AFRL-IF-RS-TR-2007-169
Final Technical Report
July 2007

HIGH PRODUCTIVITY COMPUTING SYSTEMS AND COMPETITIVENESS INITIATIVE

CACI Technologies, Inc.

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. Q455

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

AFRL-IF-RS-TR-2007-169 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/

CHRISTOPHER FLYNN
Work Unit Manager

/s/

JAMES A. COLLINS, Deputy Chief
Advanced Computing Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
JUL 2007

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
Sep 03 – Feb 07

4. TITLE AND SUBTITLE
HIGH PRODUCTIVITY COMPUTING SYSTEMS AND COMPETITIVENESS INITIATIVE

5a. CONTRACT NUMBER
F30602-00-D-02210025

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
62303E

6. AUTHOR(S)
Suzy Tichenor

5d. PROJECT NUMBER
Q455

5e. TASK NUMBER
QF

5f. WORK UNIT NUMBER
25

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
CACI Technologies, Inc.
1300 Floyd Ave.
Rome NY 13340

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington VA 22203-1714

AFRL/IFTC
525 Brooks Rd
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TR-2007-169

12. DISTRIBUTION AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 07-AFRL-07-0047

13. SUPPLEMENTARY NOTES

14. ABSTRACT
In the face of increasing competition from other market economies, the US needs to ramp up its ability to innovate—and High Performance Computing (HPC) must be a key ingredient in America's innovation capacity. The task ahead therefore must include broadening the use of HPC in US industry through collaborations in which more knowledgeable HPC users—especially leading government users—share their expertise and advanced computing resources with less-experienced and inexperienced private sector companies. We must move beyond today's "islands of innovation" by making access to HPC much easier and more pervasive. Leveraging these world-class assets to their fullest is critical if we are going to successfully harness America's innovation capacity to drive future economic growth and industrial leadership. In today's high competitive global economy, the country that wants to *out-compete* must be able to *out-compute*. This report summarizes an initiative whose objective was to attempt to assure a path for realizing dual-use applications for next generation high performance computing.

15. SUBJECT TERMS
High Productivity Computing Systems, Competitiveness, Innovation

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
U	U	U	UL	306	Christopher Flynn
					19b. TELEPHONE NUMBER (*Include area code*)

Table of Contents

1.0 Goals:

Analyze economic imperative for sustaining U.S. leadership in High Performance Computing (HPC) by examining its impact on productivity, ***innovation*** and competitiveness in the private sector user community.

2.0 Approach:

Determine whether the private sector is using HPC as aggressively as it could and should. For example, are there technical barriers related to the availability and ease-of-use of hardware and software that are stalling increased usage of this tool across the private sector? Are there business barriers to adopting/using HPC that derive from the internal cost and budgeting structures within companies, management decision-making processes, or corporate investment priorities? Explore the role of public-private sector partnerships to address barriers. Leverage government investment in HPC R&D, systems and expertise to accelerate innovation and advance competitiveness.

3.0 Methodology:

Three-tiered collaborative sifting process incorporating insights gained from

• HPC Executive Advisory Committee comprised of leading public/private sector HPC users, system manufacturers and software developers
• HPC Private Sector User Surveys
• HPC User Conferences

4.0 Key Initiative Findings:

Despite the benefits that companies and the country can accrue from applying HPC, the market remains a niche within the much larger $254 billion commercial computing market. The HPC market represents less than 3% of total commercial computing sales. This niche is further segmented between a small group of experienced users employing very large-scale systems, and a very large group using entry-level systems. However, the users at the entry level are not adopting more powerful systems, resulting in a large "missing middle" within the HPC market/user base. And while these entry level users may be able to run yesterday's problems faster, these companies will 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. *See also Council Report from the Application Software Workshop, Accelerating Innovation for Competitive Advantage: The Need for HPC Application Software Solutions (2005), available at www.compete.org/hpc for further discussion.*

There remains a large group of companies that do not use HPC to drive innovation and accelerate productivity for competitive advantage. These companies are sometimes called "never-evers" because they have "never ever" tapped into the benefits of HPC. *See also Council Report from the Application Software Workshop, Accelerating Innovation* *for Competitive Advantage: The Need for HPC Application Software Solutions (2005), available at www.compete.org/hpc, for further discussion and potential remedies.*

While HPC is essential to the business survival of experienced users, a combination of education and training barriers, along with technical barriers such as lack of easy to use, production quality, scalable application software, are preventing companies from using HPC as aggressively as they could and should. *See Appendix 4: Council on Competitiveness Study of U.S. Industrial HPC Users for further discussion.*

The small size of the HPC market reduces the amount of research and development capital available to develop next generation HPC application software, and lessens the opportunity for Independent Software Vendors (ISVs) to obtain investor funding. This discourages development of third party HPC application software. *See Appendix 2:* 2005 *HPC Users Conference Report: Accelerating Innovation for Prosperity, and Appendix 5: Study of the ISVs Serving the High Performance Computing Market, Part A: Current Market Dynamics for further discussion and potential remedies.*

The ISV business model of selling unlimited annual-use licenses on a per-user or per-processor basis is often cost-prohibitive for users that want to run problems over more CPUs. This discourages them from running larger, more complex problems on their current systems or from adding additional CPUs. See *Appendix 5: Study of the ISVs Serving the High Performance Computing Market, Part A: Current Market Dynamics, and Appendix 6: Study of the ISVs Serving the High Performance Computing Market, Part B: End User Perspectives for further discussion and potential remedies.*

Software produced by the research community is not adequate for industry's needs. It tends to be mission focused and difficult to use in a corporate production environment. *See Appendix 2: 2005 HPC Users Conference Report: Accelerating Innovation for Prosperity for further discussion. See also Council Report from the Application Software Workshop, Accelerating Innovation for Competitive Advantage: The Need for HPC Application Software Solutions (2005), available at www.compete.org/hpc for further discussion and potential remedies.*

There is a growing gap between the HPC resources and expertise resident within the DOE laboratories and those within industry, creating a strategic opportunity for partnerships. Experienced industrial HPC users have large, complex, competitively important problems that cannot be solved today in reasonable timeframes. However, some of these firms could make progress if they had access to the systems and expertise within the labs. *See also Council Report from the Application Software Workshop, Accelerating Innovation for Competitive Advantage: The Need for HPC Application Software Solutions (2005), available at www.compete.org/hpc for further discussion and potential remedies.*

Public-private partnerships like the Department of Energy's (DOE's) Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program can leverage the government's HPC resources to meet both national security needs and provide a competitive lift to the country by helping our most experienced industrial users solve their most complex and competitively important problems. When industrial firms are able to access these systems through partnership programs like INCITE, they gain a "crystal ball" look into their technical future and a head start in preparing their internal processes and programs to embrace it. This could provide a model to enhance the return on the government's investment in HPC assets at little or no additional cost. *See Appendix 2:* 2005 *HPC Users Conference Report: Accelerating Innovation for Prosperity for further discussion.*

University HPC centers and national laboratories can be "anchors" for regional economic development, but must help local, state and regional government officials understand the value of this technology. Additionally, if HPC centers want to assume this role, they must be able to support a broad spectrum of users, from entry-level to experienced, and assist with a wide range of problems. This is particularly true when companies come to the HPC center without a well-defined project or model. HPC centers must be prepared to help companies define the model and even run it on the Center's HPC system. *See Partnering for Prosperity: Industrial Partnerships through the NSF's Supercomputing Resources (2006, and Partnering for Prosperity: Industrial Partnerships through the NNSA's Academic Strategic Alliance Program (2006, available at www.compete.org/hpc for further discussion.*

Corporate supply chains may offer a mechanism by which to expand HPC usage to a broader range of industrial users. However, company opinions and practices vary tremendously. While some companies have reached out to their suppliers to introduce them to modeling and simulation with HPC, others leave it to the supplier to select the best tools and processes to meet customer requirements. *See Appendix 3: 2006 HPC Users Conference Report: Moving Beyond Islands of Innovation for further discussion.*

Optimizing global supply chains for maximum productivity and profitability is an emerging HPC application that could have a profound impact on U.S. industrial competitiveness. *See Appendix 3*: *2006 HPC Users Conference Report: Moving Beyond Islands of Innovation for further discussion.*

5.0 Initiative Accomplishments:

5.1 Launched/Managed HPC Advisory Committee

The Council created, launched and conducted annual meetings of its HPC Advisory Committee, a national "brain trust" of industrial HPC users, universities, computer hardware and software developers, and federal agencies/national laboratories that use and fund development of HPC. This Committee provided the strategic direction for the Council's HPC Initiative, providing critical insight in the drivers of industry's use of HPC, barriers to be addressed, and the role of public-private partnerships to help fill the "expertise" gap.

Advisory Committee members spoke at each of the Council's HPC Users Conferences (see below). In addition, after meetings with Dr. Raymond Orbach, DOE Under Secretary for Science, they spurred the agency to open the *Innovative and Novel Computational Impact on Theory and Experiment (INCITE)* program to industry proposals for the first time. Seven firms have passed DOE's strict peer review process since the program was expanded to include them in 2005. They have been awarded substantial time on DOE's most advanced high performance computers at the Oak Ridge National Laboratory, Lawrence Berkeley National Laboratory, Pacific Northwest National Laboratory and Argonne National Laboratory.

5.2 Conducted Four HPC User Surveys

The Council conducted four surveys of HPC Users and independent software vendors to better understand the links between use of HPC and competitiveness, and the barriers preventing more widespread adoption of this technology.

- Study of U.S. Industrial HPC Users confirmed that HPC is essential to business survival, and drew out the obstacles preventing companies from accessing and using HPC more aggressively to increase productivity and competitiveness.

- Survey of Independent Software Vendors Serving the HPC Market – Part A: Current Market Dynamics revealed that the small size of the HPC market discourages commercial development of the production quality, HPC application software that U.S. businesses need to solve their most competitively challenging problems.

- Survey of Independent Software Vendors Serving the HPC Market – Part B: End User Perspectives explored the U.S. business requirements for advanced HPC application software, and the financial and technical obstacles blocking firms from obtaining it.

- Study of Innovation, Competitiveness and HPC revealed that while many large industrial firms rely on HPC for competitive advantage, many of their suppliers have not adopted modeling and simulation with this important tool. Over time, the larger "customer" firms could be at risk if their suppliers are unable to meet demanding market requirements. It could also place these suppliers at risk if their competitors adopt HPC and are able to better meet customer needs.

These surveys are available via the Council website http://www.compete.org/hpc and are included in this report at Appendices 4-6.

5.3 Conducted Three HPC Users Conferences

The Council conducted three annual HPC User Conferences. While other HPC conferences focus on HPC products, and address R&D challenges to designing and manufacturing more powerful systems and software, these conferences were the first to focus on *applying* HPC resources to drive industrial competitiveness. Conference speakers and panelists discussed applications needs that cut across sectors; and ways the public and private sectors can partner to address barriers preventing more widespread industrial usage of the critical technology. These conferences also provided a public venue to share and discuss the findings from the HPC User Surveys.

Conference speakers and attendees represented senior public and private sector executives. Over the three years, attendance grew 20% to 250 executives from industry, academia and government. A conference report was produced after each event. They have been widely distributed to conference attendees, policy makers, and the broader Council membership, and are available via the Council website http://www.compete.org/hpc. They are included in this report as Appendices 1-3. DVD's of the 2004 and 2005 conferences are also available via the Council website http://www.compete.org/hpc.

6.0 Publications:

- 2004 HPC Users Conference Report: *Supercharging U.S. Innovation & Competitiveness* (see Appendix 1)
- 2005 HPC Users Conference Report: *Accelerating Innovation for Prosperity* (see Appendix 2)
- 2006 HPC Users Conference Report: *Moving Beyond Islands of Innovation* (see Appendix 3)
- *Council on Competitiveness Study of U.S. Industrial HPC Users* (see Appendix 4)
- *Study of the ISVs Serving the High Performance Computing Market, Part A: Current Market Dynamics* (see Appendix 5)
- *Study of the ISVs Serving the High Performance Computing Market, Part B: End User Perspectives* (see Appendix 6)
- *Council on Competitiveness Study of HPC as an Innovation Driver for Industry* (see Appendix 7)

7.0 Complimentary Activities:

7.1 Elevated High Performance Computing to the National Policy Agenda

HPC was mentioned prominently in a 2005 report by The President's Information Technology Advisory Committee, a private sector advisory committee to the President of the U.S. *Computational Science: Ensuring America's Competitiveness* referenced work completed by the Council's HPC initiative numerous times. References included data about the critical importance of high performance computing to national competitiveness and were drawn from the Council's 2004 HPC Users Conference report and 2004 HPC Users Survey. The PDF version of the report is available online at <http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf>.

