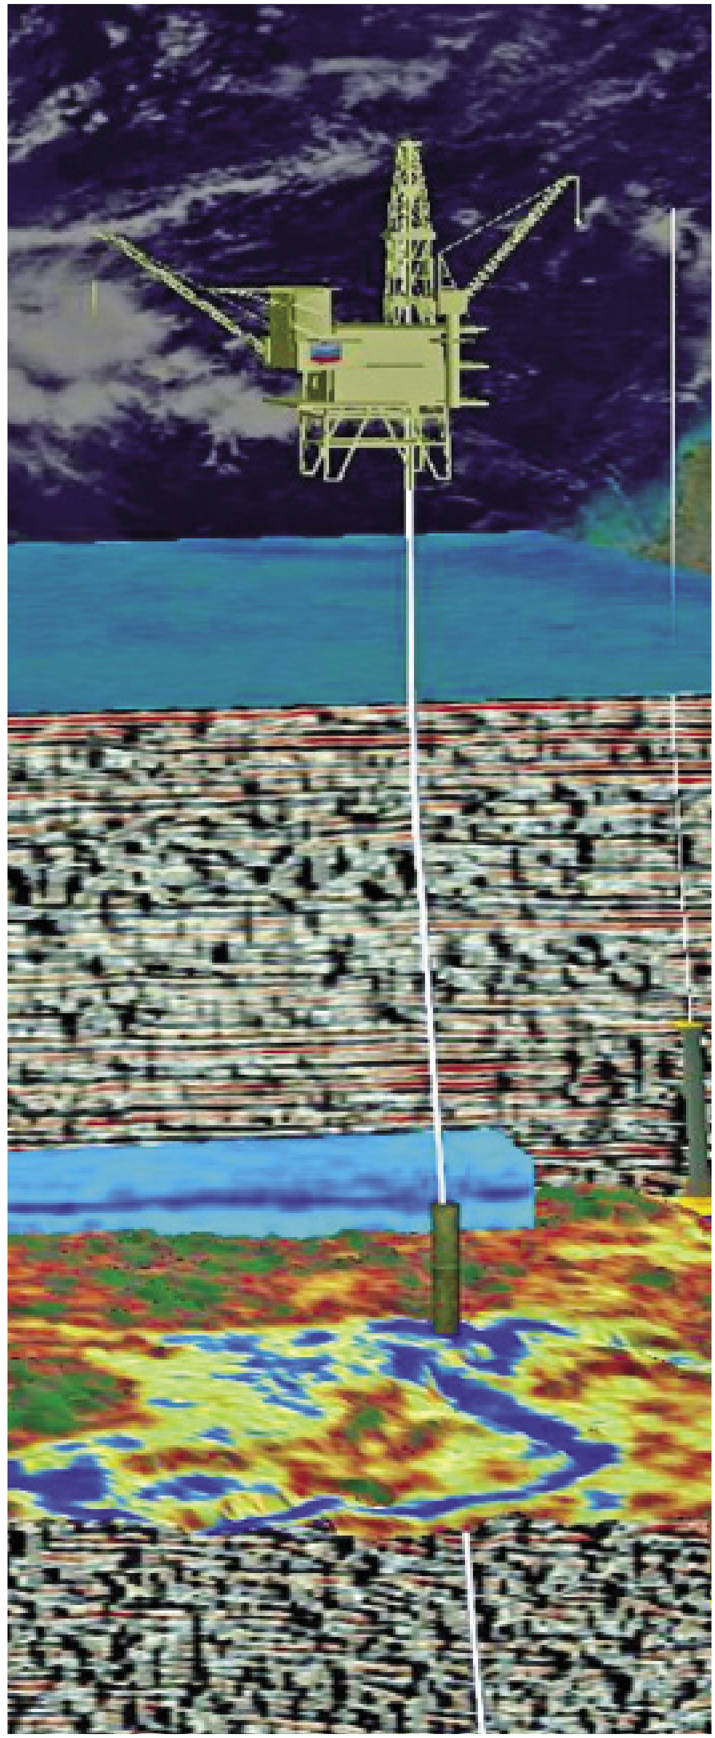

HPC simulation of a well location, tapping into the resources of an oil-filled sandy channel, colored in blue. Image courtesy of Chevron Corporation.

Federal funding for HPC should be balanced, pushing the top end of technology and expanding usage within the federal government.

The government plays a critical role in advancing usage of HPC, both by investing in the development and application of the most sophisticated HPC systems, and by using it itself as aggressively as possible where applicable. "If we cease with our imagination," cautioned Ms. Crawford, "and stop pushing computing power and the development of codes to take advantage of that, economic competitiveness fails." She also suggested that "not everyone has to get there at the same time...then we're not pushing the edge." Dr. Camp expanded, noting that there is an HPC technology maturation curve. It is therefore important for a few users to invest now in order to achieve petascale performance in the near term, even though everyone will not benefit at the same time. "Look at what it took to take advantage of tera-scale computing. That took a decade... and that was to get it into the high-end community. It took closer to a decade and a half to get it into U.S. industry."

Dr. Paul reinforced the benefit to industry of government investing at the high end. "That gives us a picture of what we will have as a commodity in 10 years. That lets us plan the commercial scale distribution." But he and others urged the government to take a "balanced approach" and not to forget "the other end." "The best thing that you could do for HPC in the long run is to have HPC be a part of everybody's compute life," Dr. Paul emphasized. "Because then people will internalize the process of computing as a way to make decisions and operate businesses. Push the top of the pyramid on the technology side to give us a glimpse of what you can do, and then widen the base as broadly as you can."

This afternoon panel discussion of the challenges and opportunities in the HPC software market set the stage for a workshop the following day to discuss this topic in more depth. The workshop, *Accelerating Innovation for Competitive Advantage: The Need for HPC Application Software Solutions*, created a more detailed roadmap of public and private sector actions to ensure a financially strong independent software industry, and the availability of the production quality HPC application software that users need to sustain competitive advantage. The Executive Summary of the workshop report is included in the Appendix. The full report is available on the Council's website, http://www.compete.org/hpc.

Conclusion

Dr. David E. Shaw, Chairman of D.E. Shaw & Co., Inc. and Co-chair of the Council's HPC Advisory Committee, wrapped up the conference by urging the participants to continue to explore mechanisms that will sustain the long-term health of HPC. Not only is HPC tied to U.S. economic competitiveness but it is also linked intimately to ensuring national security.

Dr. Shaw worried, "what kind of shape will the United States be in strategically, and in terms of our national competitiveness and national security, 10 and 15 years from now, if we are not investing in and planning for the long term by leveraging HPC to drive our innovation process?" Taking input from the conference attendees, Dr. Shaw agreed that excessive attention to quarterly earnings is short sighted. "What it probably takes," he concluded, "...is to have some courageous CEOs. People who are willing to wake up to long-term prospects, people who are willing to stand up and say, this is going to work and while it may be a risk, we are going to invest in this."

