High Performance Computing Software Workshop Report

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

Council on Competitiveness

SPONSORS

The Council on Competitiveness thanks the following co-sponsors for their generous support of the High Performance Computing Application Software Workshop:

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

Ohio Supercomputer Center

ISBN: 1-889866-29-6

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.

Design: Soulellis Studio

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

This work was supported by the U.S. Department of Energy, under grant DE-FG26-04NT42101.

TABLE OF CONTENTS

CO-CHAIRMEN'S INTRODUCTION

America finds itself at a unique crossroads, shaped by dramatic shifts in the nature of competition. With technology, talent and capital now available globally, we can no longer compete effectively on traditional cost and quality terms. Simply doing things the way we have always done them will not be enough to sustain leadership in the global market place. In the 21st century, it is innovation—the ability to create new value—that will confer a competitive advantage. Few technologies hold more promise for accelerating innovation than High Performance Computing (HPC).

In the Council's 2004 *Study of U.S. Industrial HPC Users*, nearly 100 percent of the respondents indicated that they considered HPC to be indispensable to their organizations. These companies stated that they would not exist as viable businesses and could not compete effectively without HPC tools. At the same time, these users identified a growing gap between the production quality, HPC application software that is needed and the software that is available. Some firms are highly dependent on legacy codes that are difficult to replace. Also, organizations in the private sector often are unable to scale their code to the capabilities of their hardware. As a result, productivity and competitiveness are stymied as both cutting edge innovations are stalled and users struggle to expand HPC usage by embedding it into their work processes.

In order to better understand the reasons why this gap has emerged, the Council on Competitiveness and the Ohio Supercomputer Center co-hosted a day-long workshop: *Accelerating Innovation For Competitive Advantage: The Need For HPC Application Software Solutions*. Participants examined the effectiveness of current HPC application software business models, as well as technical and resource barriers preventing the production of more usable application software. In a collaborative atmosphere, they outlined a roadmap of actions independent software vendors, universities, and government research establishments can take individually and collectively to address these barriers and ensure that our country has the application software needed to solve our most competitively important problems.

If the United States is going to continue to lead the world in significant new discoveries, new businesses and new industries, we need the breakthrough insights that emerge through the use of HPC. This will require significant advances in application software that will make HPC usable by more companies. Success is critical if we are going to harness fully America's innovation capacity for future economic growth.

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

Karen A. Holbrook
President
The Ohio State University

EXECUTIVE SUMMARY

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, with support from the Department of Energy's Office of Science, co-sponsored a daylong workshop 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 algorithms that are more comprehensive and accurate.
- The 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.

WORKSHOP REPORT

High Performance Computing Is Critical To U.S. Competitiveness

High Performance Computing (HPC) is critical to increasing the competitiveness of American businesses because it accelerates the innovation process. It is innovation—the ability to create new value—that will provide a competitive advantage for American companies in the 21st century. In today's competitive global market, HPC has become essential to accelerating the innovation of those things most important to our everyday lives, such as energy production and conservation, automobile and airplane safety, weather and climate prediction, national security, financial services, and healthcare. HPC assists companies in creating new inventions and products; in designing better, more reliable products, processes, and services; in minimizing the time to build engineering prototypes; and in streamlining production processes and reducing production costs. Along with theory and experimentation, modeling and simulation with HPC has become the third leg of science, and the path to competitive advantage.

In the Council's 2004 *Study of U.S. Industrial HPC Users*[1], nearly 100 percent of the respondents indicated that they considered HPC to be indispensable to their organizations. These companies stated that they would not exist as viable businesses and could not compete effectively without HPC tools. In addition, the No. 1 reason given in the survey for purchasing high-end computers was their ability to run very large and complex computational problems, which companies need to solve in order to have a competitive edge.

Unfortunately, the HPC market for hardware and production quality software occupies only a small niche within the much larger commercial computing market, and sales often do not generate sufficient revenue to support the research and development needed for advanced software development. Further, because it is a niche market, it must compete for entrepreneurs, investment capital, and technical skills with market segments that promise faster growth and richer rewards. In addition, the niche HPC market itself is further segmented between very experienced users exploiting the most powerful systems and a much larger set of newer users that employ smaller systems. The differing needs of these users could affect the pace of market growth.

• The peak of the HPC market, where the most powerful systems are used, is comprised of both a highly experienced research community along with long-time, highly experienced industrial users. These users want more advanced software that can scale and address their most complex problems, such as stochastic modeling. Many of the industrial users are dependent on legacy code provided by ISVs that must be updated to take advantage of today's highly parallel hardware systems. While a few codes have been modified to run on hundreds or thousands of processors, pricing models often make it prohibitively expensive to run problems at this scale. In many cases, though, the

Everyday Things

The Procter & Gamble Company is responsible for producing "everyday solutions." They manufacture the things that people use every day, such as toilet paper, diapers, coffee, potato chips, cosmetics, deodorant, laundry detergent, shampoo, and dog food. What do these products have to do with HPC? More than you think.

For example, predicting how to comb tangled hair, or whether a Tide® bottle will break if it falls from the top of a washing machine, requires solving complicated physics and chemistry problems. The only way to solve these problems is by creating accurate models and simulations of these interactions using HPC. "The problems," stated Thomas Lange, Director of Corporate R&D Modeling & Simulation for The Procter & Gamble Company, "are high deformation, high contact, non-Newtonian, elastic, plastic, viscoelastic, viscoplastic. Everything is nonlinear. Everything is very difficult."

Unfortunately, the type of problems that Procter & Gamble and similar companies are trying to solve requires sophisticated models and advanced applications, which do not currently exist. Without advances in HPC application software, explained Mr. Lange, "products don't get better…your antiperspirant won't work well, it will hurt when you shave, clothes won't be as clean, diapers will leak, and things will cost more than they should."

1 See *Council on Competitiveness Study of U.S. Industrial HPC Users*, available at www.compete.org/hpc.

application software must be rewritten from scratch to take advantage of current hardware systems. The cost to completely rewrite this legacy code or develop the new code that these users want is significantly more than most ISVs can afford. This segment of the niche HPC market is small and has been shrinking, according to market research from IDC. As a result, it cannot generate the returns ISVs need to justify investment. Participants also worried that if it continues to decline, so will the innovative breakthroughs that can only occur through advanced computation on the most powerful systems, and that are needed to successfully compete in the global market. It is these breakthroughs that drive innovation and productivity in the larger market base.

> *"We're hoping to work at two tiers: at the bottom, to increase awareness and simplify usability for [everyday] people... and then again at the top, where you have relationships with national labs and universities."*
>
> — Wood Lotz, CEO, Absoft

- Below the peak of the HPC market is a much larger set of newer users. Many of the new users in this segment often prefer that ISVs focus on adding new features, such as easier user interfaces and better visualization tools to existing software, instead of investing limited research and development dollars to enhance scalability across tens, hundreds or thousands of processors. For the most part, they are using systems with less than 32 CPUs – far less than is currently available in the market. (The exceptions are those firms who came into existence because of the availability of HPC systems and whose problems are "embarrassingly parallel," such as the computer generated animation industry. They are using thousands of processors.) While this segment of the HPC market is growing substantially, according to market research from IDC, participants were concerned that the user base is not "moving up" and advancing in its use of more powerful systems. While users may be solving yesterday's problems faster with HPC, they will not be able to solve tomorrow's more difficult problems that drive competitive growth unless they continually apply more advanced systems. Although this segment of the HPC market is growing, it is only a portion of a niche market, and does not attract large and sustained ISV investment.
- Finally, on the outside of the HPC market are the "never evers" who may not even be aware of the contribution that HPC can make to their competitiveness, or who find HPC too expensive and too difficult to use.