7.2 Brought HPC Into The Mainstream Dialogue About U.S. Competitiveness

The Council energized leading media outlets such as the New York Times, Business Week, Financial Times, the Economist, and the National Journal, as well as trade press such as Information Week, Computerworld, Federal Computer Week and HPC Wire to publish thought-provoking articles discussing the linkages between widespread use of high performance computing and the ability to compete successfully in the global economy. The Council and members of our HPC Advisory Committee members have been quoted.

7.3 Sponsored Discussion at the International Supercomputing (SC) Conference

Two Council-proposed "Birds of a Feather" (BOF) discussions were selected by the planning committee for the annual, international Supercomputing Conference in 2004 and 2005. This is the leading HPC industry conference in the world. It is held annually in the U.S. and attendance ranges from 8,000-10,000.

- *Leadership in a Global Economy: "Out-compete" means "Out-compute" (SC 04).* Members of the Council's HPC Advisory Committee participated as panelists, discussing the importance of HPC to industrial productivity and economic competitiveness. Articles about the panel and interviews with panelists appeared in a variety of journals including Computer World, ClusterWorld and the National Journal, providing additional visibility for this initiative and further reinforcing our message that high performance computing is critical to U.S. global competitiveness.

- *"The Evaporation of the HPC Application Software Market" (SC 05).* Members of the Council's HPC Advisory Committee participated as panelists, discussing the challenges of creating and maintaining HPC application software for a competitive, corporate "production" environment, the state of the ISV application software market, and the role of government, universities and national laboratories to help accelerate development of new and/or updated code. The BOF served to further publicize the results of the ISV survey, and engaged the larger supercomputing community in a discussion of the future opportunities for HPC software development.

7.4 Produced "HPC in Everyday Life" Video

The Council developed a video entitled, "HPC in Everyday Life," in collaboration with DreamWorks Animation, a major animated motion picture studio and a member of the

Council's HPC Advisory Committee. Other members of the Advisory Committee submitted text, visualizations and video clips, which were woven into the 8-minute film. From medicine to consumer products, from energy security to aerospace, this video reveals how HPC is behind many of the products and services that we take for granted in everyday life, and illustrates the importance of HPC in breakthrough innovation.

The film was unveiled at the 2005 HPC Users Conference and has subsequently been requested by government agencies, universities, vendors and users. It was also shown at the Supercomputing 05 as a part of the plenary presentation delivered by Thomas Lange, Director of Corporate R&D Modeling & Simulation at The Procter &Gamble Company. Approximately 800 people were in attendance to hear his talk.

The DVD is available to the public on the Council's website, at **http://www.compete.org/hpc/everyday_life_video.asp**

7.5 Conclusion

In the face of increasing competition from other market economies, the U.S. needs to ramp up its ability to innovate—and HPC must be a key ingredient in America's innovation capacity. The task ahead therefore must include broadening the use of HPC in U.S. industry through collaborations in which more knowledgeable HPC users—especially leading government users—share their expertise and advanced computing resources with less-experienced and inexperienced private sector companies. We must move beyond today's "islands of innovation" by making access to HPC much easier and more pervasive. Leveraging these world-class assets to their fullest is critical if we are going to successfully harness America's innovation capacity to drive future economic growth and industrial leadership. In today's highly competitive global economy, the country that wants to *out-compete* must be able to *out-compute*.

Conference Report

First Annual
High Performance Computing Users Conference

Supercharging U.S. Innovation & Competitiveness

July 13, 2004, Washington, D.C.

Council on Competitiveness

Cover Story:

The Council on Competitiveness thanks Tom White, of tomwhite.images-illustration & design, New York, NY, for the cover illustration which will be used throughout the Council's High Performance Computing Initiative.

It depicts and embodies the human mind, spirit and visionary thinking behind countless innovations. This creative and competitive nature will continue to drive the way we think and invent.

thank you

The Council on Competitiveness thanks the following cosponsors for their generous support of the First Annual High Performance Computing Users Conference: *Supercharging U.S. Innovation and Competitiveness*.

Sponsors

Defense Advanced Research Projects Agency

Department of Energy National Nuclear Security Administration

Department of Energy Office of Science

ISBN 1-889866-18-0

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, DC 20005

Tel: (202) 682-4292
Fax: (202) 682-5150
E-mail: council@compete.org

To learn more about the Council on Competitiveness, 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 contributions.

Table of Contents

12

Introduction

1st Annual High Performance Computing Users Conference, Washington, D.C., July 13, 2004

Today, with technology, talent and capital available globally the United States is facing unprecedented competitive challenges abroad. So how does the U.S. compete and prosper?

We strongly believe that innovation is the bedrock of America's competitiveness. It drives gains in productivity and market share that guarantee a rising standard of living for all Americans.

However, an economy built on innovation can only succeed if it has access to and uses the best tools to drive the innovation process. And there are few areas of technology that hold more promise for stimulating innovation and propelling competitiveness than high performance computing (HPC). Indeed, we believe that the country that out-computes will be the one that out-competes.

Yet, American businesses face a dilemma when it comes to the use of HPC to advance productivity and innovation. There's widespread recognition that both longtime and emerging HPC users face a variety of barriers to using this innovation-driving technology to its full potential.

In order to better understand these barriers, and to begin a public-private sector dialog to address them, the Council on Competitiveness—in partnership with the Defense Advanced Research Projects Agency (DARPA) and the Department of Energy's National Nuclear Security Administration and Office of Science—convened the First Annual HPC Users Conference, in Washington, D.C., on July 13, 2004.

The conference was unique. During the course of a day, senior government officials responsible for HPC funding and development, experienced business HPC users, and senior academic HPC users—more than 200 participants in all—shared their experiences and vision through engaging interactive panels, keynote addresses, and whole group discussion. They provided inspiring examples of the productivity gains possible with HPC and the "grand challenge" HPC opportunities that American companies believe can fuel their competitiveness. And they frankly identified the business and technical barriers limiting their HPC use. Conference participants were able to put their discussion in context through hearing the results of the Council's first annual National HPC Users Survey.

The message from the day's proceedings was clear. If the U.S. is going to reach the heights at which new businesses and industries are created, we need the breakthrough insights that emerge from advanced HPC applications.

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

Karen A. Holbrook
President
The Ohio State University

Executive Summary

Executive Summary

1st Annual High Performance Computing Users Conference, Washington, D.C., July 13, 2004

The 1st Annual High Performance Computing Users Conference brought together more than 200 senior government, business and academic high performance computing (HPC) users and policy makers to discuss the enormous opportunities available from more widespread use of HPC, and identify the barriers limiting private sector application of this technology in America today.

The current focus on HPC and its linkage to economic competitiveness underscores the convergence of a number of powerful global forces. These include the commoditization of HPC technology, the globalization of business operations and lightning speed competition, and the recognition of the enormous potential of HPC to transform business functions, particularly in the large services sector, as well as in advanced manufacturing.

Senior government officials confirmed that HPC resources are essential to solving problems critical to U.S. national security interests, and organizations such as the Defense Advanced Research Projects Agency (DARPA) and the Department of Energy (DOE) play a critical role as investors and leading edge users. But the full benefits of advances in high performance computing have not been captured by the private sector.

This observation was amplified in the results of the first HPC Users Survey, released at the conference. The majority of businesses surveyed identified HPC as an indispensable innovation tool, but noted that their companies were not using HPC as aggressively as possible.

Conference panelists reinforced this message. They noted that HPC provides significant competitive advantage to current industrial HPC users. These advantages include reduced design and analysis R&D costs through virtual prototyping, shorter time to market, and even the creation of new industries, such as the digital animation sector. However, they emphasized that their companies need improved HPC resources to solve a range of currently intractable, major R&D problems, and thus meet growing international competition.

Barriers to private sector HPC use were identified in three key areas:

- *Business culture: In the boardrooms of many American companies HPC isn't seen as an innovation edge, but rather a cost of doing business with no clear idea of the actual return on investment.*
- *Educational: Many industries don't have the people who can productively use HPC to its full innovation potential.*
- *Technical: These include the essential need for easier-to-use HPC software and improved mathematical models.*

In summary, there was clear agreement on the transformational capabilities of HPC to enhance business productivity and boost national economic security. Key next steps proposed for supercharging U.S. innovation with HPC included: renewing and creating new government-industry-university partnerships, fueling next-generation HPC simulations, and improved linkage between the HPC skills and knowledge taught by universities and those required by businesses.

Building on the Conference momentum and findings, the Council on Competitiveness HPC Advisory Committee is creating an Action Agenda, including a focus on new 21st century government-private sector partnerships. Partnerships will form a key part of discussions at the 2nd High Performance Computing Users Conference, planned for July 13, 2005.

Conference Proceedings

Context: Innovation is Imperative

Council on Competitiveness President Deborah L. Wince-Smith set the context for the conference with opening remarks that framed the event as part of the Council's broader innovation agenda. The ability of the U.S. to create, innovate and improve more rapidly than its global competitors will be the key to ensuring productive growth, and guaranteeing an increased standard of living for all citizens. "By shrinking 'time to insight' and 'time to solution' through the use of high performance computing, we can accelerate the innovative process in ways simply not seen in the past," Wince-Smith explained.

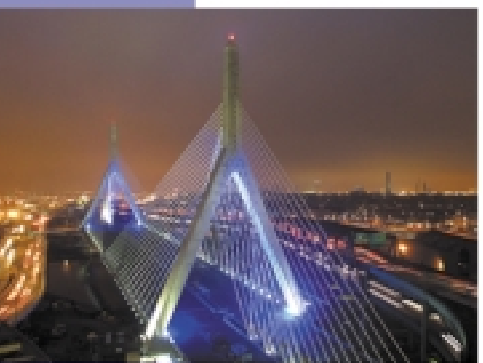

Image courtesy of HNTB.

Conference Master of Ceremonies Dr. David E. Shaw, Chairman, D. E. Shaw & Co, Inc., and Cochairman of the Council's High Performance Computing (HPC) Advisory Committee, then framed the day's discussion by outlining the questions that the more than 200 panelists and attendees had gathered to discuss:

- *What are the linkages between HPC, national security and economic competitiveness?*
- *What are the current competitive advantages of HPC to America's businesses?*
- *How could the intractable challenges industries face today be solved by increased HPC power, and how would doing this benefit their companies' competitiveness?*
- *What are the barriers impeding access to HPC by U.S. businesses?*
- *What are possible solutions to these barriers?*

In thinking about these questions, Dr. Shaw encouraged all conference participants to consider particularly "the various ways that the public sector and the private sectors can work together to more effectively harness HPC's capability for our country's competitive advantage."

The Government View: National Security, HPC and Economic Competitiveness

What's the link between government investment in HPC science and technology and the economic competitiveness of America's businesses?

The Defense Department and the Department of Energy (DOE) laboratories under the National Nuclear Security Administration and the Office of Science have long driven advances in high performance computing by assuming risk, both as investors and users, which the private sector often cannot afford to take. As investors, they have supported cutting edge, public and private sector R&D critical to advancing this technology's development and ensuring that the U.S. remains a leader in designing and manufacturing HPC systems.

As committed "power users," they purchase the most advanced systems available to gain a competitive advantage in accomplishing their missions. And as they aggressively use HPC to solve complex problems vital to national security and basic science, they push and prove out the technology, often changing it along the way. Laboratories under DOE's National Nuclear Security Administration, for example, often have been the first to purchase new HPC systems, providing not only critical early revenue, but also important user insight to HPC developers. This valuable information exchange helps hardware and software developers refine their products, and has pioneered the way for more usable and affordable systems, enabling broader adoption of this technology across the private sector to propel innovation and competitiveness.

Dr. Anthony Tether, Director of the Defense Advanced Research Projects Agency (DARPA), pointed out that HPC resources are clearly a major competitive advantage for America's national security interests. "We know in the Department of Defense and the Department of Energy the role of large computers that enable us to model very large systems to the point that we really don't have to build them until the final stage, but instead we can simulate the entire system and tell what the performance is going to be," said Dr. Tether. "That gives the U.S., from a security viewpoint, a great competitive advantage over other countries."