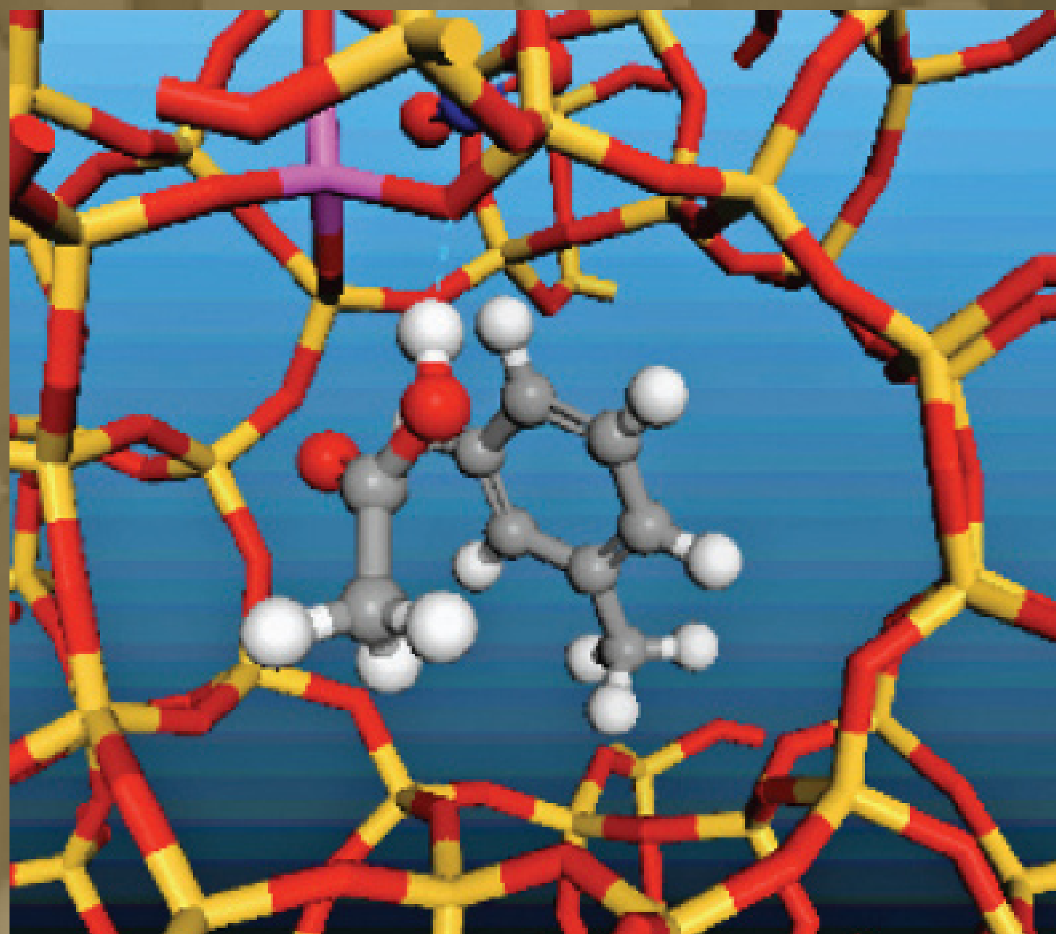

HPC simulation of an "environmentally friendly" chemical synthesis of agrochemicals and pharmaceuticals. Image courtesy of Accelrys.