As a result of its small size and user fragmentation, the HPC market has produced insufficient development funding to create the more capable HPC application software tools that research has shown to be feasible. This in turn has hampered the ability of companies to achieve competitive advantage through faster, better, cheaper products and services that are developed and produced through the aid of HPC. Finally, the resultant limits on the ability of the United States to turn scientific strength into competitive high-value products hamper the growth of our economy, creation of high-paying jobs, and the continually rising standard of living.

Surveys commissioned by the Council on Competitiveness[2] indicate that the weakness of current HPC applications software is a particular bottleneck to more aggressive use of HPC by U.S. corporations. Accordingly, on July 14, 2005, the Council on Competitiveness and the Ohio Supercomputer Center, with support from the Department of Energy's Office of Science, hosted a daylong workshop to further examine the current lack of the production quality HPC applications software that users need. By bringing together HPC users, independent software vendors, university researchers, hardware vendors, and government scientists and engineers, the Council sought both to identify the current barriers to the creation of this HPC software and specific ways to overcome them. This meeting was the first opportunity for many of these stakeholders to come together and discuss the challenges facing corporate HPC users as well as the HPC software industry that creates much of the application software.

The workshop was divided into four sessions. The first session laid out the general business model for industrial applications software, including the middleware that supports it. The second and third sessions explored the technical and resource barriers limiting progress in developing more

2 See *Council on Competitiveness Study of U.S. Industrial HPC Users* and *Council on Competitiveness Study of ISVs Serving the High Performance Computing Market* available at www.compete.org/hpc; See Appendix (page 27) for Executive Summary of the ISV Study.

advanced HPC solutions. The final session addressed the opportunities for progress, drawing on the conclusions from the preceding panels. This summary report reflects the major themes and specific recommendations that emerged during the workshop.

Users Need More Capable HPC Application Software

Solving industry's complex problems requires not only the most advanced hardware, but also sophisticated application software. Historically, users have had various options for obtaining application software. The university and national laboratory research communities have been pioneers in pushing the limits of HPC through the development of novel research codes and algorithms used to address advanced scientific research and mission-critical problems. The ISVs have commercialized some of this research software for the broader industry user community. ISVs also have developed unique application software and middleware for the HPC user market. Finally, some companies choose to develop strategic HPC application software in-house so that they are not reliant on third-party software that also is available to their competitors.

> *"Over the years, we've viewed hardware as an enabler, and the breakthroughs come from application software that can leverage the latest hardware."*
>
> — Stan Posey, HPC Industry and Business Development, Silicon Graphics, Inc.

Unfortunately, a variety of market forces and technical challenges in recent years have caused the ISVs to pull away from creating new and innovative HPC applications, and other developmental paths cannot satisfy the market need. 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. Without the necessary application software, American companies are losing their ability to aggressively use HPC to solve their most challenging problems, and are potentially ceding leadership in the global marketplace.

Repeatedly, the corporate HPC users at the workshop expressed frustration at not having the tools and applications they need. They are looking for application software with easy to use interfaces that can be placed directly into the hands of their engineers, designers, and analysts; and software that will allow better analysis and/or augment or replace traditional testing and evaluation processes. The more experienced users would like software that will run stochastic models to test out the many variables that affect a problem instead of testing the effect of a single variable.

What Is Stochastic Modeling?

In the real world the variables that determine the answer to a problem – such as temperature, pressure, materials properties, etc. – can have a range of possible values. This is one reason why a real world problem can have a range of answers. But a single run of a computer model assigns a single, exact value to each variable and delivers a single answer.

To match real world situations, many runs of the model are required. For each run, the modeler changes the values of the variables so that all runs taken together map out the expected range of the variable. The answers produced by the runs then map out the range of expected answers.

For example, in chemical reactions the temperature usually affects the results strongly, yet it can't be controlled precisely. If the nominal temperature is 90° F, the modeler might make ten runs with the temperature set at 89.5, 89.6, 89.7, ...90.4 and plot the difference in the predicted reaction products to give the chemist a sense of the real world answer.

The word stochastic comes from the Greek word stokhastikos, meaning pertaining to chance. Stochastic modeling attempts to deal with chance in this fashion.

For example, the physics and chemistry involved in making consumer products and consumer products packaging better and less costly is very complicated, and would benefit from models that take into account the complexities and uncertainties of real-life manufacturing and use (see *Everyday Things* on page 9). Unfortunately, most HPC applications run models that solve idealized and simplified problems. To get realistic answers, these models have to be run many times, spanning the range of each variable that is expected to actually occur. In other words, the models must be stochastic.

Thomas Lange, Director of Corporate R&D Modeling & Simulation for The Procter & Gamble Company, explained the importance of stochastic modeling to his company: "To get in and really understand the physics, you have to do it stochastically. And this gets to be a massive problem. To get realistic answers, we've got to run [the problem] lots of times, varying all the variables. If I could do a thousand case runs and fully characterize the variability map, I could go back to my experimenters and say, 'Not only have I done the one case, I have modeled this entire variability space. And this is not only the average response that you should expect when you do the experiment. Here is the distribution of the errors.' Then the believability of that particular piece of advice goes up by about a factor of ten."

Unfortunately, many ISVs have turned their focus away from improving the capability of HPC software and are focusing instead on the desktop computing markets, which are experiencing larger growth and where they can expect to see a better return on investment. Still others only can afford to make minor changes to existing application codes, rather than create novel applications and solutions. As a result, in many cases, the software that is available is outdated and cannot scale to the level needed to solve many industry problems. In other cases, the applications that companies require simply do not exist.

Barriers to Providing HPC Software That Users Need

To better understand why U.S businesses are having difficulty obtaining the production quality application software they need, workshop panelists discussed the pressures affecting both users and software developers. They echoed and often magnified the findings from the *Council on Competitiveness Study of ISVs Serving the High Performance Computing Market*[3] and the comments made during the afternoon software panel at the Council's HPC Users Conference[4]. Multiple themes emerged, and key barriers were identified, which help to explain why the current HPC software enterprise is not improving existing HPC software or creating new applications at the pace the market desires. These themes include market-based as well as technical issues that are highlighted in the following sections.

Market and Resource Barriers

A variety of market and resource barriers are retarding development of scalable, production quality application software

- **The Niche HPC Market Discourages Development of Third Party HPC Application Software**

  Despite the rise in global competition and the need for HPC to help accelerate innovation and increase productivity, HPC remains a niche market with much smaller revenue opportunities than the larger commercial computing market. This lowers the research and development capital available to develop next generation HPC application software and lessens the opportunity for ISVs to obtain investor funding. It also heightens the risk that some ISVs may be bought-out, which could restrict or eliminate access to their software for U.S. commercial and national security organizations that depend on it.