Dr. Tether noted that DARPA is typically six-to-eight years ahead of industry in terms of information technology use, giving the agency an over-the-horizon perspective on HPC. DARPA's history includes a crucial role in the creation of the Internet, the first personal computer, and massively parallel high performance computing.

But the full benefits of these advances, particularly the most recent ones, have not been fully embraced by the private sector. This is not only crucial for businesses, but also for broader national interests in terms of ready and economical access to HPC technologies.

"The facts are that if these high performance computers don't end up having a commercial value, the business will go away," noted Dr. Tether. He added that this larger commercial application would only happen if "HPC is easy to use." He noted also that DARPA is indirectly addressing this through its current emphasis on high productivity computing, which stresses computational productivity rather than sheer processor speed.

Speaking later in the morning, Dr. Everet H. Beckner, Deputy Administrator for Defense Programs at DOE's National Nuclear Security Administration, reiterated the singular value of HPC to national security.

He pointed out that one of the longest running applications of HPC, as well as one of the most impressive demonstrations of its capability, is its application to U.S. nuclear stockpile maintenance under the Advanced Simulation and Computing (ASC) program. As a result of the Comprehensive Test Ban Treaty, the U.S has relied on the advanced modeling and simulation capabilities of the world's most powerful high performance computers to assure the safety and reliability of the U.S. nuclear arsenal.

"We're dedicated to telling the President that the nuclear weapons work. And this rests on the backbone of HPC," said Dr. Beckner.

In his keynote luncheon address, Dr. John H. Marburger III, Science Advisor to the President and Director of the Office of Science and Technology Policy, stressed that transferring government HPC advances to American businesses requires "more than ordinary vision." He reflected on the history of HPC noting that it has always required visionaries in science and business to see and capitalize on HPC's transformative potential.

He raised a cautionary note, however, pointing to a significant gap—analogous to an enormous competitiveness opportunity—between proven HPC technologies and their private sector use. "What is happening today is that the gap is widening between the known potential of high performance computers and the capability that is actually being realized with commercial off the shelf hardware and associated software," he said.

He noted that the government's science role clearly extends to economic security. "We do feel responsible for the economic security of the nation," said Dr. Marburger. "And we believe that our economic future depends on being able to continually innovate, and we know that innovation depends on science. Therefore we think it's important to invest in science and to try and understand the whole chain of investment that leads from that investment, whether in HPC or physics, all the way to making products and selling them competitively. We want to make sure that this process works for the American people."

An HPC Vendor Perspective: The Power to Transform the Business Enterprise System

HPC is clearly critical not only to America's strategic strength but also its economic might, Dr. Paul Horn, IBM's Senior Vice President for Research, told attendees during the conference's morning keynote address. He stressed that HPC is a crucial 21st century commercial innovation tool.

"The economy in the U.S. is the biggest not because we have the most productive workers, but because we dig holes with backhoes not shovels," said Dr. Horn. He provided a myriad of examples in which HPC is central to the success of companies in sectors such as petroleum exploration and production, aerospace, life sciences, entertainment and financial services.

Dr. Horn pointed out that the conference's focus on HPC and competitiveness is a timely one that is the result of a variety of converging global forces. These include the commoditization of HPC technology, the globalization of business competition, and the recognition of the enormous potential of HPC to transform business practices, as it has scientific ones. "HPC has become the third leg of science," noted Dr. Horn. Along with experimentation and theorizing, scientists and engineers can now use high performance computing to make discoveries in simulations of real-world events. There's now in vivo, in vitro, and in silica. HPC simulations are dramatically accelerating the rate of discovery in fields from astrophysics to material science.

Dr. Horn identified three key current trends in business HPC use that reflect its growing competitive importance:

- *Computational needs of businesses continue to grow rapidly,* such as in moving to the use of full model simulations and computer-generated movies with ever more life-like images.

- *There are enormous opportunities for further productivity gains with HPC.* Companies can extract trillions of dollars in excess cost through business enterprise transformation. According to Dr. Horn, this opportunity is particularly present, and untapped, in the services sector. This sector accounts for 80 percent of the U.S. economy, but is a late-comer in the use of HPC. He stressed that HPC can be used to redesign the fundamental business operations and processes within companies. This includes using HPC to model business optimization, such as real-time technical and business integration from raw resource to finished product.

- *There is enormous potential for ongoing application of HPC to many scientific fields such as drug discovery.* "We're at the infancy of a revolution in drug discovery and this won't happen without HPC," said Dr. Horn.

HPC Drives Business Competitiveness

HPC clearly provides substantial benefits to America's scientific and national security community. But how does this translate into business economic competitiveness? In the view of Dr. James F. Decker, Principal Deputy Director, Office of Science, DOE, and the moderator of the conference's first panel, in today's intensely competitive global marketplace it's a question that warrants particular attention in thinking about the national value of HPC.

"At one time the U.S. held an unchallenged leadership position in both the development and application of supercomputers," noted Dr. Decker. "However, today other countries have aggressive governmental HPC programs that represent significant challenges to U.S. leadership in both government research programs and industry."

Dr. Decker moderated a panel of four senior HPC industry users representing the broad spectrum of aerospace, consumer products and entertainment users. What, he asked them, are the current competitive advantages of HPC to their industries? How could the intractable industrial challenges they face be aided by increased HPC power? And what impact would this have on their companies' competitiveness?

Reducing Design Costs Through Virtual Prototyping

It costs approximately one billion dollars in R&D to design and test an aircraft engine, Pratt & Whitney Senior Fellow Dr. Saadat Syed told the conference. He explained that these turbines are very complex, multicomponent systems that require extensive design and analysis testing. However, the opportunity exists to wrench costs from old-line testing through the use of high performance computing. Virtual prototyping is essential to minimizing these costs, he said.

"Physical testing is very expensive and takes a lot of time," Dr. Syed told attendees. "We use HPC to analyze these designs before we take them to physical testing. This is where HPC can have a huge impact."

This virtual prototyping has a disproportionately large, positive, competitive impact on the engine's total cost compared with other R&D investments, noted Dr. Syed. While product design traditionally accounts for about five percent of the development costs, its influence on the final product costs can be as high as 70 percent. And, he emphasized that improved computing capability will lead to better designs and reduced R&D costs.

Reducing Physical Tests for Faster Time to Market

Design & Analysis Tools Have a Large Influence on Total Product Cost

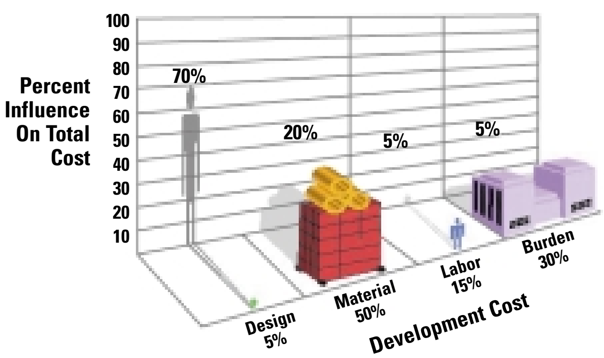

Courtesy of Pratt & Whitney.

Reinforcing Dr. Syed's comments, Doug Ball, Senior Manager, Enabling Technology and Research at The Boeing Company, provided a powerful example of how HPC simulations are drastically reducing the need for costly physical testing, and also reducing the critical issue of time to market.

The Boeing Company has reduced the number of expensive and time-consuming wind tunnel tests for the wings of its new planes by using HPC-based computational fluid dynamics simulations. In 1980, Boeing conducted 77 wind tunnel tests during the development of its 757 wings. For its 7E7 Dreamliner series currently in development, Ball estimates it will require as few as five wind tunnel tests. "If you can take five months out of the design time it's a lot cheaper," noted Ball. "And the competitive advantage of HPC comes from being the first to apply these insights in the market."

According to Ball, not only do the simulations reduce development costs and time to market, but they produce a better product. He pointed out that since air is invisible it's impossible to visualize the air fluid dynamics around an aircraft's wings in a wind tunnel. However, this can be done using simulations, thus providing more detailed and useful results.

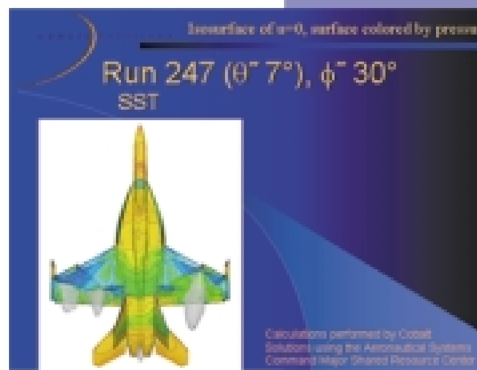

Image courtesy of The Boeing Company.

Breakthrough Insights for Manufacturers

While the aerospace industry, starting in the 1970's, was the first business sector to use HPC it has progressively moved into other sectors, including automotive, petroleum, entertainment, pharmaceutical, and notably, consumer products and packaging, observed Thomas J. Lange, Associate Director, Corporate Engineering, Head of Computer Assisted Engineering at The Procter & Gamble Company.

He pointed out that it might come as a surprise that the company that makes such well-known consumer brands as Pampers®, Tide® and Scope® spends $1.6 billion annually on R&D and has more than 24,000 active patents. And, he stressed, today HPC simulations are central to all of his company's R&D. "Explore digitally, confirm physically," is now the mantra at The Procter & Gamble Company, says Lange.

The use of HPC simulations extends from such large scale issues as modeling supply chain throughput reliability, to the virtual testing of bleach containers for their resistance to breaking when dropped, an important cost issue that involves maximizing container shape and safety while minimizing packaging costs. High performance computing also offers the promise to revolutionize manufacturing through advanced modeling and simulation of the entire process from concept through production, linking product development and manufacturing into a seamless process.

explore digitally

Lange told attendees that advanced simulations are also used to solve seemingly simple production questions. He explained how computational scientists at The Procter & Gamble Company turned to high performance computers and computational fluid dynamics equations similar to those used to analyze aircraft wings to solve the problem of the "flying Pringles®." The company's Pringles potato chips were literally flying off the conveyor belt rather than dropping into their tubular container because of irregular airflow around the uniquely designed snack food. The result was increased costs from contaminated product that had to be thrown away, not to mention a significant cleanup problem. Company scientists were able to model the air turbulence around the Pringles and adjust the manufacturing process accordingly to accommodate for this.

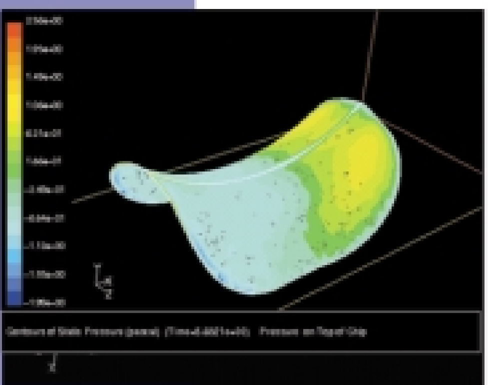

Image courtesy of
The Procter & Gamble Company.

HPC Creating New Industries

Ed Leonard, Chief Technology Officer at DreamWorks SKG, extended the application of HPC to note that the technology has created a whole new digital animation industry, which is drawing audiences around the world, and profits for the U.S. entertainment sector.

Leonard explained to attendees that the company's recent film ***Shrek 2*** would have been impossible to create without advanced HPC resources—each of the film's 130,000 frames required 80 hours of processor time, a total of a remarkable 10 million CPU hours to create a 90-minute digitally rendered film.

He explained that HPC enables filmmakers to produce computer-generated characters that behave, and more importantly for viewers, appear, like the real thing. ***Shrek 2*** wowed audiences in large part because a range of HPC hardware and software tools enabled filmmakers to control subtle aspects of lighting to characters' fur and clothing texture, creating a form of stylized realism. "People know what cloth looks like and how it should move, and if it doesn't they notice," said Leonard. "We couldn't do this without high performance computing."

Winning Globally Requires HPC-Driven Solutions

Given the level of global competition, use of high performance computing is no longer an option. It is an essential tool for developing and commercializing the next generation of high value products and services for which global consumers will pay a premium and that will capture global market share. Market leaders are and will be those that harness this capability to its full advantage.