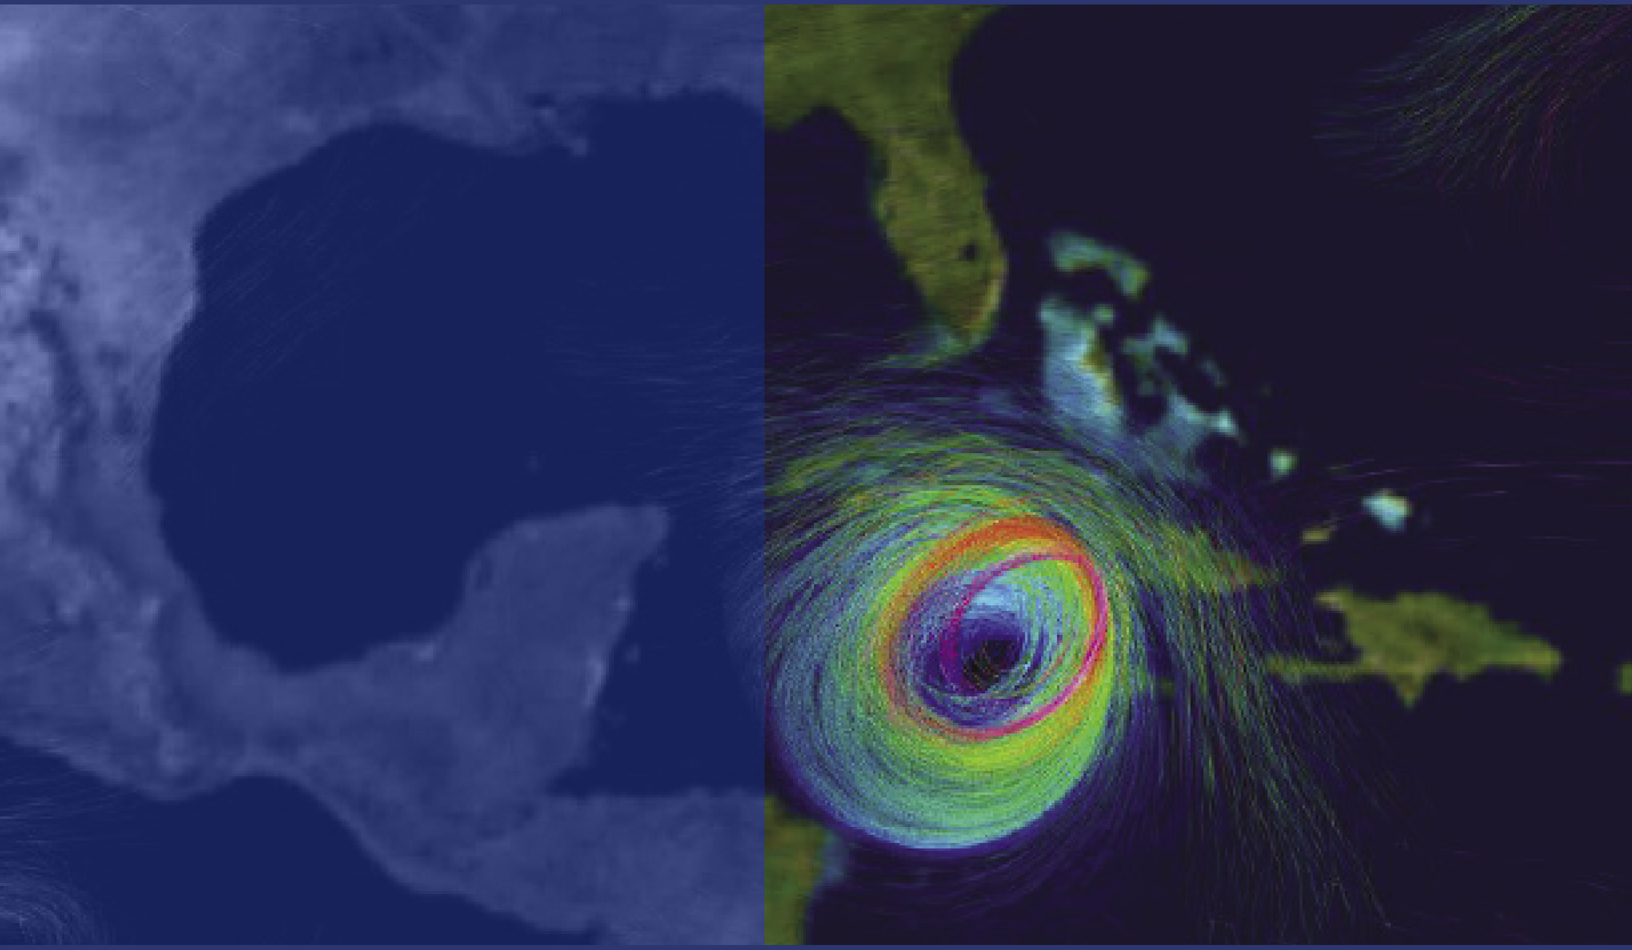

HPC simulation of Hurricane Ivan, 2004. Image courtesy of Intel Corporation.

APPENDICES

APPENDICES

Conference Agenda

Executive Summary of *Council on Competitiveness Survey of ISVs Serving the High Performance Computing Market*

Executive Summary of Workshop Report, *Accelerating Innovation for Competitive Advantage: The Need for HPC Application Software Solutions*

Council on Competitiveness High Performance Computing Advisory Committee

Council on Competitiveness Members, Affiliates and Staff

CONFERENCE AGENDA

Second Annual High Performance Computing Users Conference: Accelerating Innovation for Prosperity, July 13, 2005

7:30 a.m.	**Breakfast/Registration**
8:15 a.m.	**Welcome** • Deborah L. Wince-Smith, President, Council on Competitiveness • Dr. Karen A. Holbrook, President, The Ohio State University, and Co-Chair, Council on Competitiveness HPC Advisory Committee
8:30 a.m.	**Morning Keynote Address: A Challenge for Our Future** • Roger Enrico, former CEO, PepsiCo; current Chairman, DreamWorks Animation, SKG, Inc.
9:15 a.m.	**Panel 1: High Performance Computing - Driving Solutions to Grand Challenge Problems** Panelists will discuss cutting edge, "grand challenge" problems that require leaps in software development to successfully address them, and whose successful solution will have a significant impact on society and propel competitiveness and prosperity. **Moderator:** Dr. Jeffrey Wadsworth, Director, Oak Ridge National Laboratory **Panelists:** • Guru Bhatia, General Manager and Director, Engineering Computing Information Technology Group, Intel Corporation • Dr. Mark H. Ellisman, Professor of Neurosciences and Bioengineering, University of California, San Diego, School of Medicine and Director, the Center for Research on Biological Systems, The National Center for Microscopy and Imaging Research and the Biomedical Informatics Research Network (BIRN) Coordinating Center, University of California, San Diego • Dr. George Michaels, Associate Laboratory Director, Computational and Information Sciences Directorate, Pacific Northwest National Laboratory • Dr. Michael Zyda, Director, GamePipe Laboratory, Viterbi School of Engineering, University of Southern California Information Sciences Institute
10:45 a.m.	**Break**
11:15a.m.	**Results of the Study of Independent Software Vendors Serving the HPC Market** Study results will be released revealing the state of independent software vendors (ISVs) globally, which industries/companies are dependent on which ISVs, user and ISV perspectives of the strengths and weaknesses of the ISV industry, and areas where partnerships could help to advance software development. • Dr. Earl Joseph, Research Vice President, High Performance Systems Program, International Data Corporation
12:15 p.m.	**Luncheon/Keynote:** • Dr. Arden L. Bement, Jr., Director, National Science Foundation
1:45 p.m.	**Panel 2: High Performance Computing: Bridging the Software Gap** Industry, university and government panelists will discuss the challenges of using, maintaining, and creating application software suitable for a competitive, corporate "production" environment, and the role of universities and national laboratories to help accelerate development of new and/or updated code. The panel format will be a roundtable discussion **Moderator:** Dr. Graham B. Spanier, President, The Pennsylvania State University **Panelists**: • Paul Bemis, Vice President Marketing, Fluent Inc. • Dr. William Camp, Director, Computation, Computers, Information and Mathematics, Sandia National Laboratories • Dona Crawford, Associate Director Computation, Lawrence Livermore National Laboratory • Loren Miller, Director, IT Research, Development & Engineering, The Goodyear Tire & Rubber Company • Dr. Donald Paul, Vice President and Chief Technology Officer, Chevron Corporation • Dr. Reza Sadeghi, Vice President, Enterprise Computing, MSC.Software Corporation
3:30 p.m.	**Next Steps/Building the Strategy** • Dr. David E. Shaw, Chairman, D.E. Shaw & Co., Inc. and Co-chair, Council on Competitiveness HPC Advisory Committee
4:00 p.m.	**Adjourn**

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

This study is the first independent, extensive assessment of the landscape and market dynamics surrounding Independent Software Vendors (ISVs) that serve high performance computing (HPC) users.[1] An important impetus for undertaking this study was the July 2004 "Council on Competitiveness Study of U.S. Industrial HPC Users," sponsored by the Defense Advanced Research Projects Agency (DARPA). The study found, among other things, that 97 percent of the U.S. businesses surveyed could not exist, or could not compete effectively, without the use of high performance computing (HPC). This study and the Council's yearly HPC Users Conference identified application software issues as a significant barrier preventing more aggressive use of HPC across the private sector.

To meet their HPC needs, American businesses — and key areas of the U.S. Government and the scientific research community — rely on a diverse range of commercially available software from ISVs. A serious gap exists between the needs of HPC users and the capabilities of ISV applications. High-end HPC users want to exploit the problem-solving power of contemporary HPC computer servers with hundreds, thousands or (soon) tens of thousands of processors for competitive advantage, yet few ISV applications today "scale" beyond 100 processors and many of the most-used ones scale to only a few processors in practice.

A serious gap exists between the needs of HPC users and the capabilities of ISV applications.

It is important to understand that the ISV organizations are not at fault here. The business model for HPC-specific application software has all but evaporated in the last decade. As for-profit companies (in most cases), they focus their software development primarily on the much larger and more lucrative technical computing markets for desktop systems (workstations, PCs, Macs) and smaller servers. IDC market research shows that the HPC portion of the technical server market often represents less than five percent of their overall revenues, and in some cases this figure is less than one percent. Even if they could afford this investment, the motivation for major rewrites is generally inadequate because the HPC market is too small to reward this investment. For business reasons, the needs of HPC users are often an important but secondary concern.