> *"We can't say 'niche' often enough. I think this community sometimes misunderstands what a niche is...Our challenge is further exacerbated by the fact that this is a very small market, with a multiplicity of requirements that are disproportionate to its size."*
>
> — Chris Doehlert, President & CEO, Etnus LLC

As for-profit companies, many of which are public and subject to quarterly earnings reporting requirements, the ISVs often must pursue the broader commercial computing market with its larger revenue opportunity. In addition, according to the Council's 2005 *Study of ISVs Serving the High Performance Computing Market*, more than a third of the ISVs that responded qualify as small businesses, earning less than $5 million annually. Even if these companies were to invest 10 percent of their revenue in research and development, this would be sufficient to add only a few new features to already existing HPC applications. Fundamental rewrites of application software or the

3 See Appendix (page 27) for Executive Summary of the Study. See *Council on Competitiveness Study of the ISVs Serving the High Performance Computing Market*, available at www.compete.org/hpc.

4 See *Conference Report from the 2nd Annual HPC Users Conference*, available at www.compete.org/hpc.

creation of new software require far more capital than is available.

The niche status of the HPC market also reduces the potential for outside investor funding for the ISVs that serve the HPC community, eliminating another possible source of research and development resources. "Right now, if you want to glaze over the eyes of a venture capitalist, mention that you're gong to do a software play in HPC," warned Scott Metcalf, President and CEO of PathScale, Inc. "You'll put them to sleep instantly – if they don't actually just physically throw you out of the room."

Additionally, unlike software developers who serve the broad commercial computing market and can keep their prices low by amortizing their research and development costs across a large user base, the ISVs that serve the HPC market do not have a correspondingly large customer set against which they can amortize their expenses. As a result, the ISVs must shoulder the cost of a major software redesign or spread it across a small user community, significantly increasing the price of the HPC application software. According to Chris Doehlert, President & CEO of Etnus LLC, these pressures are felt particularly by the independent software tools vendors, who create "niche tools for a subset of customers within this niche market."

Another potential consequence of the niche HPC market and the correspondingly small size of the ISVs could be the exposure of the ISVs to potential buy-outs by either domestic or foreign entities who might remove the software from the market for their own proprietary use. This loss of knowledge would weaken our HPC infrastructure, compromising U.S industrial and military advantage. Mr. Doehlert reminded the workshop participants "your greatest security against that is acting in ways that allow me to keep making a viable and faster growing business.... Don't just fund research – purchase the products that you really need at fair prices."

An opportunity may be emerging to grow the HPC user base dramatically and increase the revenue available to advance application development: There is a "tsunami" of data output that is expanding the need for HPC among current users as well as potential new users. New sensors are capturing huge volumes of data that must be processed, analyzed and managed. The oil and gas industry for example, anticipates a thousand-fold increase in sensors. In addition, the cost of data storage has fallen, leading to larger data repositories that can be used for more in-depth analysis and better simulations. New information-driven applications could take advantage of advanced networking as well as distributed computing, and could represent a significant potential market for HPC. "Data increasing by three orders of magnitude in the next ten years, and then some other means of doing something with the data, those things are going to dominate traditional simulation, just dominate it," predicted Dr. Donald Paul, Vice President and Chief Technology Officer of Chevron Corporation.

> *"It's a matter of how you cast competition...small, private companies [worth] five or ten million dollars are eminently acquirable without constraint by anybody in the world. I can't stop that, so I have to exist with it."*
>
> —David Turek, Vice President, Deep Computing, IBM

> *"Huge data-intense problems are going to change the balance of hardware and software architectures."*
>
> — Dr. William Camp, Director, Computation, Computers, Information and Mathematics, Sandia National Laboratories

- **Traditional ISV Business Models are Perpetuating the Niche Market**

In today's market, ISVs typically sell unlimited annual-use licenses, but often on a per-user or per-processor basis. Users often find that they cannot afford the additional software licensing fees that accrue when they want to run their problems over more CPUs. This discourages them from running larger, more complex problems on their current systems or from adding additional CPUs. As a result, users don't exploit the applications to the desired and needed degree. Mr. Lange explained that the current ISV business practice of pricing per CPU is preventing

Procter & Gamble from running the stochastic models his firm needs to understand better the science underpinning their products. "My main reason [for not running stochastic models] is I can't afford the license cost that it takes to run that many cases. The pricing model doesn't match my current computing hardware environment right now. When I pay per processor, that doesn't get it done. I understand that we're in transition. And so there isn't any animosity on my part. But the reality is the pricing model does not allow me to use [our HPC system] to the degree I would like."

In addition, annual license fees can be prohibitively expensive for the small and mid-sized companies that dominate the U.S business landscape – the same companies that represent the potential growth market for HPC. Most cannot afford HPC systems and software, nor do they want do deal with the complexity. Many only want access to HPC systems and software on a periodic basis for exploratory and/or intermittent use. In addition, they often need access to expertise. "Small and medium sized companies can't hire the really expensive people and can't afford the upfront costs," observed Paul Bemis, Vice President of Marketing at Fluent Inc. "They're what we call 'never evers'. They've never ever used the software before, so they have no proof point that it's going to improve productivity. This has to change." The current ISV business model of selling annual software licenses does not meet the needs of this untapped market. "We need to figure out a way to deliver [HPC] cost effectively in a pay-as-you-go model that allows them to experiment with it as small, up-front investments." said Mr. Bemis.

"The analogy is you're going to rent a car for the weekend. You go to Hertz…and they say, sorry, all we offer are annual leases. But you only need the car for the weekend. Sorry, they say again, all we offer are annual leases. That is the ISV community today. It needs to change. We need to be asking, 'How much do you need and how long do you need it?' "

— Paul Bemis, Vice President Marketing, Fluent Inc.

An interesting service model has evolved in the oil and gas industry that may provide an analog. This industry is dominated by small, independent businesses that are the primary producers of oil and gas in this country. These small companies rely on service firms to perform their complex seismic computations and reservoir simulations. The service firms, populated with experienced engineers and scientists, have bundled their knowledge of the industry, their technical knowledge to both create and use HPC software, and access to HPC systems and created very profitable consulting businesses. This business model produces much higher margins than software development alone because it delivers what the small oil and gas companies really want: access to industry expertise and advanced HPC tools in order to make decisions, without having to invest in the underlying technology. In addition to enhancing ISV revenue potential, it permits even the smallest oil and gas company to match the technical sophistication of a larger competitor.

- **Many ISVs Cannot Afford to Acquire Needed Parallel Systems or Support Multiple HPC Platforms**

The small size of many ISVs severely limits their ability to acquire even a single large parallel system, much less several systems from different vendors. Yet this access is essential for developing codes that will run on large systems and for testing codes on different vendor systems.

Hardware vendors that once donated machines to ISVs for software development, testing and support can no longer afford to do so because of declining profit margins. Instead, ISVs must now set aside significant funds to acquire them. This capital expenditure comes at a price. Even the largest ISV may be forced to choose between purchasing hardware or hiring the people they need to tackle the next software development challenge.

"We're not used to setting aside millions of dollars a year to build HPC systems. We would rather go hire the next Ph.D. to extend the reach of Virtual Product Development and address the scalability requirements of our solutions."

— Dr. Reza Sadeghi, Vice President, Enterprise Computing, MSC.Software Corporation

In addition, limited funds make it difficult for many ISVs to support multiple computer platforms. They therefore tend to focus on the cluster segment of the HPC market because the decreasing price of processors has driven many HPC users to this particular architecture. Some users expressed concern that alternate platforms that may be better suited for some problems are being ignored.