All of the panelists emphasized that their companies need improved HPC resources to help solve a range of currently intractable, major R&D problems in order to meet growing international competition.

"My vision is to take an engine to verification testing without doing physical engine testing. Some people laugh at the idea," said Pratt & Whitney's Dr. Syed. However, he was far from alone among the panel's HPC leaders in seeing the possibility for dramatic HPC-driven competitiveness advances.

"We've pretty much reached the limits of what planes are going to look like. They're no longer just going to be tubes with wings," said The Boeing Company's Doug Ball. He noted that in order to achieve this conceptual leap, the 21st century HPC challenge is to dramatically reduce the time required for HPC simulations, and thus the testing of various iterations of new models.

It's this type of leap to the next level of HPC innovation that panelists said will provide American businesses with a significant, international competitive advantage.

Ball pointed out that the international commercial aircraft market for the next 20 years is estimated at about one trillion dollars. However, there's intense competition for this market. Dr. Syed noted that the European Union has committed billions of euros of public funding to achieve "a world-class European aeronautics industry (that) leads in global markets for aircraft, engines, and equipment by 2020." Already Pratt & Whitney's market share has declined 10 percent vis-a-vis European competitors since 1992.

Ed Leonard, of DreamWorks SKG, noted that global competition is also a key factor in the digital film industry with emerging competition from countries such as India and China. In Leonard's view, to meet this competition, "We need better development tools to make HPC computing faster, easier, and more reliable." For example, in order to develop several films simultaneously and thus have a richer product pipeline, he said DreamWorks requires better cluster-management software that provides intelligent real-time configuration and prioritization of processors.

Similarly, The Proctor & Gamble Company's Lange noted there are whole realms of product design R&D possibility still untapped because of the need for HPC simulations with greater detail and integration ability.

His vision for future virtual design includes the ability to simulate a person's biomechanical interaction with a product, such as removing the top of a plastic container. It's an advance that would greatly accelerate the design analysis process, and thus time to market. However, Lange noted, to make this a reality will require the development of complete biomechanical simulations that mimic muscle, fat, skeletal and skin behavior that can interact with simulated products.

First National HPC Users Survey: Essential Tool Needs Broader Usage

The critical link between HPC and business competitiveness identified by the conference's first panel was reinforced by the results of the first High Performance Computing Users Survey. It was released at the conference by the survey's leader, Dr. Earl Joseph, Research Vice President, High Performance Systems Program, International Data Corporation.

The survey—an initiative of the Council on Competitiveness and DARPA—was commissioned to provide a detailed, birds-eye view of HPC business use in America. Why do American companies acquire HPC resources? How are they using them? What is HPC's impact on their competitiveness? And what's the industry's outlook for future HPC use, including the potential value of high performance computers that are dramatically faster and easier to use?

The survey's respondents were veteran chief technology officers, chief information officers and production and research managers from companies representing the full range of industries using HPC today, including IT and electronics, petroleum, chemical, pharmaceutical, aerospace, automotive and entertainment. These companies' experience with HPC ranged from two to thirty years, with an average of 15 years.

The survey's core finding? HPC is an indispensable business tool to dozens of America's largest companies, but this powerful innovation driver is not used as aggressively as possible because of a variety of technical, cost and business culture hurdles.

Dr. Joseph highlighted the major findings to conference attendees:

- ***High Performance Computing Is Essential To Business Survival***

Respondents were near unanimous in indicating that HPC is so central to their work that their organizations couldn't function without it. The number-one reason given for purchasing HPCs is their unique ability to run very large and complex computational problems that these companies must successfully address to maintain competitive advantage.

- ***Companies Are Realizing A Range Of Impressive Bottom Line And Business Benefits From Using HPC***

The competitive benefits of HPC use identified included shortened product development cycles and faster time to market—in some cases more than 50 percent faster—and an associated reduction in R&D and production costs.

- ***Dramatically More Powerful And Easier To Use Computers Would Deliver Strategic Competitive Benefits And Could Add Billions To A Company's Bottom Line***

Those respondents who were able to quantify the potential benefits from access to more powerful and easier-to-use

HPC systems suggested bottom-line benefits of from tens-of-millions to billions of dollars. Given the possibility of HPC systems 100 times faster and ten times easier-to-use, respondents' projections included the comments that: "We could save $1 billion from a faster cycle time," "We could test two-generations-out models that we are researching today," and "We would look to rewrite the entire science underlying the current technology and methodology we are using."

• *Companies Are Failing To Use HPC As Aggressively As Possible*

A majority of respondents indicated that they are not using HPC as aggressively as possible. In some cases this is due to a variety of business and technical barriers. In others, the HPC resources simply don't exist today to solve the specific problems.

• *Business And Technical Barriers Are Inhibiting The Use Of Supercomputing*

"The largest single factor preventing more aggressive use of HPC tools is the lack of computational scientists either within a company or externally who are able to apply HPC tools to the company's problems," Dr. Joseph told attendees. A closely related issue is the current difficulty of use of both hardware and software that limits the broader application of HPC in many businesses.

The survey also found that despite proven returns on investment, "respondents noted that upper management often does not appreciate the value of HPC hardware, software and tools. As a result, HPC is often viewed as a cost instead of an investment, and many sites find it difficult to acquire internal funding to acquire HPC resources."

• *Companies Don't Have The HPC Tools They Want And Need*

A majority of the executives interviewed said that there are existing HPC tools that they would like to own or access but don't because of a number of financial and business culture reasons. However, a third said that they need hardware or software systems more powerful or capable than any available on the market today.

Key Conclusions: HPC by the Numbers

Here are some highlights, by the numbers:

- *Percentage of industrial and business users surveyed for whom HPC tools are indispensable:* **97**
- *Percentage of users surveyed for whom HPC tools shorten product development time and enable them to tackle more and larger problems:* **70**
- *Percentage who say they can't afford to purchase the HPC tools they'd like:* **67**
- *Percentage with major important technical and/or scientific problems that require more advanced HPC tools:* **67**
- *Percentage who said they didn't think their business was using HPC as aggressively as possible:* **58**
- *Percentage for whom improved ease-of-use would increase their company's use of HPC:* **56**

Not all respondents answered every question.
To read the full National HPC Users Survey visit www.compete.org.

Barriers to HPC Use in Industry

The barriers to the use of HPC by American businesses uncovered by the National HPC Users Survey, were reiterated and further elaborated on by the afternoon panel moderated by Dr. David B. Nelson, Director, White House National Coordination Office for Information Technology Research and Development.

The four leadership panelists represented a diverse range of fields, from industrial to engineering and the financial and banking services to academia. When introducing the panel, Dr. Nelson summarized the barriers they would be addressing into three broad categories:

- *Business culture barriers*
- *Educational/training barriers*
- *Technical barriers*

But he also encouraged all conference participants to focus on not only the barriers, but also ways to overcome them. “This conference is in a sense a problem solving and an opportunity realizing event,” observed Dr. Nelson. “Barriers can be easy ways to say ‘no.’ But we should find ways to say ‘yes’.”

Business Culture Barriers

Panelists identified two overarching business culture barriers that face those promoting greater HPC use in their companies: the question of whether HPC is viewed as an investment or cost, and the related issue of identifying return-on-investment.

• *HPC: Investment or cost?*

In the boardrooms of many American companies, HPC isn’t seen as an innovation edge, but rather a cost of doing business that’s viewed as “an enormous hole in the pocket book,” said Daniel Wolgemuth, Senior Vice President/Chief Information Officer HNTB LTD. This company is one of America’s top 20 design engineering firms and uses HPC in the design and analysis of large tunnels, bridges and stadiums.

His comments echoed the findings of the HPC Users Survey which noted that for many companies, especially those with less experience using HPC to solve their business problems, senior executives view these resources as a cost rather than an investment. According to Wolgemuth, this attitude represents a major mindset hurdle that must be addressed if we are to realize increased use of HPC to boost competitiveness.

Indeed, he said, for many companies taking advantage of HPC means rethinking the integration of HPC into their business and profit models. For example, the current project pricing models used in many consulting companies are time-based. As such, there’s reluctance on the part of some team members to rapidly accelerate the design and delivery process with HPC.

“In the consulting world, faster isn’t always [seen internally as] better,” noted Wolgemuth. “So in the professional services marketplace, if we’re going to really innovate we have to come up with new HPC-related pricing models that are value based.”

• *The return-on-investment factor*

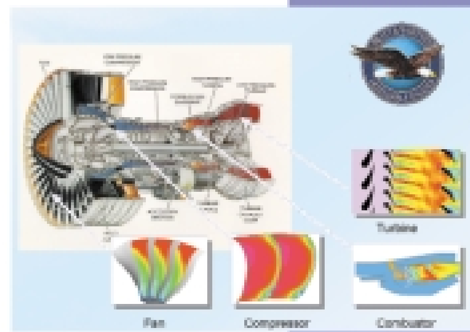

Image courtesy of Pratt & Whitney.

Part of the reason that senior executives view HPC as a cost rather than an investment is that, in most cases, they’re not given a quantifiable sense of their return on millions of dollars of investment in HPC infrastructure and staffing, said Jeffrey Birnbaum, Managing Director, Global Head of Enterprise Computing, for Morgan Stanley.

“Return on investment is one of the big issues if we’re going to make this (accelerated vision of HPC) a reality,” he noted. Morgan Stanley uses HPC extensively for market analysis and modeling for derivatives and other trading.

In Birnbaum’s view, what he terms “legacy thinking”—a psychological tie to doing things in a traditional way—can be a major hurdle to using the latest HPC resources for competitive business advantage. He noted that this is a particularly important issue given that America’s major emerging competitors—India and China—don’t have the same, potentially stunting, legacy of HPC use, or nonuse.

Missing the Right Stuff: Educational and Training Barriers

Today, many industries don’t have the people who can productively use HPC to its full innovation potential, said Dr. Stan Ahalt, Executive Director of the Ohio Supercomputer Center.

In Dr. Ahalt’s view part of the reason for this is that educators need to emphasize greater vision and imagination when it comes to applying HPC to real-world problems—they need to encourage students to think big. “We don’t teach our students to think about what problems they could

investment vs. cost

solve if they had an infinite amount of computational time and an infinite amount of storage," said Dr. Ahalt, noting that these students will have access to much larger HPC systems than available today.

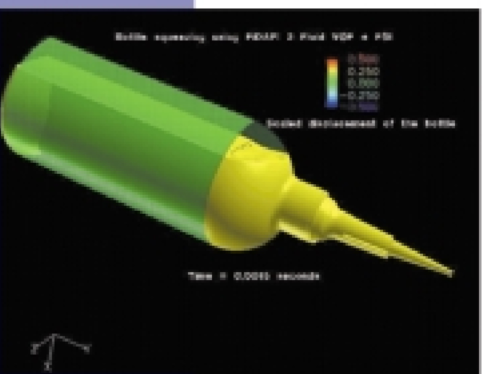

Image courtesy of The Procter & Gamble Company.

The panelists identified a two-pronged shortfall in human resources for HPC use in businesses, (though there was a range of opinion over the extent of the problem):

- *Not enough people in the HPC educational/training pipeline;*
- *A poor match between the education provided by universities and the particular HPC skills and multidisciplinary perspective required by industry.*

"Today's engineering undergraduates are rarely taught passé programming languages and there's no introduction to distributed parallel programming. As a result, graduate student programming skills are poor," said Dr. John O. Hallquist, president of Livermore Software Technology Corporation, a leader in the development of the simulation software used in automobile crash modeling.

According to Dr. Ahalt, the result of this lack of parallel programming skills in industry "is that engineers settle on sub-optimal designs."

It's a vicious circle in which an unclear sense of the value of human resources reduces the real value of HPC ownership and use, and thus exacerbates the perception of HPC hardware and software as a cost rather than an investment, noted Morgan Stanley's Jeffrey Birnbaum.

HPC Technical Barriers

The panelists identified key technical barriers that fall into three broad categories:

• *Need for improved mathematical models and basic science*

Dr. Hallquist noted that mathematical models are the foundation for all high performance computing since they convert scientific insights about physical processes into a form that can be used for the development of computer codes and software.