The business model for HPC-specific application software has all but evaporated in the last decade.

For U.S. industries that need to out-compete their non-U.S. competitors by out-computing them, the limited scalability of today's application software can present a major barrier. In practice, it means that large, complex, competitively important problems, such as those encountered in designing new cars and airplanes and pharmaceuticals, or increasing the yield from oil reservoirs often cannot be solved today in reasonable timeframes. While yesterday's problems may run faster, companies find it difficult to solve the new, cutting edge problems that will propel them to the head of the competitiveness pack. In effect, they are standing still. And standing still is falling behind.

For U.S. industries that need to out-compete their non-U.S. competitors by out-computing them, the limited scalability of today's application software can present a major barrier.

1 See *Council on Competitiveness Study of the ISVs Serving the High Performance Computing Market*, available at www.compete.org/hpc.

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

Key Findings

The business model for HPC-specific application software has all but evaporated in the last decade

ISVs typically cannot afford to spend the time and money that would be needed to rewrite their application software.

As for-profit companies (in most cases), ISV organizations focus their software development primarily on the much larger technical computing markets for desktop systems (workstations, PCs, Macs) and small servers. The technical HPC computing market often represents less than five percent of their overall revenues, and in some cases this figure is less than one percent. Software development is expensive and labor-intensive, and most ISVs are small to medium sized companies. Even when business in their mainstream markets is doing well, ISVs typically cannot afford to spend the time and money that would be needed to rewrite their application software to meet the more-demanding requirements of the small market of HPC users. For business reasons, the needs of HPC users are often an important but secondary concern. Given the shape of their markets — high-volume and revenues from sales to smaller technical systems, relatively low revenue from the high end part of the technical computing pyramid — the return on investment for developing highly scalable codes for HPC users usually does not justify the expenditures or risks.

"We have customers asking for this, so it should be a priority. But we need money and then a person dedicated to this task, plus bigger hardware to develop and test our applications on."

"We just have too much to do. We would need more time in the day to address the needs of HPC users."

ISV applications are important for improving and maintaining U.S. business competitiveness, but they can exploit only a fraction of the available problem-solving power of today's high-performance computers (HPC).

Few ISV applications today can take advantage of more than 128 processors.

Contemporary HPC computer servers can be equipped with hundreds, thousands or (soon) tens of thousands of powerful processors, yet few ISV applications today can take advantage of more than 128 processors. Some of the important applications for the automotive and aerospace industries cannot currently scale beyond 1-4 processors. Advanced computational tools play a major role in U.S. industrial competitiveness by assisting companies in bringing new and/or more capable products to market more quickly than their competitors around the world. Although scalable computer architectures such as clusters have allowed US and other companies to amass "mind boggling" amounts of raw computation power within their budgets, large classes of application programs have not been able to take significant advantage of this power. Increasing the scalability of ISV applica-

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

tions could enable industries that rely on HPC to improve product success, quality and time-to-market substantially, but in many cases this would require ISV organizations to rethink and fundamentally rewrite their software.

"Many ISV codes don't scale beyond 32 or 64 processors, sometimes fewer, at a time when the largest HPC systems have 1,000 or even 10,000 processors. In fact, in the area of structural analysis, many of the widely used applications barely scale to eight processors. This severely limits the size of the problem that can be addressed within a reasonable amount of time."

"Better algorithms need to be developed to scale applications for HPC users."

"As biological data volumes continue to escalate, researchers need more capable ways of exploring, analyzing and annotating this data."

For many applications, the ISVs know how to improve scalability but have no plans to do so.

Changes in market dynamics, especially the adoption of clusters, have allowed most ISVs to grow revenue with only normal feature enhancements ("technology updates"). Even if an ISV had the resources for a major rewrite, the ISV might choose to spend that R&D money on other projects, rather than on increasing scalability for a small part of the total market.

"We have made some significant strides in modifying our application for HPC, but we can't justify investing more."

When the task is scaling to hundreds of processors, ISVs representing about 37% of codes that could be scaled have no plans to upgrade the scalability of their products. This figure increases to 44% when the goal is scaling to thousands of processors, and to 60% for tens of thousands of processors.

IDC has found from other research in the HPC sector that the underlying problems ISV applications address vary greatly in complexity, and for this reason it is easier to scale up some applications than others. ISV applications that are able to scale today to large numbers of processors in many cases do so because the underlying problems they address are relatively easy to parallelize ("embarrassingly parallel"). Some of the most complex and consequential problems are far more difficult to scale to large numbers of processors.

"We already have enough creativity. What we need to do this is more time and human resources."

"We have made some significant strides in modifying our application for HPC, but we can't justify investing more."

"We need to see a business need from our customers."

"Show me the business case."

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

The open-source community is not now, nor has it been a significant source of new application software for HPC

The vast majority of ISV applications (88%) are supplied by for-profit businesses or come from universities (7-8%). Only about 3 percent of the applications are "open source" codes. Note that most open source software is middleware and not application software.

Most of the applications (73%) were developed by the ISV organizations themselves, although one out of every four (24%) was born in a national laboratory or university. Only 3 percent of the applications are based on open source software. (See Figure 1.)

Most ISVs need to pursue profitable growth and can ill afford investments of time or money that are unlikely to contribute to this goal.

Original Source of ISV Application

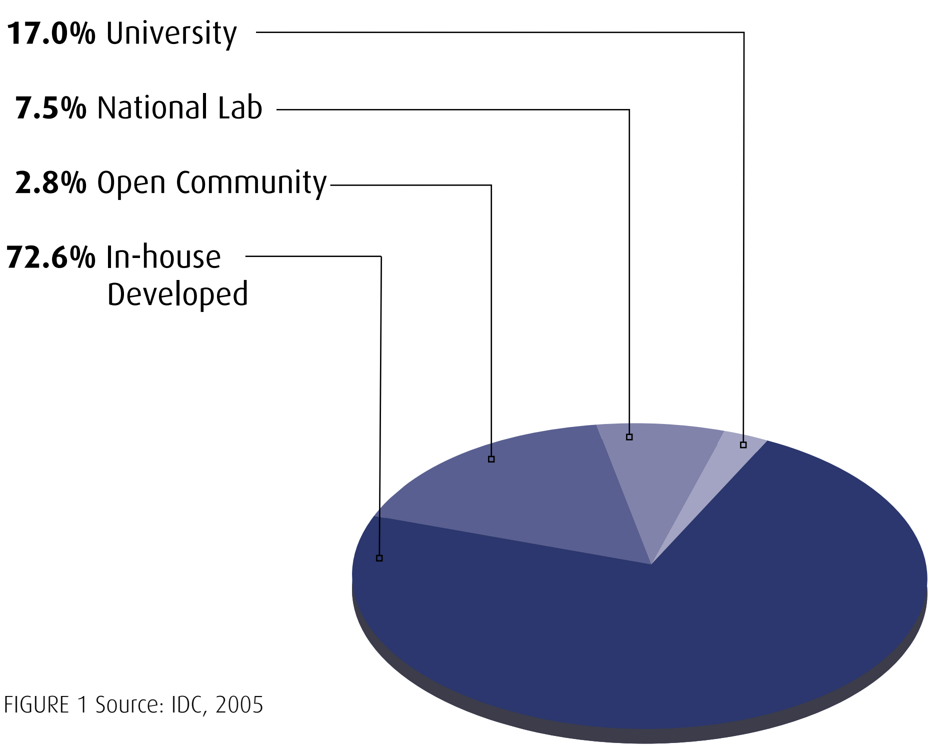

FIGURE 1 Source: IDC, 2005

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

The majority of the ISV applications (88%) are supplied by for-profit businesses. By contrast, only 7-8 percent come from universities, and an even smaller number (3%) are open-source codes. (See Figure 2.) This preponderance of for-profit applications means that most ISVs need to pursue profitable growth and can ill afford investments of time or money that are unlikely to contribute to this goal.