> *"I think the fundamental problem is that we're all driven to using the cheapest system, which today is some variant of a PC-cluster ...ISVs have to support that. And that makes it almost impossible for them also to design for a better system."*
>
> — Dr. Robert Lucas, Computational Sciences Division Director, University of Southern California Information Sciences Institute

- **A Shortage of Skilled Talent Is Holding Back HPC Software Development**

All of the software development communities represented at the workshop – national laboratories, universities, ISVs, hardware developers, and users – are in need of more skilled workers. However, a close examination of the current education pipeline presents a troubling picture regarding the availability of next-generation scientists and engineers. American-born students are choosing to pursue computational science and engineering less frequently than in the past. The research institutions face an added challenge in that the development time of HPC application codes may exceed the length of time that a graduate student spends at their facility.

> *"The people we need to architect these codes are unique individuals – they have to understand the science domain, computational science, computer science, and then figure out how to make these big applications run on a parallel computer and to scale."*
>
> — John Morrison, Division Leader for Computing, Communications & Networking, Los Alamos National Laboratory

Recently, foreign-born students who have obtained degrees in science and engineering disciplines at U.S. universities have begun returning to their home countries after graduation. Stricter visa requirements enacted since 9-11 have made it more difficult for those foreign-born students who want to stay. Because H-1B visa applicants must maintain continuity of employment for two years with the same employer, small and medium sized ISV companies are finding that these students would rather work for larger, more stable corporations. This situation is creating a shortage of skilled talent that is holding back both the development of HPC applications and the use of HPC. Further complicating the issue for the laboratories and industries that do sensitive work for the government is the requirement that the workers must be American citizens and cleared for top-secret access. The scarcity of available talent, particularly when limited to American-born workers, significantly hampers both the development and adoption of HPC application software.

On a positive note, trends in desktop computing may bring new talent to bear on this problem. Multi-core processors – technology that allows multiple processors to work together on a single chip—already are finding their way to the desktop computers of many workers, and could stimulate development of more parallel software. Multi-core processors provide significant benefits to industry, including shorter communication paths and improved memory bandwidth. As this technology proliferates in the future, every desktop computer is likely to become a parallel computer. The broader commercial software development community will be forced to address the same challenge that the HPC community has struggled with for years: creating applications that work in a parallel environment. This should bring extraordinary energy, money and *talent* to the software development process, resulting in new tools and new solutions for HPC users, too.

Technical Barriers

In addition to market and resource barriers, HPC users and software developers also face a number of technical barriers that are impeding development of more advanced HPC application software and middleware.

- **Software from the Research Community is Not Adequate for Industry's Needs**

While national laboratories and universities spend hundreds of millions of dollars every year on fundamental HPC software research, it tends to be mission focused and difficult to use in a corporate production environment. Research software typically is developed by Ph.D.s for use by highly experienced researchers who are willing to put in extraordinary time and effort to make it work. Research organizations do not develop their applications with industry's needs in mind. Unfortunately, the exchange between national labs, the corporate user community and the ISVs is not as robust as it could be, resulting in many missed opportunities to transition software from the research community to commercial use.

> *"I believe there are literally thousands of new applications out there that do things that the current ISV applications aren't yet capable of. Yet, it appears to be extremely difficult on both sides to move those applications into the ISV community."*
>
> — Dr. William Camp, Director, Computation, Computers, Information and Mathematics, Sandia National Laboratories

Further compounding the problem is a shrinking federal budget that has reduced the amount of government funding available to support joint programs between national labs and the ISVs. Such programs could translate research software onto the desktop for an operational environment within industry. However it is unlikely that the laboratories will reallocate significant portions of their budgets to focus on these kinds of joint endeavors. They could be criticized for misspending funds since there is an assumption that they have been given "just enough" to fulfill their missions.

Additionally, research software developed within the universities is often an abstraction of a real problem, and may be licensed as open source. Therefore, it has not been subjected to the rigorous testing demanded by industry, and may not allow for proprietary extensions to the code that companies require. Consequently, the software developed by these research organizations, while providing an important foundation for commercial software, is often inadequate for immediate use by industry.

> *"Start digging as to where these software packages came from, an awful lot of them have their roots in government-funded development that was commercialized. [But] one of the things in my view that is institutionally different now than it was 30 years ago or even 20 years ago - the amount of resources going into those kinds of activities has been steadily drying up."*
>
> — Dr. Phillip Colella, Senior Staff Scientist, Lawrence Berkeley National Laboratory

- **Some Open Source License Models Inhibit Commercialization of HPC Research Software**

Much of today's ISV HPC software has been derived from prior research codes, and it is important to preserve this software commercialization path. Most HPC research codes are produced by the national laboratories or universities and licensed as open source, i.e., the source code is freely available and may be copied or altered. However, several participants stated that the terms of some open source license models discourage software developers from improving and commercializing these codes, thereby hindering software innovation.

For example, the General Public License (GPL) requires that derivative works also bear the GPL designation, so they automatically become open source as well. This makes ISVs and other software developers reluctant to invest in development, because they cannot keep their improvements proprietary or derive competitive advantage from them. David Turek, Vice President of Deep Computing at IBM, shared the frustration of many software developers: "It's very difficult when we make major investments to try to produce differentiated technology, and somebody says to me 'Really nice. Now can you put it on everybody else's platform in the world and make it cheap or free?' What we have to do given this circumstance is to work with other platform vendors on new business models that satisfy the needs of the end users while also making our own technology investments economically viable."

On the other hand, the Berkeley Software Distribution (BSD) license allows developers to incorporate the open source code into derivative works that are not open source, thereby allowing the derivative works to become proprietary. The participants believed that most successful HPC software that has transitioned from government and universities to industry bore BSD type open source licenses. The software developers then created their own version of the applications but were not required to give away their derivative intellectual property.

> *""When I look at the successful packages that have come out [of the national labs], they all came out with a BSD type of open source license."*
>
> — Al Geist, Corporate Fellow, Oak Ridge National Laboratory

It should be noted that companies who modify GPL-licensed code for internal use *only* are not affected by the requirement to make derivative works open-source; it only applies when the code is distributed to third parties. Also, despite the limitations of the GPL license, a number of ISVs are building businesses based on GPL code. The ISVs are making money providing expert services to users of the codes. However, it should be noted that these ISVs are not in the HPC market.

Many workshop participants expressed general skepticism about the overall appropriateness of open source to HPC applications software and middleware. They noted that open source means "community-supported" or "unsupported" depending on the size of the market for the code. Mr. Doehlert pointed out again that HPC is still a niche market and that support could be limited. "You take a niche market – in my case, for tools. You're going to boil it down to nothing. And there's not going to be somebody who can support it...Open source needs to be applied where it makes sense and the scale could make sense."

Users also pointed out that companies need some entity to guarantee that the software "works", and that open source does not result in "industrial strength" application software entering the market... "something that I would trust my checkbook against," said Mr. Lange. Dr. Paul concurred, noting that users in a production environment in particular are leery of open source applications and would only use them as long as "somebody with a big balance sheet is between them and the source."

• Legacy Code Requirements and Hardware Limitations are Stalling Software Development

For many ISVs, the requirement to maintain and add features to legacy codes is tying up limited development dollars that could be spent updating old code for greater parallelism or creating new applications. While some software vendors do have applications that can scale to hundreds or even a thousand processors, end users often are reluctant to switch from their legacy codes. Changing the way the problem is run might change the results. For example, some companies that use established applications, like MSC.Software Corporation's structural dynamics code NASTRAN, want to be able to re-run a model from 30 years ago and achieve identical results. Re-coding legacy calculations to achieve greater parallelism can change the answers by several percent. Consistency in results is often more important than improved performance. "I believe the industry needs to strongly consider modeling the *real world*. Modeling the real world requires that we enhance our software platforms and that our customers change how they perform Virtual Build, Test and Review," said Dr. Reza Sadeghi, Vice President, Enterprise Computing for MSC.Software Corporation. "Most industries are very comfortable with the way they have done things and will not easily change" because they are afraid the answer may change.