He identified a range of modeling issues which are currently barriers to more advanced simulations using his company's finite element software. These include load balancing in parallel computers (the ability to share the workload efficiently among hundreds, or thousands, of processors), and the level of accuracy of highly complex, multicomponent physics models such as automobile crash simulations.

"We need to be able to run a model twice and get exactly the same answer. The ultimate goal is to reduce the number of physical prototypes to zero," he said.

• *Legacy software inhibiting usage*

American businesses are relying on a diverse range of proprietary in-house and commercially available outsourced software to meet their HPC needs. Legacy codes, especially third party software legacy codes, are seen by some as a significant obstacle to more widespread use. Many important HPC application codes have been in use for decades. But whichever route they're taking, a major barrier for all is the applicability, compatibility and ease-of-use of their HPC software.

"It's HPC software that needs dramatic improvement," said Dr. Ahalt. "So if we can solve the basic problem of better applied parallel software, the entire spectrum of HPC applications from basic science to industry will reap the benefits."

He added that the HPC software crunch extends beyond the technical applications to issues of HPC load balancing, optimization, cluster management and advanced visualization.

- ***The right fit: Codes must be scaled for broader industrial use***

The emphasis in the government HPC community on the development of codes for parallel application on thousands of processors has created a barrier for many potential business users.

"To date the concentration in HPC has been on high-level, sophisticated programs rather than simpler, more broadly applicable applications," said Dr. Ahalt. He notes that while the codes developed by DoD, NSF and DOE often involve thousands of processors, most industrial applications involve less than 64 processors. "Right scaling" the codes, he suggested, is thus a significant and sometimes insurmountable obstacle for industrial HPC users.

What Drives Business HPC Decisions?

During the course of the conference, panelists and speakers identified four key business issues, which, although not "barriers," form the framework within which American businesses are making decisions about HPC purchase and use.

Cost is Core

Cost issues are the driving force in the business use of HPC. "I can't tell you the number of times people feel constrained because of the costs of HPC," noted Jeffrey Birnbaum of Morgan Stanley.

Questions of cost apply not only to the hardware, but also to the software, and to facilities issues such as HPC housing and cooling, noted The Boeing Company's Doug Ball. Indeed, he said, for some users the cost of simulation software dwarfs the cost of the hardware. In this sense, cost is a critical issue in limiting the transition from physical testing to the use of advanced simulation.

According to IBM's Dr. Horn cost is an issue that's recognized by HPC suppliers. "We need to build the technology at a low enough cost to allow our customers to use it."

Productivity, not Speed

The traditional HPC metric for measuring performance is raw processor speed. Supercomputers on the Top 500 Supercomputer List are ranked based on flops, the number of floating point operations per second. "We've got to find a better way of defining performance other than flops," said one conference attendee voicing a commonly held sentiment. For business users the key performance metrics are HPC price, suitability to their specific applications, and the speed at which the machine can run their specific applications. Indeed, in many cases, application specificity and experience with a particular hardware architecture are more important than cost or potential speed using a new hardware platform.

HPC is a Tool

"The key thing for our engineers is using the HPC as a tool—they don't care about computational fluid dynamics," said The Boeing Company's Ball. Easier to use systems not only benefit experienced HPC users, but also invite a larger pool of individuals to use this tool, ultimately expanding potential applications, noted Dr. Ahalt.

Choices, not Answers

HPC users emphasized that in a business setting HPC doesn't provide singular solutions, but rather additional critical information on which to base business decisions.

"There's always someone who knows the answer. The challenge is to get the other 10,000 people to agree. The benefit of modeling is in helping to get people to agree. It aids executives' intuitive decision making," said Procter & Gamble's Thomas Lange.

Next Steps: Supercharging U.S. Innovation with HPC

"What is our vision for the future? What are the next steps?" asked Dr. Holbrook, President, The Ohio State University and CoChairman, Council on Competitiveness HPC Advisory Committee, in leading the conference's closing discussion on the next steps.

In providing a context for this discussion, she noted that in her view "What we've heard is that there's very clear agreement on the transformational capabilities of HPC and our ability to innovate to give our businesses, our scientists and engineers, and our nation a competitive advantage to enhance productivity and national security."

Participant suggestions for overcoming hurdles to HPC use and capitalizing on its potential to fuel U.S. economic competitiveness fell into five broad categories:

Strengthening Government, Industry, University Partnerships

Panelists, speakers and attendees emphasized that the main priority is the need for renewed and reinvigorated partnerships between businesses, government, and academia. What shape should such partnerships take in the 21st century?

As the DOE's Dr. Decker stated, "I think one of the important issues in advancing HPC in this country is how do we improve the connections between government and industry to help facilitate development in both sectors." These partnerships include the need to reinforce, re-establish, or create links and government-industry partnerships between software and hardware development to avoid the "not-invented-here syndrome" in many businesses, said The Boeing Company's Doug Ball.

In terms of interagency linkages within the federal government, Dr. Marburger, Science Advisor to the President and Director of the Office of Science and Technology Policy noted that High End Computing (HEC) is one of six major areas within the Networking and Information Technology R&D (NITRD) group, one of the President's major interagency coordination efforts funded at more than $2 billion annually. He noted that NITRD has formed a High End Computing Revitalization Task Force which produced a report, released in May 2004, outlining among other things an interagency roadmap for high-end computing core technologies and an accessibility improvement plan.

Improving HPC Education and Training

While views differed as to the extent of the problem, there was a widely shared perspective that there's need for improved linkage between the HPC skills and knowledge taught by universities and those required by private sector companies. At the same time, companies must maintain adequate internal human resources to maintain HPC innovation. "There's no way around spending money on people," said Dr. Beckner of DOE's National Nuclear Security Administration.

Conference participants also noted the need to emphasize multidisciplinary training and an increased understanding of parallel programming methods. This multidisciplinary training includes not only the ability to understand the scientific and technical aspects of a problem, but also how to work collaboratively with a diverse range of partners for whom HPC is a "black box" that purely provides a service on the way to a product. "Even with our most technical folks there needs to be the ability to communicate with the artistic folks," said Ed Leonard of DreamWorks SKG.

Fueling Next Generation HPC Simulations

The ongoing and targeted role of university and government agency research in developing ever more advanced mathematical models and codes is crucial to making the leap to a next level of innovative HPC use by American businesses. "It's not the lack of sheer crunching power, it's the lack of ability. We need the math and models or all the computer cycles won't help," said one attendee during the general discussion on solutions.

Participants felt it's critical that government and academia continue to play a leadership role in advancing HPC applications, demonstrating proof-of-concept, and efficiently sharing these advances with industry. "It's very important for government agencies to demonstrate that whole system simulations can be done," noted Dr. Syed of Pratt & Whitney.

This new science and engineering insight extends to the entirety of the HPC envelope, including faster, more reliable networking, scalable parallel algorithms and the need for ever greater fidelity of models. The issue of model validation is crucial to the regulatory approval of HPC models as an alternative to physical testing.

Imaging New Business and Value Models

Businesses must change the ways they think about HPC in their value models. And this change in thinking must be championed internally by HPC users and clearly communicated through "upreach" to the boardroom level. "As HPC leaders, it's our job to translate the HPC dream into something that works in the boardrooms across the United States. Can it be done? Absolutely, " asserted HNTB LTD's Daniel Wolgemuth.

Next Generation Access Scenarios for Business HPC

Image courtesy of DreamWorks SKG.

There were three proposed scenarios for increasing HPC use, each briefly outlined here:

• *Computing on demand*

Also described as "utility computing," this scenario envisions a new market model. Rather than maintaining costly, permanent internal resources, businesses such as digital animation companies would be able to access HPC resources on an "as needed" basis, primarily to meet peak demand. However, some conference participants thought this model could involve seemingly prohibitive security concerns for many companies.

• *Competitive computing*

As with computing on demand, this scenario also sees a broadening of competitive HPC use by lowering the cost of computing through maximizing the use of internal computing resources. This involves operating HPC resources in what Morgan Stanley's Jeffrey Birnbaum termed a "stateless compute environment," in which HPC assets act as a large dynamic pool available to a broad range of users, including the potential to sell cycles to recoup investment. This relies on the development of easy-to-use software that can be broadly applied.

• *Blue collar computing*

The most ambitious and wide-ranging vision of change in HPC use to fuel innovation was presented by Dr. Stan Ahalt, of the Ohio Supercomputing Center. He proposed an HPC community R&D emphasis on improved programming models that exist in the zone between dominant current use (up to about 64 CPUs) and the "heroic" scale problems tackled by national agencies such as DOE and DoD. This zone of 'blue collar computing,' he said, offers a logical step to sparking "the next long lived productivity expansion" and requires that we "focus our innovations, advances, and education on the entire application spectrum." He added that while "the benefits will be primarily economic, the heroes…will also reap the benefits."

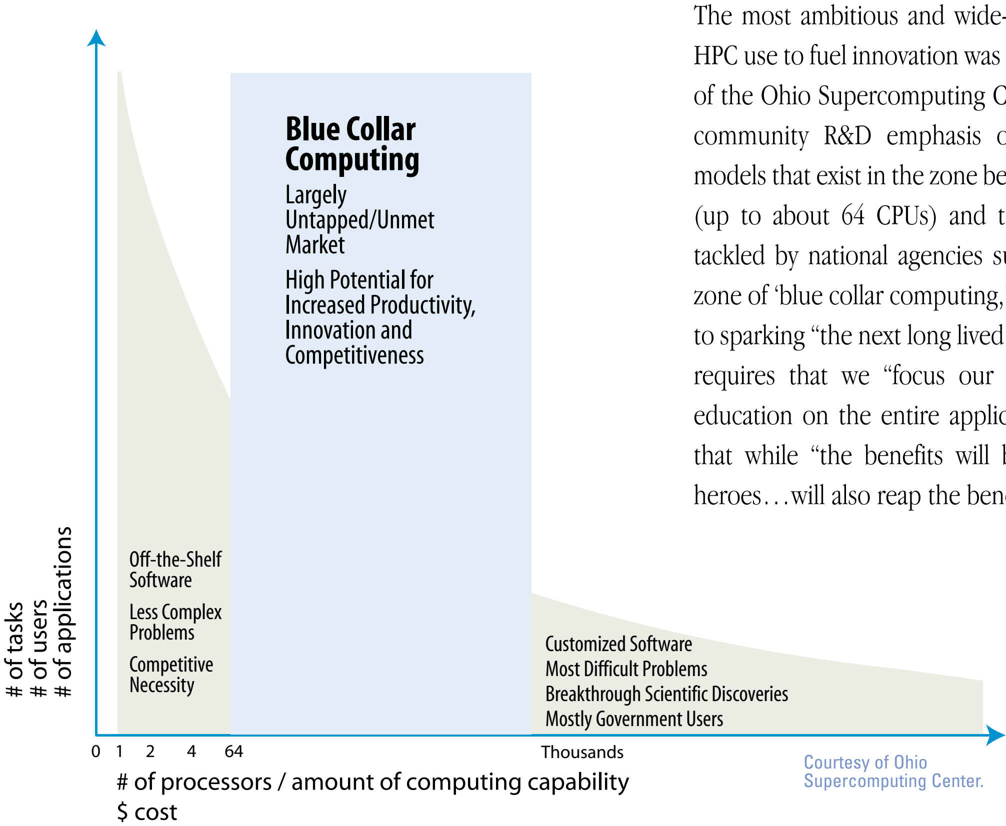

Courtesy of Ohio Supercomputing Center.

Looking to the Future

Based on the results of the 1st Annual High Performance Computing Users Conference, the Council on Competitiveness' HPC Advisory Committee has the clear direction to identify opportunities for public-private partnerships, said cochair Dr. David Shaw. This is now at the heart of the Advisory Committee's work in creating an Action Agenda. And it will form a key part of discussions at the 2nd High Performance Computing Users Conference, planned for July 13, 2005.

In closing, Advisory Committee cochair Dr. Karen Holbrook reiterated a statement by Dr. Marburger made during his keynote address that she believed encapsulated the day's events: "We are approaching a 'tipping point' beyond which entirely new applications of computing will bring a new wave of transformations in our industrial ways of life, and further disrupt older ways of doing business," said Dr. Marburger. "We need to spread the word about the new capabilities and build confidence in the new visions to motivate private and public investment in them…. In today's globally competitive economy we cannot afford to leave this to others."