Percentage of Applications by Type of Ownership

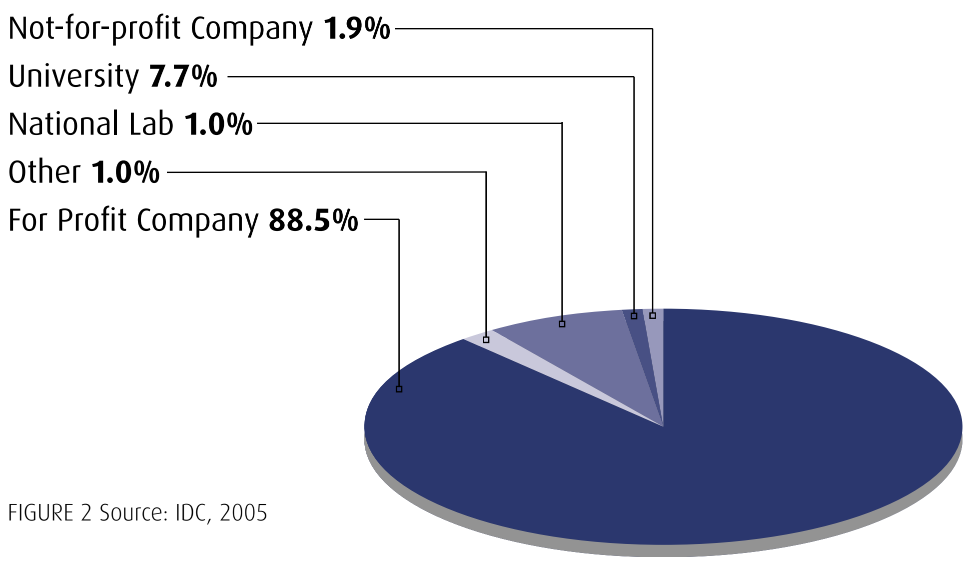

FIGURE 2 Source: IDC, 2005

If current development timeframes continue, the majority of ISV codes will not be able to take full advantage of petascale systems until three to five years after they are introduced.

There is a lack of readiness for petascale systems

Three-quarters (74%) of the ISV applications are "legacy applications" that are more than five years old, and seven out of eight (87%) are at least three years old. Fewer than half (46%) of the ISV applications scale even to hundreds of processors today, and 40% of the applications have no immediate plans to scale to this level. Very few codes scale to thousands of processors today or are being aimed at this level of scalability. If current development timeframes continue, the majority of ISV codes will not be able to take full advantage of petascale systems until three to five years after they are introduced.

"To keep up with HPC hardware, there need to be better software developer tools."

"We would need to extend into additional programming languages."

"We'd have to take a whole new approach to our software code."

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

Market forces alone will not address this problem and need to be supplemented with external funding and expertise

Historically, HPC hardware vendors operated on large margins and invested substantial human and financial resources in collaborating with application ISVs to improve the performance of application software on their HPC hardware products. In today's commoditized, lower-margin market for HPC hardware, neither HPC hardware vendors nor the ISV organizations themselves can afford to make major new R&D investments to fundamentally rewrite application software to take advantage of highly scalable systems. Market forces alone will not address the gap between HPC users' needs and ISV application software capabilities. Market forces need to be supplemented with external funding support and expertise to improvethe scalability of ISV software that is needed for improving the competitiveness of U.S. businesses.

Overall annual sales revenues (all products and services) of organizations offering ISV applications show a bifurcated pattern, with strong representation (29%) in the $1-5 million range and in the $50 million and up realm. (See Table 1.) Few ISV applications (3%) are associated with organizations in the $25-50 million range.

Number of ISV Applications and Companies by Total Company Revenue

	Companies		Applications	
Total Company Revenue	**Number**	**Percent**	**Number**	**Percent**
Under $1M	6	11.1%	9	8.2%
$1M to $5M	10	18.5%	27	24.5%
$5M to $10M	7	13.0%	11	10.0%
$10M to $25M	5	9.3%	10	9.1%
$25M to $50M	3	5.6%	3	2.7%
Over $50M	11	20.4%	32	29.1%
No response	12	22.2%	18	16.4%
Total	**54**	**100%**	**110**	**100%**

TABLE 1 Source: IDC, 2005

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

Revenue for fundamental rewrites is generally not available.

Most technical ISVs lack the funding and/or the business case to provide fundamental rewrites of their codes. Technical server markets are very small relative to most commercial software market segments, and the capability computing segment is only a small portion of that. For example, a "hot" computer game can generate $250 million of revenue, whereas a large technical ISV only earns about $50 million of revenue per year across all products. Furthermore, over a third of the ISVs that provided total revenue figures qualify as small businesses, earning less than $5 million a year. Even if an ISV invests the industry average of 10 percent of revenue in R&D, this amount is usually only sufficient to add selected features and cover testing and certification on a large number of different computers. Revenue for fundamental rewrites is generally not available.

"It would be great to have a stable five-year funding horizon to meet these HPC requirements."

"It's about time and money. To scale up for HPC, we'd need to reduce the risks."

Most ISV organizations would be willing to partner with outside parties to accelerate progress

Five out of six (83%) of the respondents said they would be open to developing partnerships with other organizations, and when the "maybe" responses are added in, the percentage climbs to 98%. The preferred partners were other code developers (25%), government labs (25%) and universities (22%). (See Table 2.)

Partners ISVs Selected as Potentially Most Useful, by Application

Partner Type	Number of Applications for Which the Partner Would Be Useful	Percent of Overall Responses
Other code developers	61	25.2%
Government labs	60	24.8%
Universities	53	21.9%
Buyers	43	17.8%
Investors	25	10.3%
Total:	**242**	**100%**

TABLE 2 Source: IDC, 2005 Note: Multiple responses permitted.

EXECUTIVE SUMMARY: ISV STUDY

Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software

In past studies, respondents have sometimes indicated resistance to the idea of collaborating with the U.S. Government, believing that government collaborations may impose unwanted conditions and requirements ("strings"). In sharp contrast to this history, all 104 ISV respondents were open to the possibility of working with the government, and 93 of them (89%) gave a definite yes. (See Figure 3.)