> *"Legacy code does put a huge burden on us for backward compatibility."*
>
> — Wood Lotz, CEO, Absoft

Hardware limitations also are slowing down development of advanced applications. The communications bandwidth—the maximum rate at which data can be transmitted—of PC clusters is limited, making it impossible to scale many applications beyond a few processors, and discouraging ISVs from developing parallel software. Also, data intensive computing is rising rapidly and for these applications, data communications may be as important as the speed of the processor.

Limitations in memory bandwidth, input/output (I/O), and even network bandwidth also are slowing user demand for parallel software and therefore reducing the business rationale for ISVs to invest heavily in advanced applications. "People are smart," said Dr. Robert Lucas, Computational Sciences Division Director, University of Southern California Information Sciences Institute. "They know which parts of their code are bottlenecks, and they decide, 'It's just not profitable [to parallelize]. I can't benefit if I am memory bandwidth-limited or I/O bandwidth-limited or perhaps network bandwidth-limited.' ...People do these calculations. They're fundamental."

Bandwidth Matters

For many codes, the limit on data transfer rates, called bandwidth, determines the execution time of the problem – the processors are mostly idle waiting for data. Limited bandwidth creates bottlenecks in (1) moving data between processor and local memory, (2) moving data to other processors or memory within the computer, (3) moving data to/from external storage (called I/O), and (4) moving data through external networks. Each of these paths has a different bandwidth, determined by the computer designer. Achieving high bandwidth is expensive and difficult, so there are cost pressures towards lower bandwidth. Scaling up codes to use more processors tends to expose bandwidth limitations - an important reason why many codes do not scale well.

Fields such as bioinformatics, pharmaceuticals, and aerospace are driving the need for new solutions in the area of bandwidth. Generating these solutions is becoming essential in a world where the individual user can create terabytes of data. Ultimately, it may force changes in both hardware and software architecture.

- **Current HPC Software Needs Better Interfaces So That It Can Be Integrated Into The Business Workflow**

Businesses can increase dramatically the value of their HPC systems when usage is expanded beyond the Ph.D. researchers and integrated into the workflow of the organization. (Workflow practices denote the organizations, processes, and sequencing that companies use to design, engineer, produce, and support their products and services; workflow practices are derived from internal factors, such as time-to-market or location of facilities, and external factors, such as regulations.) This means that HPC applications and tools must be distributed at the desktop to the "non-experts," i.e., the designers, engineers, analysts and others more directly involved in the production of an organization's product or service. The key to putting HPC applications and tools into the hands of the non-Ph.D.s at the desktop is better user interfaces that are easy to operate.

"How do you get applications into the hands of people who are not going to be the experts? That raises the general question of, how do you set the balance between functionality and interfaces? And when does one trump the other? I think a lot of people might say that the application with the best interface will win almost every time because of the fit into the workflow."

— Dr. Donald Paul, Vice President and Chief Technology Officer, Chevron Corporation

Integration of HPC into workflow practices becomes even more important as competitive pressures drive companies to shorten product cycles. For example, American car manufacturers recently have driven their design and manufacturing cycle for new cars from more than four years to less than two years, largely in response to global competition. Particularly in high-technology businesses, most of the profits flow to the new products that first satisfy customer needs. As little as a six-month delay in product introduction may allow competitors to take market share and force substantial cuts in profit margins.

In this environment, even one day's delay in setting up and running a HPC problem can significantly impact the workflow. Yet, because of the primitive user interfaces in use today, it can take longer to set up a problem than to run it, particularly as the number of processors increases. When users are under time-to-market pressures, they will only run smaller problems whose solutions may not provide the needed competitive boost. As Stan Posey of Silicon Graphics, Inc., said: "I talk to a lot of engineers... I'm always asking the same question: 'Why won't you run your computational fluid dynamics (CFD) job to 100 processors?' Part of the response I get is that 'It doesn't do me any good to get an instantaneous response in my solver when I'm still building the next mesh.' More investment is still needed in automatic mesh generation to help the business case for highly parallel CFD in practice."

The difficulty in producing a good user interface is often underestimated. Many of the workshop participants suggested that the competencies and cultures of people who develop effective human-digital interfaces are quite different from the people who develop the underlying science of the application. "We underestimate how difficult it is to produce a user interface that is good," said Mr. Lange. "The best user interfaces that we use in our world for statistics and for discrete event simulations were actually written by people who were not engineers, who were not scientists, and were not users originally. They were artists and entertainment people...more game-like in their orientations."

> *"If I were an ISV, the first thing I would do if I were going to bring out a first product is make sure that I had my application interface characteristics well-defined so it could plug into any user interface that a client may have already installed locally,"*
>
> — Dr. George Michaels, Associate Laboratory Director, Computational and Information Sciences Directorate, Pacific Northwest National Laboratory

Integrating HPC applications into workflows through better user interfaces is key to expanding the market for HPC. "Many of the companies that you want to interact with already have workflow systems put in place to help them manage their activities," pointed out Dr. George Michaels, Associate Laboratory Director, Computational and Information Sciences Directorate, Pacific Northwest National Laboratory (PNNL). "And so you want to be able to plug into the back side of that the same way that people figure out how to plug into Microsoft Office or a Web browser-based workflow management system."

Into The Field

The oil and gas industry was an early adopter of high-performance computing, using it to conduct critical reservoir simulations. In the early days, these tools were extremely expensive, complex and difficult to use and consequently, companies like Chevron Corporation placed them in the hands of the company's top researchers.

However, as computing became faster, cheaper and easier to use, Chevron began to look for a way to make these tools available to the thousands of engineers who were out in the field everyday. By arming these engineers with advanced, user-friendly HPC applications, Chevron has been able to increase its productivity and competitiveness. In the process, Chevron also has created a new constituency of users that demand this technology to do their jobs.

Perhaps more importantly, because of their "on the ground" insight into the search for oil and gas, these users are providing Chevron's research and software development teams with new ideas and data to build more advanced, accurate, and sophisticated reservoir models. According to Dr. Donald Paul, Vice President and Chief Technology Officer for Chevron, "When we distributed [HPC tools] to every field engineer, we got ideas for research we would have never come up with. We would absolutely never have seen them because the people in the field say 'Yes, this is the most sophisticated model that you could ever have of a reservoir. Unfortunately, it's not actually a real one.' And that feedback provides insights that no research group will ever have."

PNNL did just that when the William R. Wiley Environmental Molecular Sciences Laboratory (EMSL) integrated its entire workflow into the Molecular Science Software Suite (MS3). MS3 includes NWCHEM, a computational chemistry package that is designed to run on high performance parallel supercomputers. NWCHEM and all of its modeling and simulation tools have now been downloaded and installed at more than 1,500 different sites nationwide, which has expanded the market for and application space in computational chemistry. More applications are being developed that will use this common problem solving environment. As noted above, the more the HPC market grows, the more incentive the ISVs will have to create the application software that users want and need.