Appendix

Conference Agenda

Executive Summary of HPC Users Survey

Council on Competitiveness
High Performance Computing
Advisory Committee

Council on Competitiveness
Members, Affiliates, and Staff

FIRST ANNUAL

High Performance Computing Users Conference:

SUPERCHARGING U.S. INNOVATION & COMPETITIVENESS

July 13, 2004

AGENDA

7:30 a.m. Breakfast

8:15 a.m. **Welcome**

Deborah L. Wince-Smith, President, Council on Competitiveness

Dr. David E. Shaw, Chairman, D.E. Shaw & Co., Inc. and CoChair, Council on Competitiveness HPC Advisory Committee

Dr. Anthony Tether, Director, Defense Advanced Research Projects Agency

Morning Keynote Address: Competing through Computing Power: Leveraging HPC for Global Economic Leadership

Dr. Paul Horn, Senior Vice President for Research, IBM

Note that time is allocated within following sessions for questions and answers

9:15 a.m. **Panel 1: HPC: Key to Solving Industry's Intractable Problems**

Industry panelists will discuss cutting-edge business challenges that require more advanced computing capabilities, and the impact on competitiveness if these challenges can be successfully addressed.

Moderator: Dr. James F. Decker, Principal Deputy Director, Office of Science, Department of Energy

Doug Ball, Manager of Enabling Technology and Research, The Boeing Company

Thomas J. Lange, Associate Director, Corporate Engineering, Head of CAE, The Procter & Gamble Company

Ed Leonard, Chief Technology Officer, DreamWorks SKG

Saadat Syed, Senior Fellow, Pratt & Whitney

10:45 a.m. **Break**

11:00 a.m. **HPC User Trends: Results of the First National High Performance Computing Users Survey**

Survey results will be released, revealing trends in why companies acquire HPC, how they are using HPC, the impact on their competitiveness, and industry's outlook for future use.

Dr. Earl Joseph, Research Vice President, High Performance Systems Program, International Data Corporation

12:00 p.m. Dr. Everet H. Beckner, Deputy Administrator for Defense Programs, National Nuclear Security Administration, U. S. Department of Energy

12:15 p.m. **Luncheon**

Luncheon Address: Global Challenges: HPC Solutions
Why leadership in HPC is important to achieving both U.S. national security and economic competitiveness goals.

John H. Marburger III, Science Advisor to the President and Director of the Office of Science and Technology Policy

2:00 p.m. **Panel 2: Barriers to Industry HPC Usage**
Industry and university panelists will discuss barriers that companies face in acquiring and applying HPC to solve industry challenges and ways they are successfully addressing them, if in fact they are.

Moderator: Dr. David B. Nelson, Director, White House National Coordination Office (NCO) for Information Technology Research and Development (IT R&D)

Dr. Stan Ahalt, Executive Director, Ohio Supercomputer Center

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

Dr. John O. Hallquist, Founder and President, Livermore Software Technology Corporation

Daniel Wolgemuth, Senior Vice President/Chief Information Officer, HNTB LTD.

3:30 p.m. **Next Steps/Comments from Attendees**
Dr. Karen A. Holbrook, President, The Ohio State University, CoChairman, Council on Competitiveness HPC Advisory Committee

4:00 p.m. **Adjourn**

Executive Summary

Commissioned by the Council on Competitiveness from IDC,
and sponsored by the Defense Advanced Research Projects Agency

Executive Summary

This study, commissioned by the Council on Competitiveness (COC) from IDC, and sponsored by the Defense Advanced Research Projects Agency, explores the usage and impact of high-performance computing (HPC) resources in industry and other business sectors — including currently available HPC computers and potential future computers assumed to be dramatically faster and easier to use. The study asked about both capacity-class computers, purchased primarily to address many small and medium-sized problems, and capability-class computers, purchased mainly to tackle the largest, most daunting individual problems. The 33 participants in this study are seasoned private-sector chief technology officers (CTOs), chief information officers (CIOs), and production and research managers representing a wide range of business segments that employ HPC today — from leading aerospace, automotive, petroleum, electronics, pharmaceutical, life sciences, and software companies to financial services, transportation logistics, and entertainment firms. Complete survey results are available on the Council on Competitiveness web site *www.compete.org*.

Major Findings

Nearly 100% of the respondents indicated HPC tools are indispensable

High-Performance Computing Is Essential to Business Survival

High-performance computing is not only a key tool to increasing competitiveness, it is also a tool that is essential to business survival. Nearly 100% of the respondents indicated that HPC tools are indispensable, stating that they would not exist as a viable business without them or that they simply could not compete effectively. A majority (70%) of the respondents indicated that HPC is so important that their organizations could not function without it.

Risks from Not Having Access to HPC Computers

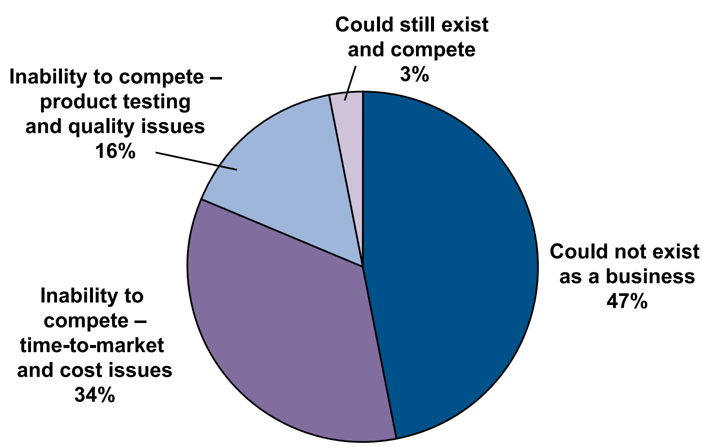

Figure 6 *Source: IDC, 2004*

"There is no other way for us to complete our work. We would not exist."

Typical comments include:

"There is no other way for us to complete our work. We would not exist."

"The time to market would prohibit our business from existing."

"We would not be able to stay technologically ahead of other competing nations."

The number 1 reason given for purchasing high-end computers is their unique ability to run very large and very complex computational problems that companies must successfully address to maintain their competitive advantage. In addition to running these large-scale problems, the majority of respondents are also able to harness the computer power to run a larger number of smaller-scale, important problems than they were able to run in the past.

Companies Are Realizing a Range of Financial and Business Benefits from Using HPC

"It drives innovation, R&D effectiveness, and productivity."

Companies described a range of impressive competitiveness benefits realized from using high-performance computing. Approximately one-quarter of the respondents were able to quantify the ROI to their organizations, in some cases in the millions of dollars. Strategic competitive benefits included gains such as shortened product development cycles and faster time to market (in some cases more than 50% faster), not to mention the resultant reduced costs, all of which can improve a company's bottom line.

"It has been a continuous stream of revenue to our bottom line, giving us the ability to look into other development areas."

"It drives innovation, R&D effectiveness, and productivity."

Companies Are Failing to Use HPC as Aggressively as Possible

Despite the acknowledged importance of high-performance computing to business competitiveness, a majority of respondents acknowledged that they are not using HPC as aggressively as possible. Two-thirds of the respondents indicated that they have important problems that they simply can't solve today. The remaining third said that they need more powerful systems to achieve more effective solutions. Reasons for both vary. In some cases, systems with the needed capability are on the market but companies face obstacles in owning or accessing them or in using them to their fullest capability. These barriers are discussed below. In other cases, the systems required simply don't exist.

Do Organizations Have Important Computational Problems They Can't Solve Today?

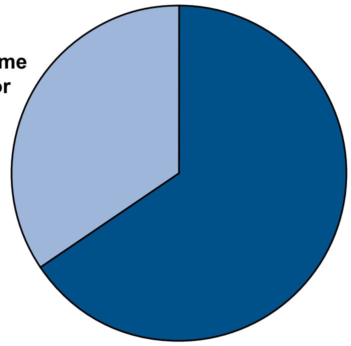

FIGURE 11 *Source: IDC, 2004*

Examples of current unsolved problems include modeling block engine assembly in full detail, simulating vehicle rollover, real-time processing of data from remote sensors, protein folding, and coordinating databases across tens of thousands of servers.

Business and Technical Barriers Are Inhibiting the Use of Supercomputing

The largest single factor preventing more aggressive use of HPC is the lack of computational scientists.

Top Factors Holding Back Organizations from Using HPC Tools More Aggressively

- *Availability of internal or external people to apply the tools to our problems*
- *Ease of use (hardware and software)*
- *Easier to get decision on investment that reduces costs now versus future*
- *Cost of HPC tools (hardware, software) versus other business investments required*
- *Decision makers do not grasp HPC impact versus other business pressures*
- *Scalability of commercial ISV software*

TABLE 8 *Source: IDC, 2004*

Respondents noted a range of reasons that HPC is not used more aggressively. The largest single factor is the lack of computational scientists — human experts (internal or external) who can apply HPC tools to the problems in question — and the budget to hire them. In most cases, the concern was the lack of resources to hire people, but in a few cases, it was the lack of available talent in the marketplace. Closely related is the ease-of-use issue; most industrial sites require software compatibility in their HPC servers and the cost to change or rewrite software is frequently seen as prohibitive.

Despite the often proven returns from using high-performance computing, respondents noted that upper management often does not appreciate the value of HPC hardware and software tools. As a result, HPC is often viewed as a cost instead of an investment, and many sites find it difficult to obtain internal funding to acquire additional HPC resources. More than half of the respondents expect their budgets for all HPC tools will decline (43%) or remain the same (17%) over the next two years.

Companies Don't Have the HPC Tools They Want and Need

When asked if there are currently available HPC tools they would like to own or access, a majority of the respondents answered in the affirmative. Relatively even numbers of respondents pointed to currently available software and hardware tools they would like to own or access. However, 31.6% stated that there are either hardware or software tools missing in the market today, and 21% said that they need hardware systems that are more powerful than any available on the market today.

Most Companies Do Not Rely on Remote Access to HPC

When respondents were questioned about their methods of accessing HPC resources, most responded that they use on-site purchased or leased HPC systems instead of accessing them remotely at partner or external provider sites. And most do not expect to outsource their most complex (and therefore most competitively sensitive) problems in the future. Security is an important inhibiting factor for some companies.

Dramatically More Powerful and Easier-to-Use-Computers Would Deliver Strategic, Competitive Benefits

"It would make these tools available to a much wider array of scientists who have good ideas but may not have programming skills."

When respondents were asked what they could accomplish with systems 100 times more powerful and/or 10 times easier to use, their replies again reflected the strategic importance of HPC to competitiveness. They saw opportunities to simulate larger, more accurate models and tackle completely new problems that they cannot address today, resulting in the ability to produce higher quality products, achieve faster time to market, and improve their financial performance.

When asked what could be accomplished if the "ease-of-use" barrier were addressed with systems that are 10 times easier to program, respondents overwhelmingly indicated that they could develop more powerful applications and fundamentally rewrite their current codes.

Not surprisingly, they also indicated that they could shorten design cycles and time to market, a natural by-product of better applications. In addition, more easily programmable systems would enable a wider universe of researchers, scientists, inventors, designers, manufacturers, and mathematicians to use high-performance computing to solve their problems, extending the benefits of these systems more broadly across the private sector for increased industrial and national competitiveness.

"We could test two-generations-out models that we are researching today."

"It would increase revenues for the company and market share."

"We would look to rewrite the entire science underlying the current technology and methodology we are using."

"It would make these tools available to a much wider array of scientists who have good ideas but may not have programming skills."

"We save $1 billion from a faster product cycle."

Dramatically More Powerful and Easier-to-Use-Computers Could Add Billions to the Bottom Line

Although not all respondents were able to quantify the potential benefits from access to more powerful and easier-to-use systems, those who could suggested bottom-line improvements from tens of millions to billions of dollars, an enormous increase over the positive financial benefits users are already achieving today.

"We save $1 billion from a faster product cycle."

"I can't release [the amount], but it is in the billions a year."