Willingness to Collaborate with U.S. Government, by Application

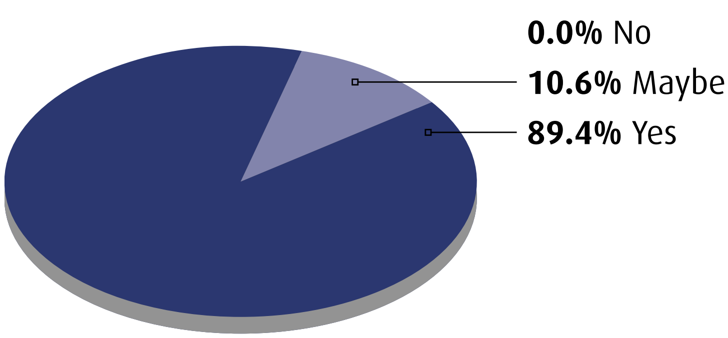

FIGURE 3 Source: IDC, 2005

"We need access to the newest hardware platforms, to machines with 10,000 processors."

"There needs to be stronger cooperation between HPC software, hardware and code developers."

"We'd also need more field research and input from user community."

"We need long term access to large systems with 10,000 plus processors, and we can't afford them."

"We need access to the newest hardware platforms, to machines with 10,000 processors."

"We need technical expertise and access to more experts in our field."

EXECUTIVE SUMMARY: SOFTWARE WORKSHOP

Accelerating Innovation for Competitive Advantage:
The Need for HPC Application Software Solutions
July 14, 2005

High Performance Computing (HPC) has become essential to accelerating innovation and assisting companies in the creation of new inventions, better designs, and more reliable products, processes, and services at lower costs.

Solving industry's complex problems requires not only the most advanced hardware, but also sophisticated application software. Unfortunately, a variety of market forces and technical challenges in recent years have caused software developers to turn their focus away from creating new and innovative HPC applications. As a result, the need for application software and middleware has become a pacing item in the private sector's ability to harness the full potential of HPC.

On July 14, 2005, the Council on Competitiveness and the Ohio Supercomputer Center co-sponsored a daylong workshop[1] to examine the reasons behind the current lack of production quality HPC application software. Workshop participants included HPC users, Independent Software Vendors (ISVs), university researchers, hardware vendors, and government scientists and engineers. They noted that the HPC market remains a niche within a much larger commercial computing market, inhibiting the development of the HPC application software that industry needs to achieve competitive advantage. For example:

- The small size of the market discourages development of new application software;
- The traditional ISV business model, which charges a license fee per processor in the system, is perpetuating the niche market;
- Because many ISVs qualify as small businesses, they often cannot afford to support multiple HPC platforms or to acquire large-scale parallel systems to test new software;
- A shortage of skilled talent also is holding back HPC software development.

Technical barriers are also slowing the development of advanced application software. Participants noted that:

- Software from the research community often is not adequate for industry's needs;
- In addition, some open source license models inhibit commercialization of HPC research software;
- Legacy code requirements and hardware limitations are stalling software development;
- Current HPC software needs better interfaces so that it can be integrated into the business workflow.

The workshop participants strongly agreed that a vibrant, growing HPC market is the best long-term guarantee that production quality HPC application software will be available. They also concurred that wider usage and adoption of HPC is key to future economic growth and the nation's ability to maintain its leadership position in the global marketplace.

The challenge that emerged during the workshop discussion was how best to stimulate HPC market growth, given the diverse and divergent needs of the current and potential HPC user community.

- The market can grow *deeper* if it becomes more valuable to current users. Greater depth can be achieved through the development of intuitive user interfaces for non-expert users; support of comprehensive design and engineering through integrated models that simultaneously compute the various processes that determine the solution to a problem; updating legacy codes and scalable applications to solve more complex problems on larger parallel computers; and developing new models and

1 See Workshop Report, ***Accelerating Innovation for Competitive Advantage: The Need for HPC Application Software Solutions***, available at www.compete.org/hpc

EXECUTIVE SUMMARY: SOFTWARE WORKSHOP

Accelerating Innovation for Competitive Advantage:
The Need for HPC Application Software Solutions
July 14, 2005

algorithms that are more comprehensive and accurate.

- The industrial HPC market can grow *broader* with the inclusion of new users in emerging fields such as bioinformatics, which will require new data intensive applications and more intuitive user interfaces to aid the new user to set up a problem and interpret the results. New software licensing models that encourage new users to try HPC without risking large amounts of money or time would also help to broaden the market.
- The HPC market can evolve by encouraging ISVs to develop different business models around their software products that increase their value to users and help to extend the market beyond its current state.

The workshop also recognized that the HPC ISVs are collectively at a transition stage or inflection point, where their future will consist of either watching their markets decline, or seizing opportunities for transitions that grow the market.

By the conclusion of the workshop, a number of recommendations for both the public and private sectors emerged, that could help stimulate growth in the HPC market and overcome the barriers preventing development of advanced HPC application software.

Recommendations for ISVs:

- ISVs and user companies should partner to conduct experiments to test the business value of HPC.
- HPC ISVs, in consultation with their customers, should consider alternate business models in order to increase revenue and financial stability and encourage HPC market growth.
- HPC ISVs should make their priority the development of easy-to-use interfaces so that HPC applications can be integrated into organizational workflows.

Recommendations for Universities:

- The nation's universities should increase their educational programs in computational science at the undergraduate and graduate levels to meet the need for skilled technical workers.
- The university and laboratory research communities should enhance their understanding of ISV needs and requirements so that they can leverage their own software research and education agendas to assist ISVs where appropriate.

Recommendations for the Federal Government:

- The government should modify its research support practices to provide sustained (multi-year) funding for research teams to develop mature research codes and algorithms and should encourage commercialization of suitable codes.
- The federal government should revise its H-1B visa continuity of employment requirements, so that ISVs can attract the talent they need.
- The federal government should carefully monitor foreign acquisitions of key ISVs.
- Where open source HPC research codes are being developed, terms of government grants and contracts should more seriously consider BSD model licenses, to enable ISVs to build commercial products on the codes without jeopardizing the ISVs' privately created intellectual property.

Implementing these recommendations entails considerable risk to the ISVs, the users, the research community, the government, and investors. However, failure to take action may inhibit competitive advances by U.S. companies and place them in jeopardy should other countries or companies capitalize on the potential of HPC first.

COUNCIL ON COMPETITIVENESS HIGH PERFORMANCE COMPUTING ADVISORY COMMITTEE

Industry Co-Chair
David E. Shaw
Chairman
D.E. Shaw & Co., Inc.