Recommendations

By the conclusion of the Software Workshop, participants agreed that a vibrant, growing HPC market is the best long-term guarantee that production quality HPC application software will be available. Participants also agreed that wider usage of HPC is key to future economic growth and the nation's ability to maintain a preeminent position in the global marketplace. However, they noted that the differing needs within the HPC user community could affect the pace of market growth.

Workshop participants suggested the following recommendations, to help stimulate long term HPC market growth through near-term accomplishments. Some recommendations are directed to the private sector, and some to the public sector. Recommendations directed to the public sector are not meant to supplant mission critical, national security requirements. Rather, it is hoped that in the process of meeting its mission critical objectives, public sector work can be leveraged to assist the private sector and provide an additional lift to U.S. competitiveness.

Recommendations for ISVs

In order to encourage parallel processing and the use of more sophisticated HPC applications, the ISV industry may need to consider changes to the pricing arrangements in their existing business model, or possibly a paradigm shift to a service-based model, providing the users with software, expertise and access to the appropriate hardware.

ISVs and user companies should partner to conduct experiments to test the business value of HPC. Participants felt very strongly that the HPC community should come together to run experiments that would demonstrate to current and potential users the business benefits of using HPC, and also help ISVs to explore different pricing models in a controlled environment. These experiments should be built around resolving specific challenges and should consider including the following features:

- Alternative ISV software licensing terms and pricing models to encourage scaling and stochastic runs without incurring excessive cost for the user;
- Large industrial projects with goals of substantial cost-savings and shortening of time-to-market;
- Substantial computer time on large, suitably configured systems at modest cost.

If ISVs do not have larger parallel computers, hardware vendors or research supercomputer centers should be encouraged to participate. In addition, federal agencies such as the Department of Energy – through its Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program – and the National Science Foundation should provide computer resources through their computer centers if needed.

Such experiments will help to mitigate some of the current risks and uncertainty for ISVs and HPC users by:

- Testing the elasticity of different license-pricing models;
- Demonstrating to current users the business value of scaling problems to run on larger systems;
- Lowering the barrier to entry for new users by allowing them to try out HPC at reasonable prices and with the assistance of experts.

Both hardware and software vendors and private and public users will benefit from such partnerships.

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. The fundamental expansion of the market for HPC is not only important to the ISV community but also as a driver for national economic development. Therefore, new business models must be developed to push this expansion and deliver HPC applications to a new group of users. In order to encourage parallel processing and the use of more advanced HPC applications, the ISV industry may need to consider changes to their existing pricing models, or possibly a paradigm shift to a service-based business model, providing the users with software, expertise, and access to the appropriate hardware.

> *"We shouldn't underestimate the value to the ISV if you can make a service model work, because what the ISV is doing over time is accreting a disproportionate share of the intellectual capital, two-legged capital. Pretty soon, the ISV has a pool of talent that's connected to the technology...they know how it's used. They know what problems count. They know what the solutions are worth. That collection of knowledge reaches some critical point and you really have something else you can sell."*
>
> — Dr. Donald Paul, Vice President and Chief Technology Officer, Chevron Corporation

A service-oriented model would allow ISVs to combine their code development with technical support and offer a higher value technical services business that would apply HPC to industry's problems. These services would include running and interpreting the client's codes, and working with the client's engineers, designers, or analysts to apply the results to improve products and processes. This model has been successfully applied to seismic analysis and reservoir management in the oil and gas industry.

Another potential business model could be a pricing mechanism that allows users to pay for periodic usage at a reasonable rate. Alternatively, a hybrid pricing mechanism offering users a mix of annual licenses with surge capacity, pay-as-you-go usage could meet the needs of some users. These models would encourage new users to test HPC, and current users to run larger problems with more CPUs. After the workshop, Fluent announced a new pay-as-you-go software program that permits users to set up computational fluid dynamics problems on their desktop and run them remotely on Fluent's HPC system. Fluent then returns the solution to the user at a very low price point. In the end, the market will decide which models can be sustained and which cannot.

HPC ISVs should make their priority the development of easy-to-use interfaces so that HPC applications can be integrated into organizational workflows. The workshop emphasized that for HPC to be widely accepted, it must be usable by engineers or analysts who are not highly-trained computational scientists, and must fit into existing workflow practices. Because a good user interface is critical to the broad adoption of HPC applications by non-experts, ISVs should consider reaching into the design community for help in creating effective human digital interfaces; in effect, outsourcing this part of the development process to those better "wired" for it. In addition, it was suggested that the ISVs begin their design process by defining interface characteristics up front, allowing the applications to be delivered as plug-and-play, thereby broadening the market. To match workflow practices, ISVs should make it as easy as possible to link together different codes or code modules. This would allow designers, engineers and analysts to integrate their work without having to reformulate data or write special data translators in moving from one stage of the analysis to the next. ISVs also should work closely with their customers, seeking early feedback so that they can refine their ideas and products to better meet market needs.

> *"The availability of easy-to-use, shrink-wrapped HPC software will lead to profound gains in commerce, industrial productivity, scientific discovery, and greater national security."*
>
> — Dr. Stanley Ahalt, Executive Director, Ohio Supercomputer Center

Recommendations for Universities

The nation's universities play a key role in developing our future workforce, and computational scientists will be highly sought after in the innovation economy. Yet only a few universities have modified their organization or curriculum to provide integrated education in computational science.

Universities should increase their educational programs in computational science at the undergraduate and graduate levels to meet the needs for skilled technical workers. The *Council on Competitiveness 2004 Survey of Industrial HPC Users* revealed that industry is unable to hire the skilled workers trained in the use of HPC applications that they need. ISVs have expressed similar concerns, citing their difficulty in hiring skilled workers to develop and support their codes and next generation application development. Programs such as the U.S. Department of Energy Computational Science Graduate Fellowship are dependent on the universities to provide quality education, but only a few universities offer comprehensive curricula in computational science, and even fewer have departments or schools organized around computational science. The President's Information Technology Advisory Committee has recommended that "universities and the federal government's research and development agencies must make coordinated, fundamental, structural changes that affirm the integral role of computational science in addressing the 21st century's most important problems."[5]

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 ISV needs where appropriate. The workshop participants were strongly in favor of creating new mechanisms that facilitate knowledge exchange between the research community and the ISVs. Increasing communication opportunities will help to facilitate technology and information transfer in both directions, allowing the national laboratories to transition their software to the ISVs for broad application and the ISVs to communicate their research needs to the labs for possible future research projects. In addition, new communications mechanisms could provide new partnership and research and development opportunities. As a first step toward accomplishing this goal, private sector organizations should sponsor periodic technical meetings aimed at fostering better communication of needs and research advancements among ISVs, users, government laboratories, and university-based researchers. These should:

- Be organized around specific technical themes;
- Include some meetings that feature emerging areas such as data-intensive computing;
- Provide face-to-face exchange of needs and research advancements;
- Offer ways to develop follow-on partnerships to realize opportunities.

> *"The government budgets aren't big enough to do it alone. The industry is under ever more competitive global pressures. The start-up companies can't find the venture capital to invest in this space and so partnerships, to me, seem to be the only viable way out."*
>
> — Scott Metcalf, President and CEO, PathScale, Inc.

Recommendations for Government

The government, through its policies and procurement practices, has the opportunity to stimulate software development and leverage its investment and expertise in HPC software research and development for national competitive gain.