High Performance Computing Advisory Committee

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

University CoChair
Karen A. Holbrook
President, The Ohio State University

Private Sector

Scott Kahn
Chief Science Officer
Accelrys

John F. Sparks
Director
Engineering & Technology Programs
Aerojet

J. D. Thompson
Technical Services
ATK Thiokol

William Madia
Executive Vice President
Laboratory Operations
Battelle Memorial Institute

Hank Kafka
Vice President
Architecture and Emerging Technology
BellSouth Corporation

Frederick H. Hausheer, M.D.
Chairman and Chief Executive Officer
BioNumerik Pharmaceuticals

Suresh N. Shukla
HPC Service Manager
The Boeing Company

Don Paul
Chief Technology Officer
ChevronTexaco

Thomas Coleman
Director, Cornell Theory Center
Cornell University

Rich Grzybowski
Director, Systems Engineering
Corning, Inc.

Peter Ungaro
Vice 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 SKG

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

Paul Bemis
Vice President Marketing
Fluent Inc.

Vincent Scarafino
Manager
Numerically Intensive Computing
Ford Motor Company

Mark Bennett
IR&D Program Manger
General Dynamics Electric Boat

Sharan Kalwani
CAE HPC Manager
General Motors Corp.

Steve Squires
Chief Science Officer
Hewlett-Packard Company

Daniel Wolgemuth
Senior Vice President/Chief Information Officer
HNTB Technology
HNTB LTD.

David Turek
Vice President, Deep Computing
IBM Corporation

Earl Joseph
Research Vice President
IDC

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

Ben Bennett
Director, HPC Programs Office
Intel Corporation

Rene Copeland
Director, Industrial Sector Sales
Linux Networx

Eric J. Pitcher
Vice President, Product Marketing
Linux Networx

Grant Cook, Jr.
Senior Scientist
Livermore Software
Technology Corporation

Pradeep Raj
Sr. Manager
Advanced Development Programs
Lockheed Martin Aeronautics Co.

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

Irene M. Qualters
Vice President
Research Information Services
Merck & Co., Inc.

Jeffrey Birnbaum
Managing Director
Global Head of Enterprise Computing
Morgan Stanley

Timothy Lyons
Executive Director
Office of the CTO
Morgan Stanley

Reza Sadeghi
Vice President
Solver Development
MSC.Software

Stan Ahalt
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

Tom Lange
Chief Technologist for Reliability
Engineering
The Procter & Gamble Company

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

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

Daniel Reed
Director
Institute for Renaissance Computing
University of North Carolina

Public Sector

Robert Graybill
Program Manager
Defense Advanced Research
Projects Agency (DARPA)

Dimitri Kusnezov
Director
Advanced Simulation and Computing
Department of Energy

Daniel Hitchcock
Senior Technical Advisor
Office of Advanced Scientific
Computing
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 (CCN)
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
Chief Scientist & Director
Bioinformatics & Computational
Biology
Pacific Northwest National
Laboratory

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

David Nelson
Director
White House National Coordination
Office for Information Technology,
Research and Development

Council on Competitiveness Membership

Council on Competitiveness Chairs

Chairman
F. Duane Ackerman
BellSouth Corporation

University Vice Chairman
G. Wayne Clough
Georgia Institute of Technology

Immediate Past Chairman
Raymond V. Gilmartin
Merck & Co., Inc.

Executive Committee

David Baltimore
California Institute of Technology

Alfred R. Berkeley III
Community of Science

Robert R. Bishop
Silicon Graphics, Inc.

Molly Corbett Broad
University of North Carolina

William R. Brody
Johns Hopkins University

Vance D. Coffman
Lockheed Martin Corporation

Jared L. Cohon
Carnegie Mellon University

John J. DeGioia
Georgetown University

Gary T. DiCamillo
TAC Worldwide Companies

Robert C. Dynes
University of California

John Edwardson
CDW Corporation

William R. Hambrecht
W. R. Hambrecht & Co., LLC

Sheryl Handler
Ab Initio

John L. Hennessy
Stanford University

Charles O. Holliday, Jr.
E.I. DuPont de Nemours & Company

Shirley Ann Jackson
Rensselaer Polytechnic Institute

Henry A. McKinnell, Jr.
Pfizer Inc

John P. Morgridge
Cisco Systems, Inc.

Samuel J. Palmisano
IBM Corporation

Vikram S. Pandit
Morgan Stanley

Michael E. Porter
Harvard University

Luis Proenza
University of Akron

Patricia F. Russo
Lucent Technologies Inc.

W.J. Sanders III
Advanced Micro Devices, Inc.

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

Ray Stata
Analog Devices, Inc.

Lawrence H. Summers
Harvard University

Lawrence Weber
W2Group

General Membership

Michael F. Adams
The University of Georgia

Michael T. Aiken
University of Illinois, Urbana Champaign

Robert A. Altenkirch
New Jersey Institute of Technology

Richard Anthes
University Corporation for Atmospheric Research

Gerard J. Arpey
AMR Corporation and American Airlines, Inc.

Morton Bahr
Communications Workers of America, AFL-CIO, CLC

William F. Ballhaus, Jr.
Aerospace Corporation

Steven Ballmer
Microsoft Corporation

Craig Barrett
Intel Corporation

Richard A. Bendis
Innovation Philadelphia

Robert Berdahl
University of California, Berkeley

James Blanchard
Synovus Financial Corporation

Lee C. Bollinger
Columbia University

Richard L. Byyny
University of Colorado, Boulder

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

Daniel A. Carp
Eastman Kodak Company

Carol A. Cartwright
Kent State University

John T. Casteen, III
University of Virginia

George Chamillard
Teradyne, Inc.

Marsha Chandler
University of California, San Diego

Ralph J. Cicerone
University of California, Irvine

Mary Sue Coleman
University of Michigan

Lewis Collens
Illinois Institute of Technology

France A. Cordova
University of California, Riverside

Denis A. Cortese, M.D.
Mayo Clinic

Margaret B. Cozzens
Colorado Institute of Technology

Michael M. Crow
Arizona State University

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

Susan Davis
Susan Davis International

Lurita Alexis Doan
New Technology Management

Peter R. Dolan
Bristol-Myers Squibb Company

Rodger B. Dowdell, Jr.
American Power Conversion Corp.

Michael T. Duke
Wal-Mart Stores, Inc.

David B. Eisenhaure
SatCon Technology Corporation

Richard J. Elkus, Jr.
Voyan Technology

Larry R. Faulkner
University of Texas at Austin

Carleton S. Fiorina
Hewlett-Packard Company

Marye Anne Fox
North Carolina State University

Robert M. Gates
Texas A & M University

E. Gordon Gee
Vanderbilt University

James W. Griffith
The Timken Company

Ajit Gupta
Speedera Networks

Peter Halpin
World Resources Company

David C. Hardesty, Jr.
West Virginia University

Jack Harding
ESilicon Corporation

Michael R. Haverty
Kansas City Southern

Victoria Franchetti Haynes
RTI International

Robert Hemenway
University of Kansas

Adam W. Herbert
Indiana University

Karen A. Holbrook
The Ohio State University

Irwin M. Jacobs
QUALCOMM, Inc.

Jerry Jasinowski
National Association of Manufacturers

Martin C. Jischke
Purdue University

Nannerl Keohane
Duke University

Carl F. Kohrt
Battelle Memorial Institute

Raymond R. Kwong
SCRAM Technologies, Inc.

Diana MacArthur
Dynamac International

Edward A. Malloy
University of Notre Dame

Thomas J. Malone
Milliken & Company

Martin G. McGuinn
Mellon Financial Corp.

Phillip Merrick
WebMethods, Inc.

Paul Meyer
Voxiva, Inc.

Charlene Miller
Global Directors, LLC

Peggy Gordon Miller
South Dakota State University

Mark P. Mills
Digital Power Capital, LLC

Clayton Daniel Mote, Jr.
University of Maryland

Anne M. Mulcahy
Xerox Corporation

Diana S. Natalicio
University of Texas, El Paso

Mark A. Nordenberg
University of Pittsburgh

Constantine Papadakis
Drexel University

Peter G. Peterson
Blackstone Group

Harold J. Raveché
Stevens Institute of Technology

Steven S Reinemund
Pepsi-Cola Company

Judith Rodin
University of Pennsylvania

Allen B. Rosenstein
Pioneer Magnetics

David Roselle
University of Delaware

Duane J. Roth
Alliance Pharmaceutical Corp.

James E. Rottsolk
Cray Inc.

Michael Ruettgers
EMC Corporation

Timothy P. Ryan
University of New Orleans

Kenan Sahin
TIAX, LLC

Henry Samueli
Broadcom Corporation

Pamela Sedmak
Ernst Sedmak, LLC

Kevin Sharer
Amgen

Paul Shirley
Qynergy Corporation

Sanford Shugart
Valencia Community College

Lou Anna K. Simon
Michigan State University

Albert J. Simone
Rochester Institute of Technology

John B. Simpson
State University of New York at Buffalo

David Skorton
University of Iowa

L. Dennis Smith
University of Nebraska

Graham B. Spanier
The Pennsylvania State University

Susan S. Stautberg
Partner Com Corporation

Charles W. Steger
Virginia Polytechnic Institute and State University

Andrew L. Stern
Service Employees International Union, AFL-CIO, CLC

Harry C. Stonecipher
The Boeing Company

Matthew J. Szulik
Red Hat, Inc.

Sidney Taurel
Eli Lilly and Company

Henri A. Termeer
Genzyme Corporation

Lydia W. Thomas
Mitretek Systems Inc.

Steve VanAusdle
Walla Walla Community College

Larry N. Vanderhoef
University of California, Davis

G. Richard Wagoner, Jr.
General Motors

Edie Weiner
Weiner, Edrich, Brown, Inc

William C. Weldon
Johnson & Johnson

Gerald J. White
BV Solutions Group, Inc.

John D. Wiley
University of Wisconsin Madison

James Woodward
University of North Carolina Charlotte

James Wright
Dartmouth College

Mark Wrighton
Washington University

Henry T. Yang
University of California Santa Barbara

Paul A. Yarossi
HNTB Corporation

Theodore Zampetis
Shiloh Industries

National Affiliates

American Association for the Advancement of Science

American Association of Engineering Societies

American Council on Education

American Electronics Association

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

Center for National Software Studies

Community Learning & Information Network

Computer Systems Policy Project

The Conference Board, Inc.

Council for Chemical Research

Council on Governmental Relations

Infotonics Technology Center Inc.

Iowa Business Council

IEEE-USA

National Association of Management and Technical Assistance Centers

National Center for Manufacturing Sciences

National Technology Transfer Center

Council Staff

Deborah L. Wince-Smith
President

David Attis
Director
Policy Studies

William Bates
Vice President
Congressional Outreach

Erich Bloch
Distinguished Fellow

Jennifer Sue Bond
Vice President
International Affairs

C. William Booher, Jr.
Senior Advisor
and Chief Operating Officer

Claudette Davis
Office Manager

Marcy Eisenberg
Research Associate, Policy Studies

John Engler
Distinguished Fellow

Chad Evans
Vice President
National Innovation Initiative

Daniel Goldin
Distinguished Fellow

Christopher Hayter
Director
Policy Studies

Mohamed Khan
Vice President
Information Systems

Randall T. Kempner
Vice President
Regional Innovation

Melissa Meurer
Executive Assistant

Suzy Tichenor
Vice President and Director
High Performance Computing
Initative

Debra van Opstal
Senior Vice President
Policy and Programs

Amanda Welch
Manager
Policy Studies

Council on Competitiveness

The Council on Competitiveness is a nonpartisan, nonprofit organization whose members are corporate chief executives, university presidents and labor leaders dedicated to setting an action agenda to drive U.S. economic competitiveness and leadership in global markets. The Council helps shape the national debate on competitiveness by concentrating on a few critical issues including national and regional innovation, competitiveness and security globalization, workforce development and the benchmarking of U.S. economic performance against other countries.

Council on
Competitiveness

Council on Competitiveness
1500 K Street NW, Suite 850
Washington, D.C. 20005
(202) 682-4292
Fax: (202) 682-5150
www.compete.org

Second Annual High Performance Computing Users Conference Report

Accelerating Innovation for Prosperity
July 13, 2005
Washington, D.C.