University Co-Chair
Karen Holbrook
President
The Ohio State University

PRIVATE SECTOR

Lalitha Subramanian
Director, Contract Research Services
Accelrys

John F. Sparks
Director, Engineering & Technology Programs
Aerojet

J. D. Thompson
Technical Services
ATK Thiokol Propulsion

William Madia
Executive Vice President, Laboratory Operations
Battelle Memorial Institute

Hank Kafka
Vice President,
Architecture and Emerging Technology
BellSouth

Frederick H. Hausheer
Chairman and Chief Executive Officer
BioNumerik Pharmaceuticals

Suresh N. Shukla
HPC Service Manager
The Boeing Company

Donald Paul
Vice President and Chief Technology Officer
Chevron Corporation

Keshav Pingali
Associate Director, Cornell Theory Center
Cornell University

Rich Grzybowski
Director, Systems Engineering
Corning, Inc.

Peter Ungaro
President
Cray Inc.

James Shanks
President, CDW Government
CDW Corporation

John Picklo
Manager,
Mainframes and High Performance Computing
DaimlerChrysler

Ed Leonard
Chief Technology Officer
DreamWorks Animation SKG

David Pensak
Chief Computer Scientist
E.I. DuPont de Nemours & Company

Paul Bemis
Vice President Marketing
Fluent Inc.

Vincent Scarafino
Manager, Numerically Intensive Computing
Ford Motor Company

Stewart Moore
Principal Engineer
General Dynamics Electric Boat

Sharan Kalwani
CAE HPC Manager
General Motors

Stephen Squires
Chief Science Officer
Hewlett-Packard Company

Cindy Gurney
Incubation Center Director/AVP, HNTB Technology
HNTB LTD.

David Turek
Vice President, Deep Computing
IBM

Earl Joseph
Research Vice President,
High Performance Systems
IDC

Guru Bhatia
General Manager and Director for IT Engineering Computing
Intel Corporation

Ben Bennett
Director, HPC Programs Office
Intel Corporation

Eric J. Pitcher
Vice President, Technical Marketing
Linux Networx

Pradeep Raj
Senior Manager,
Advanced Development Programs
Lockheed Martin Aeronautics Company

Ed Seidel
Director, Center for Applied Information Technology and Learning
Louisiana State University

Jeffrey Birnbaum
Managing Director,
Global Head of Enterprise Computing
Morgan Stanley

Timothy Lyons
Executive Director,
Office of the Chief Technology Officer
Morgan Stanley

Reza Sadeghi
Vice President, Enterprise Computing
MSC.Software Corporation

Stan Ahalt
Executive Director
Ohio Supercomputer Center

Kevin Morooney
Senior Director, Academic Services and Emerging Technologies
The Pennsylvania State University

Beverly Clayton
Executive Director
Pittsburgh Supercomputing Center

Greg Brandeau
Vice President of Technology
Pixar Animation Studios

Saadat Syed
Senior Fellow
Pratt & Whitney

Thomas Lange
Director, Corporate R&D Modeling & Simulation
The Procter & Gamble Company

Michael Vildibill
Director, High-End Computing High Performance & Technical Computing
Sun Microsystems, Inc

COUNCIL ON COMPETITIVENESS HIGH PERFORMANCE COMPUTING ADVISORY COMMITTEE

Anthony Robbins
President of SGI Federal
Senior Vice President, North American Field Operations
SGI

Daniel Reed
Vice-Chancellor for Information Technology and CIO
University of North Carolina

George Vradenburg, III
President
Vradenburg Foundation

PUBLIC SECTOR

Robert Graybill
Program Manager
Defense Advanced Research Projects Agency (DARPA)
U.S. Department of Defense

Dimitri Kusnezov
Director, Office of Advanced Simulation and Computing
U.S. Department of Energy

Daniel Hitchcock
Senior Technical Advisor, Office of Advanced Scientific Computing Research
U.S. Department of Energy

Peter Freeman
Assistant Director for CISE
National Science Foundation

Dona Crawford
Associate Director Computation
Lawrence Livermore National Laboratory

John Morrison
Division Leader, Computing, Communications, & Networking
Los Alamos National Laboratory

Stephen Younger
Senior Fellow, Theoretical Division
Los Alamos National Laboratory

Thomas Zacharia
Associate Laboratory Director, Computing and Computational Sciences Directorate
Oak Ridge National Laboratory

George S. Michaels
Associate Laboratory Director, Computational and Information Sciences Directorate
Pacific Northwest National Laboratory

William Camp
Director, Computation, Computers, Information and Mathematics
Sandia National Laboratories

COUNCIL ON COMPETITIVENESS MEMBERS

COUNCIL ON COMPETITIVENESS BOARD

Chairman
Charles O. Holliday, Jr.
Dupont

Chairman Emeritus
F. Duane Ackerman
BellSouth Corporation

Vice Chairman
G. Wayne Clough
Georgia Institute of Technology

President
Deborah L. Wince-Smith

EXECUTIVE COMMITTEE

Robert R. Bishop
Silicon Graphics, Inc.

Molly Corbett Broad
University of North Carolina

William R. Brody
Johns Hopkins University

Jared L. Cohon
Carnegie Mellon University

John J. DeGioia
Georgetown University

Gary T. DiCamillo
TAC Worldwide Companies

Robert C. Dynes
University of California

John A. Edwardson
CDW Corporation

John M. Engler
National Association of Manufacturers

William D. Green
Accenture

Sheryl Handler
Ab Initio

John L. Hennessy
Stanford University

Susan Hockfield
Massachusetts Institute of Technology

Shirley Ann Jackson
Rensselaer Polytechnic Institute

Henry A. McKinnell, Jr.
Pfizer Inc

John P. Morgridge
Cisco Systems, Inc.

Samuel J. Palmisano
IBM Corporation

Michael E. Porter
Harvard University

Luis M. Proenza
University of Akron

Patricia F. Russo
Lucent Technologies Inc.

Kenan Sahin
TIAX LLC

David E. Shaw
D. E. Shaw & Co., Inc.

Lawrence H. Summers
Harvard University

Lawrence Weber
W2Group

DISTINGUISHED FELLOWS

Erich Bloch
Washington Advisory Group

Daniel S. Goldin
The Intellisis Corporation

Alan P. Larson
Former Under Secretary of State for Economic, Business, and Agricultural Affairs

Thomas Ridge
Former Secretary
Department of Homeland Security

GENERAL MEMBERSHIP

Michael F. Adams
The University of Georgia

Richard A. Anthes
University Corporation for Atmospheric Research

William F. Ballhaus, Jr.
Aerospace Corporation

Steven A. Ballmer
Microsoft Corporation

David Baltimore
California Institute of Technology

Craig R. Barrett
Intel Corporation

Richard A. Bendis
Innovation Philadelphia

Robert J. Birgeneau
University of California, Berkeley

James H. Blanchard
Synovus Financial Corporation

Lee C. Bollinger
Columbia University

David L. Boren
The University of Oklahoma

Richard H. Brodhead
Duke University

Michael J. Burns
Dana Corporation

George Campbell, Jr.
The Cooper Union for the Advancement of Science and Art

Curtis R. Carlson
SRI International

Albert Carnesale
University of California, Los Angeles

Carol A. Cartwright
Kent State University

John T. Casteen, III
University of Virginia

Larry Cohen
Communications Workers of America, AFL-CIO, CLC

Mary Sue Coleman
University of Michigan

France A. Cordova
University of California, Riverside

Denis A. Cortese
Mayo Clinic

Margaret B. Cozzens
Colorado Institute of Technology

Michael M. Crow
Arizona State University

Ruth A. David
Analytic Services, Inc. (ANSER)

COUNCIL ON COMPETITIVENESS MEMBERS

Denise D. Denton
University of California, Santa Cruz

Peter R. Dolan
Bristol-Myers Squibb Company

Michael V. Drake
University of California, Irvine

Michael T. Duke
Wal-Mart Stores, Inc.