> *""In each of these cases [where software made the transition from research to production quality HPC applications] there was a long term investment by the government to give these packages a chanced to mature."*
>
> — Al Geist, Corporate Fellow, Oak Ridge National Laboratory

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. Many of today's most successful ISV codes have grown out of government-funded research or engineering codes. Long-term government research investments gave these application packages time to mature, and government policies encouraged their developers to transfer the applications to the private sector. However, in recent years, this mode of technol-

5 See *Report to the President on Computational Science: Ensuring America's Competitiveness* http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf.

ogy transfer seems to have declined. There appear to be at least two reasons for this. In basic and academic research, constrained funding and increasing difficulty of code development have made it difficult to bring codes to a maturity level that would interest ISVs to commercialize them. In more applied research, government codes are being brought to great capability and maturity, but their transfer is inhibited by national security restraints on releasing codes and failure to encourage code developers to move into the ISV community through government grants for start-ups, technology transfer leaves of absence and innovative contracting mechanisms.

The federal government should revise its H-1B visa continuity of employment requirements so that ISVs can attract the talent they need. Currently, the holder of an H-1B visa must remain with a single employer for two years. Yet the workloads in small or medium-sized businesses fluctuate more rapidly than in larger businesses, making it difficult for small companies to guarantee this length of continued employment, or foreign workers to accept employment under an H-1B visa. This makes it particularly difficult for ISVs, many of which are small firms, to tap into the talent pool of H-1B visa holders. Given the critical need for more production quality HPC applications software, the value of the H-1B program would be increased by allowing holders to move between companies as long as they have not been terminated for cause, instead of requiring them to maintain continuity of employment for two years with a single employer. This will make it more attractive to work for a small start-up company or small ISV, even though it may not be as stable as a larger, established company.

The federal government should carefully monitor foreign acquisitions of key ISVs. Even though most ISVs have small sales and capitalization, their codes may be indispensable to U.S. industry and national security. Should a foreign entity decide to purchase any of these ISVs and remove their software from the market, the public and private sectors potentially could be denied codes that they use to solve even their most basic problems. The Committee on Foreign Investment in the United States (CFIUS) has responsibility within the government for monitoring sales of companies to foreign-controlled entities, determining any national security implications, and recommending appropriate government action. CFIUS or a similar entity should begin to monitor sales 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. In the past, research codes have been transferred successfully from the government laboratories and universities to industrial use in HPC applications. For the most part, these were released with a BSD-type open source license. This allowed commercial users and ISVs to develop their own version of the applications without having to "give-away" their intellectual property. Workshop participants stated that recent trends towards licensing research codes under the GPL designation (see pages 16-17 for discussion) may inhibit the flow of research codes into commercial products and urged that this not happen without careful consideration.

> *"I don't think that it's always the case that blindly following open source is necessarily in the country's best interest in terms of building a competitive structure. And I say that as a member of the Free Software Foundation, who makes a lot of money off of both BSD and GPL code. I'm an insider, but I'm still going to throw rocks at it because it's not always in our national interest."*
>
> — Scott Metcalf, President and CEO, PathScale, Inc.

The government and its contractors, including universities and laboratories, need to recognize when to stop funding internal development of research software at the end of its research purpose. At that point the funders should "declare success" and transition the software to commercial or community supported open source status. It's not uncommon for institutions that pioneered software to continue to invest in it well beyond the time that their codes are unique and warrant locally funded development. This is an inefficient use of government research funds, and it takes away potential sales from ISVs that are struggling to survive and grow in what are often very small markets.

Conclusion

The workshop confirmed the importance of HPC to U.S. productivity and competitiveness as an aid to design, engineering, and analysis of products, processes, and services. It also reinforced the fact that HPC is a small niche market within the much larger computer market with tough competition for investment capital, skilled workers, and market size.

Despite its potential to accelerate innovation and increase productivity, the workshop concluded that HPC is not making the contributions that it could due to technical and resource barriers, thereby limiting U.S. competitive advances. Corporate HPC mostly depends on ISVs to supply critical software, yet the ISVs are typically small and financially fragile. As a result, many of them have been slow to scale up their codes to take advantage of modern parallel computers, slow to introduce easier-to-use interfaces for non-expert users, and slow to match corporate workflow needs by offering integrated suites of codes. Instead, their limited funds often are spent on maintaining legacy codes and making small improvements demanded by customers. Under investor pressure to grow their revenue, some ISVs are moving away from the HPC market to pursue the much larger commercial market, with its higher potential for financial returns.

Corporate HPC users have been slow to demand or exploit scalable parallel code, in part due to the difficulty of using current software. They need better interfaces so that they can expand usage beyond their Ph.D. researchers and integrate HPC into their business workflow for enhanced productivity. Some corporate practices, however, such as demanding ongoing support for legacy problems, usurp limited ISV research and development funds and constrain ISV innovation.

Although the government-funded research community has made great strides in scaling up research codes, few of its advances have appeared in corporate HPC software. This is partly because research codes are developed for and used by highly trained scientists with little thought for the private sector's needs. Also, these codes are usually developed to solve very limited idealized problems and are poorly matched to corporate needs for software that is easy to use, suitable for real-world problems, and includes robust error-handling routines. In addition, there are no ready mechanisms for the research community, ISVs, and corporate users to share needs, advances, and opportunities in ways that would encourage the flow of models, algorithms and code from the research community into software that meets corporate needs.

All segments of the HPC community – corporate users, ISVs, and researchers – are hampered by a lack of skilled workers. This lack results from curriculum gaps in university education, greater allure of other segments of the IT industry, falling enrollments in math and science courses generally, and increasing immigration barriers faced by foreign students and skilled IT workers.

The workshop offers recommendations to the private sector, to universities and the research community, and to government, all aimed at overcoming the identified barriers and improving the value of HPC. These recommendations include partnering agreements that would mitigate the risk of trying new business models, mechanisms to broaden and deepen the use of HPC, educational improvements, and adjustments that would encourage the flow of research advances into corporate HPC.

> *"By leveraging the investment and expertise of both the public and private sectors, we can harness the innovative capacity of the high-performance computing community and lift our country to a new level of competitive advantage."*
>
> — Dr. Stanley Ahalt, Executive Director, Ohio Supercomputer Center

Implementing these recommendations entails considerable risk to the ISVs, the users, the research community, the government, and investors, but 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.

Conversely, by leveraging HPC to solve complex problems, America can unleash a new era of innovation-driven growth, creating new industries and markets, fueling wealth creation and profits, and generating high-value, higher-paying jobs that will raise the standard of living for all citizens.