Roger A. Enrico
DreamWorks Animation
SKG Inc.

Marye Anne Fox
University of California, San Diego

Robert M. Gates
Texas A & M University

E. Gordon Gee
Vanderbilt University

Thomas F. George
University of Missouri, St. Louis

James W. Griffith
The Timken Company

Amy Gutmann
University of Pennsylvania

Peter Halpin
World Resources Company

David C. Hardesty, Jr.
West Virginia University

Jack Harding
eSilicon Corporation

Robert Hemenway
The University of Kansas

Adam W. Herbert
Indiana University

Richard Herman
University of Illinois, Urbana Champaign

Karen A. Holbrook
Ohio State University

Jeffery R. Immelt
General Electric Company

Irwin M. Jacobs
QUALCOMM, Inc.

John I. Jenkins
University of Notre Dame

Martin C. Jischke
Purdue University

Carl F. Kohrt
Battelle Memorial Institute

Alan G. Lafley
The Procter & Gamble Company

Douglas J. McCarron
United Brotherhood of Carpenters and Joiners of America

Martin G. McGuinn
Mellon Financial Corporation

W. James McNerney, Jr.
The Boeing Company

Peggy Gordon Miller
South Dakota State University

Richard K. Miller
Franklin W. Olin College of Engineering

James B. Milliken
University of Nebraska

Clayton Daniel Mote, Jr.
University of Maryland

Anne M. Mulcahy
Xerox Corporation

Diana S. Natalicio
University of Texas, El Paso

Mark A. Nordenberg
University of Pittsburgh

Peter O'Donnell, Jr.
O'Donnell Foundation

Constantine Papadakis
Drexel University

Antonio M. Perez
Eastman Kodak Company

Peter G. Peterson
Blackstone Group

James H. Possehl
Republic Financial Corporation

Harold J. Raveché
Stevens Institute of Technology

Steven S Reinemund
Pepsi-Cola Company

Robert L. Reynolds
Fidelity Investments

Michael C. Ruettgers
EMC Corporation

Henry Samueli
Broadcom Corporation

Carl J. Schramm
Ewing Marion Kauffman Foundation

Pamela Sedmak
Ernst Sedmak, LLC

Ivan G. Seidenberg
Verizon Communications

Kevin Sharer
Amgen

Sanford C. Shugart
Valencia Community College

Lou Anna K. Simon
Michigan State University

Albert J. Simone
Rochester Institute of Technology

John B. Simpson
State University of New York at Buffalo

David J. Skorton
University of Iowa

Andrew A. Sorenson
University of South Carolina

Graham B. Spanier
The Pennsylvania State University

Ray Stata
Analog Devices, Inc.

Susan S. Stautberg
Partner Com Corporation

Charles W. Steger
Virginia Polytechnic Institute and State University

Andrew L. Stern
Service Employees International Union

John D. Stobo
University of Texas Medical Branch

COUNCIL ON COMPETITIVENESS MEMBERS, AFFILIATES & STAFF

Matthew J. Szulik
Red Hat, Inc.

Henri A. Termeer
Genzyme Corporation

Lydia Waters Thomas
Mitretek Systems Inc.

Lee T. Todd, Jr.
University of Kentucky

Peter J. Ungaro
Cray Inc.

Steven L. VanAusdle
Walla Walla Community College

Larry N. Vanderhoef
University of California, Davis

George Vradenburg, III
Vradenburg Foundation

G. Richard Wagoner, Jr.
General Motors

Edie Weiner
Weiner, Edrich, Brown, Inc

William C. Weldon
Johnson & Johnson

Jack M. Wilson
The University of Massachusetts

James H. Woodward
University of North Carolina, Charlotte

James Wright
Dartmouth College

Mark S. Wrighton
Washington University

Henry T. Yang
University of California, Santa Barbara

Paul A. Yarossi
HNTB Corporation

Steven G. Zylstra
Pittsburgh Technology Council

NATIONAL AFFILIATES

American Association for the Advancement of Science

American Electronics Association

American Mathematical Society

American Petroleum Institute

American Society for Engineering Education

American Society for Quality

ASME International

Association of American Universities

Association of University Related Research Parks

BITS' Financial Services Roundtable

Community Learning & Information Network

Computer Systems Policy Project

The Conference Board, Inc.

Council for Chemical Research

Council on Governmental Relations

IEEE-USA

Iowa Business Council

The Japan Society

JumpStart, Inc

National Association of Management and Technical Assistance Centers

National Center for Manufacturing Sciences

National Center for Women & Information Technology

NorTech

Oak Ridge Associated Universities

SMC3

*Current as of January 17, 2006

COUNCIL ON COMPETITIVENESS STAFF

Deborah L. Wince-Smith
President

C. William Booher, Jr.
Chief Operating Officer

Debra Van Opstal
Senior Vice President

William C. Bates
Vice President, Congressional Outreach

Merryl Burpoe
Vice President, Global Initiative

Chad Evans
Vice President, National Innovation Initiative

Randall T. Kempner
Vice President, Regional Innovation

Mohamed Khan
Vice President, Information Systems

Betsy Thurston
Vice President for Strategic Development

Suzy Tichenor
Vice President and Director, High Performance Computing Initiative

Jennifer Bond
Senior Advisor, Global Initiative

David Attis
Director, Policy Studies

Melyssa Fratkin
Policy Director, High Performance Computing Initiative

Julie Jordan
Policy Director, International

Gabrielle Trebat
Director, International Policy Studies

Marcy Jones
Executive Assistant

Mary Marchal
Senior Research Associate

Jennifer Mateik
Controller

Lars-eric Rödén
Research Associate,
National Innovation Initiative

Amanda Welch
Manager, Policy Studies

About the Council on Competitiveness

The Council is the only bipartisan, non-profit, national organization that brings together thought leaders from all elements of America's private sector—business, academia, and labor—to promote economic growth and raise the standard of living for all Americans. As the only national organization whose membership is comprised exclusively of CEOs, university presidents, and labor leaders, the Council is strategically positioned to ensure that regional, national, and world economic agendas anticipate and respond to the demands of a competitive environment for global trade, commerce, and prosperity.

Council on Competitiveness
1500 K Street, NW, Suite 850, Washington DC 20005 www.compete.org