APPENDICES

Workshop Agenda

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

Council on Competitiveness High Performance Computing Advisory Committee

Council on Competitiveness Members, Affiliates and Staff

WORKSHOP AGENDA

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

8:00 a.m. **Registration and Breakfast**

8:30 a.m. **Opening Remarks and Setting the Stage**

- **Introduction:** Suzy Tichenor, Vice President and Director, High Productivity Computing Project, Council on Competitiveness
- **Opening Remarks:** Deborah L. Wince-Smith, President, Council on Competitiveness
- **Setting the Stage:** Dr. Stanley Ahalt, Executive Director, Ohio Supercomputer Center and Workshop Chair

9:00 a.m. **Session 1: Effectiveness of Current HPC Application Software Business Model/s**
Moderator: Dr. Donald Paul, Vice President and Chief Technology Officer, Chevron Corporation

Panelists:
- Chris Doehlert, President & CEO, Etnus LLC
- Thomas Lange, Director, Corporate R&D Modeling & Simulation, The Procter & Gamble Company
- Dr. George Michaels, Associate Laboratory Director, Computational and Information Sciences Directorate, Pacific Northwest National Laboratory

10:00 a.m. **Session 2: Technical Barriers to Progress**
Moderator: Dr. William Camp, Director, Computation, Computers, Information and Mathematics, Sandia National Laboratories

Panelists:
- Wood Lotz, CEO, Absoft
- Dr. Robert Lucas, Computational Sciences Division Director, University of Southern California Information Sciences Institute
- Stan Posey, HPC Industry and Business Development, Silicon Graphics, Inc.
- Dr. Reza Sadeghi, Vice President, Enterprise Computing, MSC.Software Corporation

11:00 a.m. **Break**

11:30 a.m. **Session 3: Resource Barriers to Progress**
Moderator: Dr. Phillip Colella, Senior Staff Scientist, Lawrence Berkeley National Laboratory

Panelists:
- Scott Metcalf, President and CEO, PathScale, Inc.
- John Morrison, Division Leader, Computing, Communications & Networking, Los Alamos National Laboratory
- David Turek, Vice President, Deep Computing, IBM

12:30 p.m. **Lunch**

1:30 p.m. **Session 4: Opportunities for Progress**
Moderator: Dr. Stanley Ahalt, Executive Director, Ohio Supercomputer Center and Workshop Chair

Panelists:
- Paul Bemis, Vice President Marketing, Fluent Inc.
- Al Geist, Corporate Fellow, Oak Ridge National Laboratory
- Wood Lotz, CEO, Absoft
- Dr. Donald Paul, Vice President and Chief Technology Officer, Chevron Corporation

3:00 p.m. **Summary and Path Forward:**
Dr. Stanley Ahalt, Executive Director, Ohio Supercomputer Center and Workshop Chair

3:30 p.m. **Adjourn**

EXECUTIVE SUMMARY

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

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

Key Findings

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

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

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."

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

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).

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 applications 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.

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.

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

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

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

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

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.)

Original Source of ISV Application

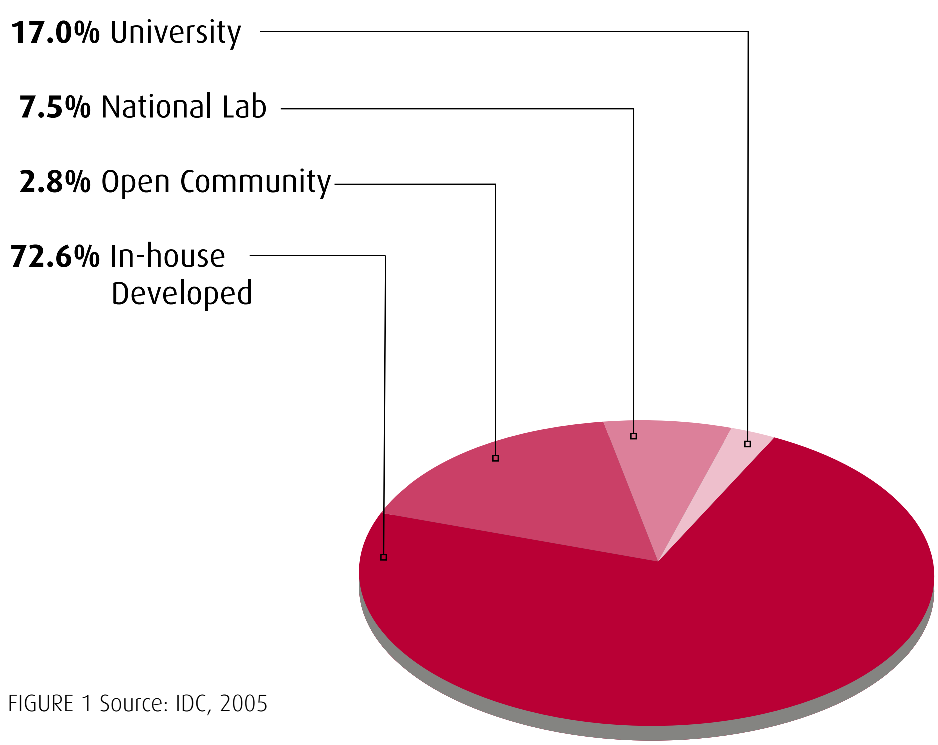

FIGURE 1 Source: IDC, 2005

EXECUTIVE SUMMARY

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

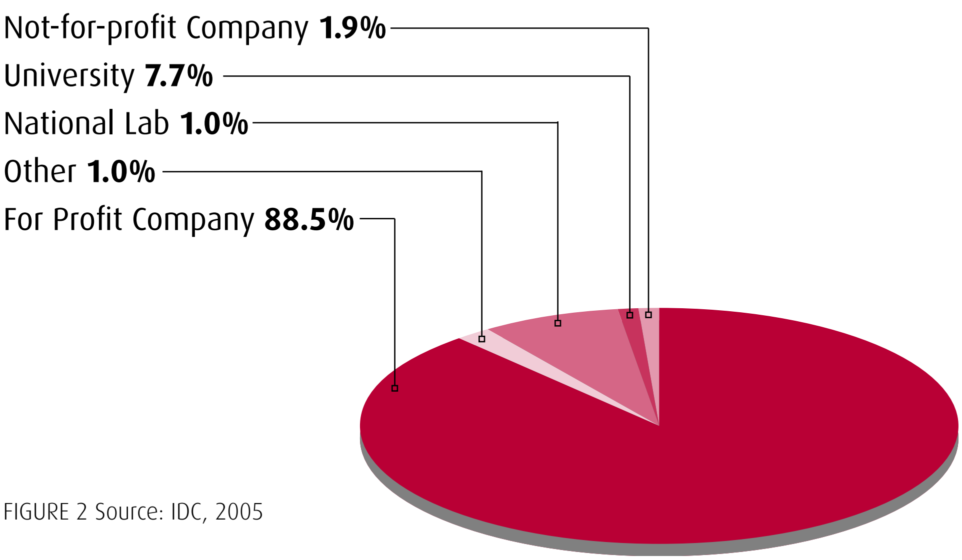

FIGURE 2 Source: IDC, 2005

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.

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

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

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

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.

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

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.)

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

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

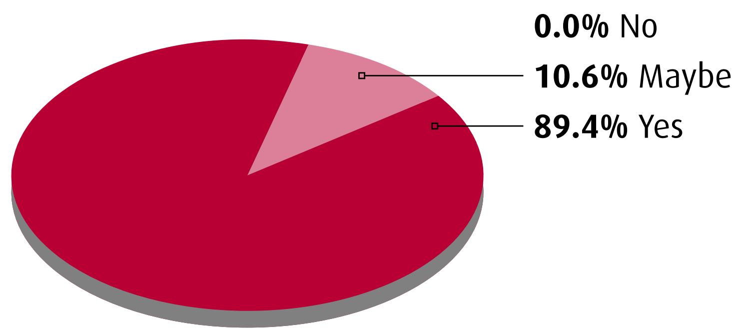

FIGURE 3 Source:

"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."

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, Inc.

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

COUNCIL ON COMPETITIVENESS MEMBERS

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

COUNCIL ON COMPETITIVENESS AFFILIATES & STAFF

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