Council on Competitiveness

SPONSORS

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

Defense Advanced Research Projects Agency

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

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

National Science Foundation

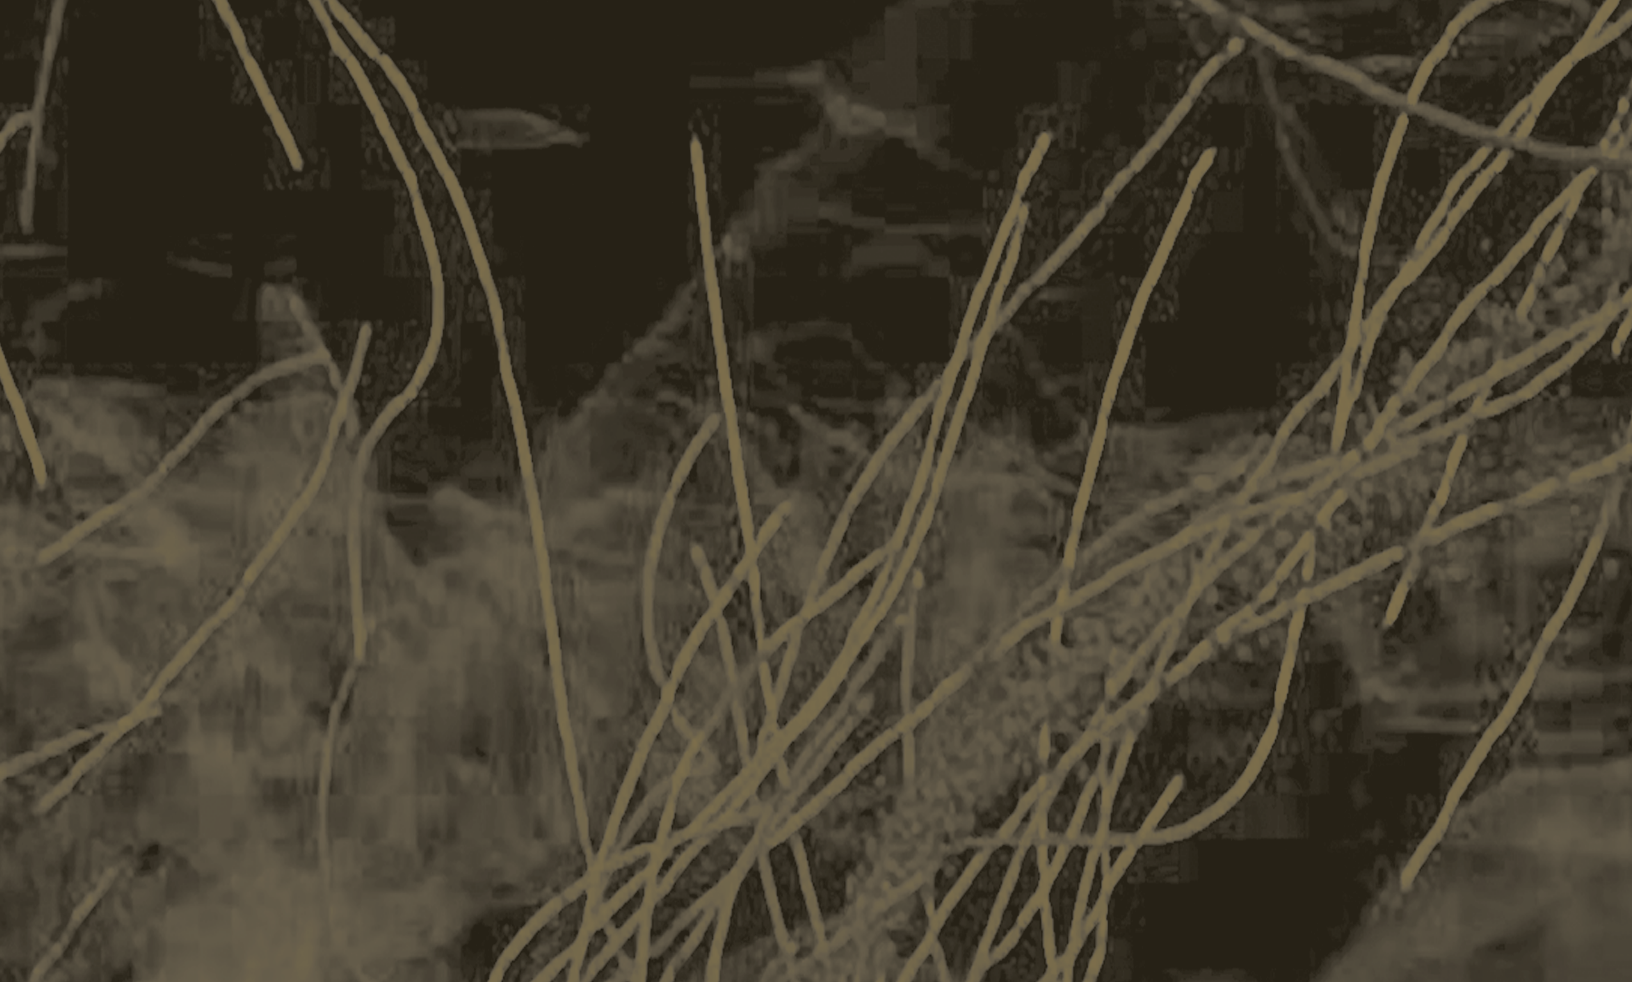

ISBN: 1-889866-28-8

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

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

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

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

Design: Soulellis Studio

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

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

Printed in the United States of America.

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

TABLE OF CONTENTS

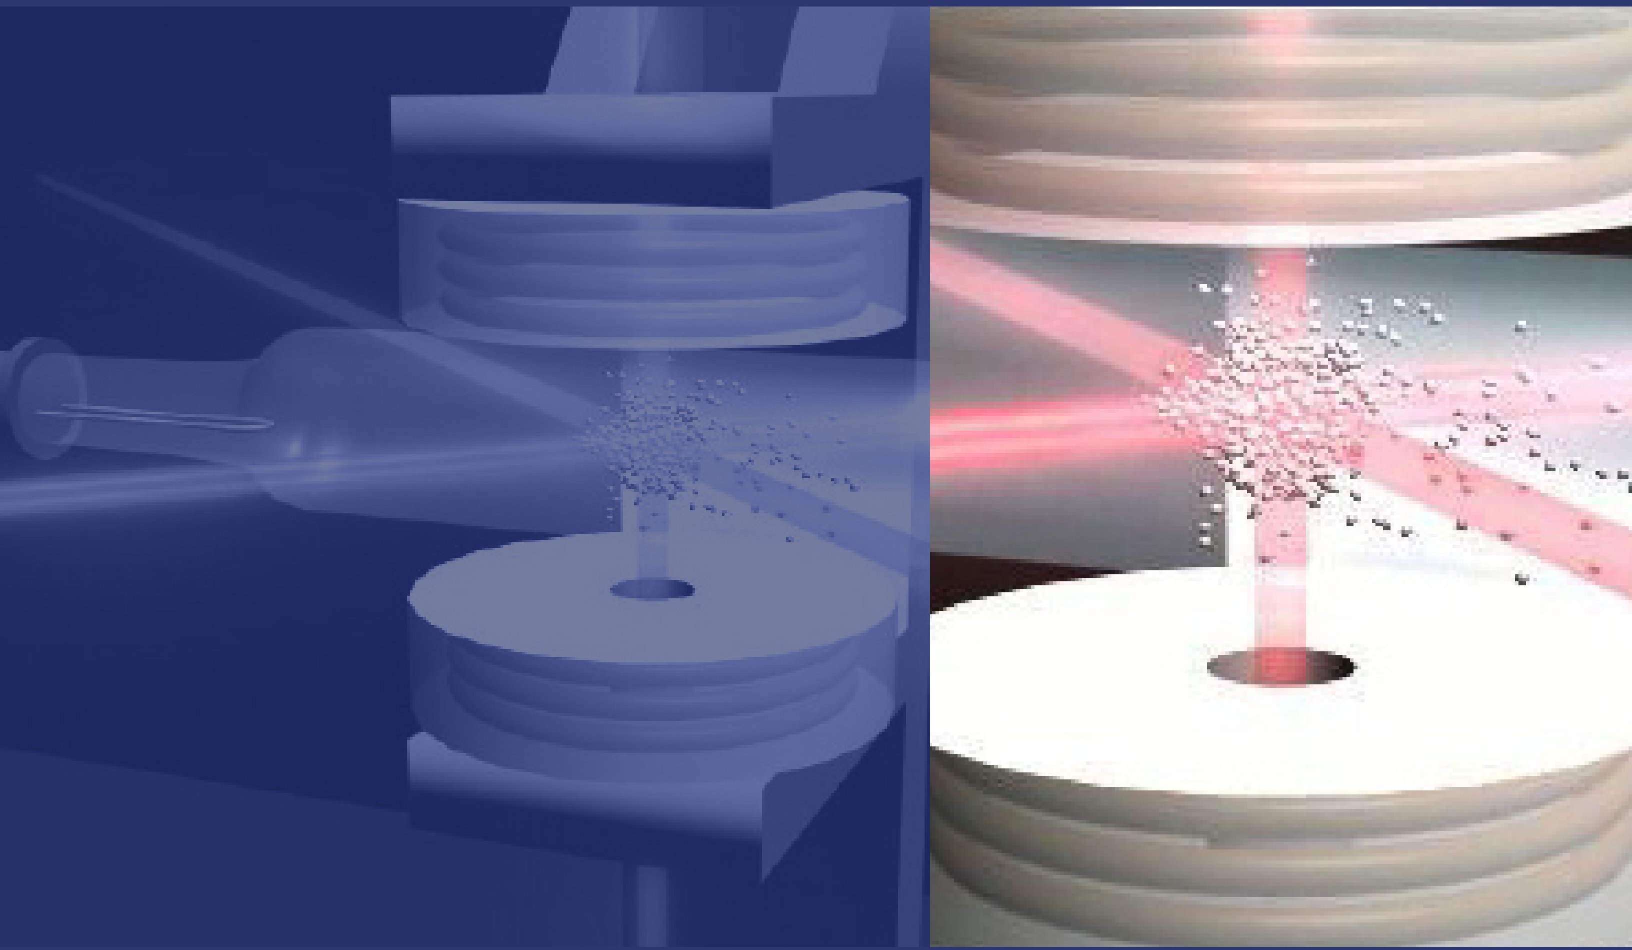

HPC simulation of atom trapping. Image courtesy of the Ohio Supercomputer Center.

PROCEEDINGS

EXECUTIVE SUMMARY

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

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

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

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

The conference began with the unveiling of "HPC in Everyday Life," an 8-minute video narrated by the penguins from the hit movie *Madagascar.* The video, created by DreamWorks Animation SKG, Inc. in collaboration with the Council, depicts how HPC is behind many of the products and services that we take for granted in everyday life. From medicine to consumer products, and energy security to aerospace, this video illustrates the importance of HPC in breakthrough innovation, and demonstrates the need for continued support.

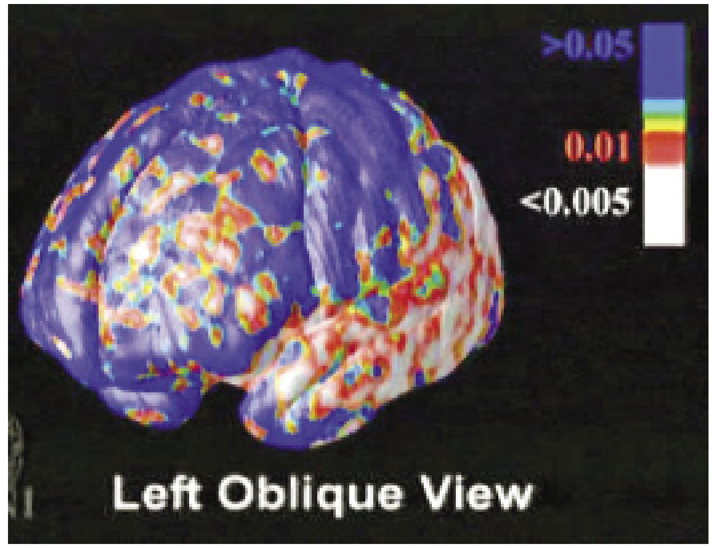

HPC simulation of Alzheimer's disease spreading through the brain of a living patient. Image courtesy of Silicon Graphics, Inc.

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

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

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

From this discussion, three significant themes emerged:

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

Three recommendations were suggested to address these issues:

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

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

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

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

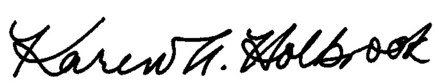

Karen A. Holbrook
President
The Ohio State University

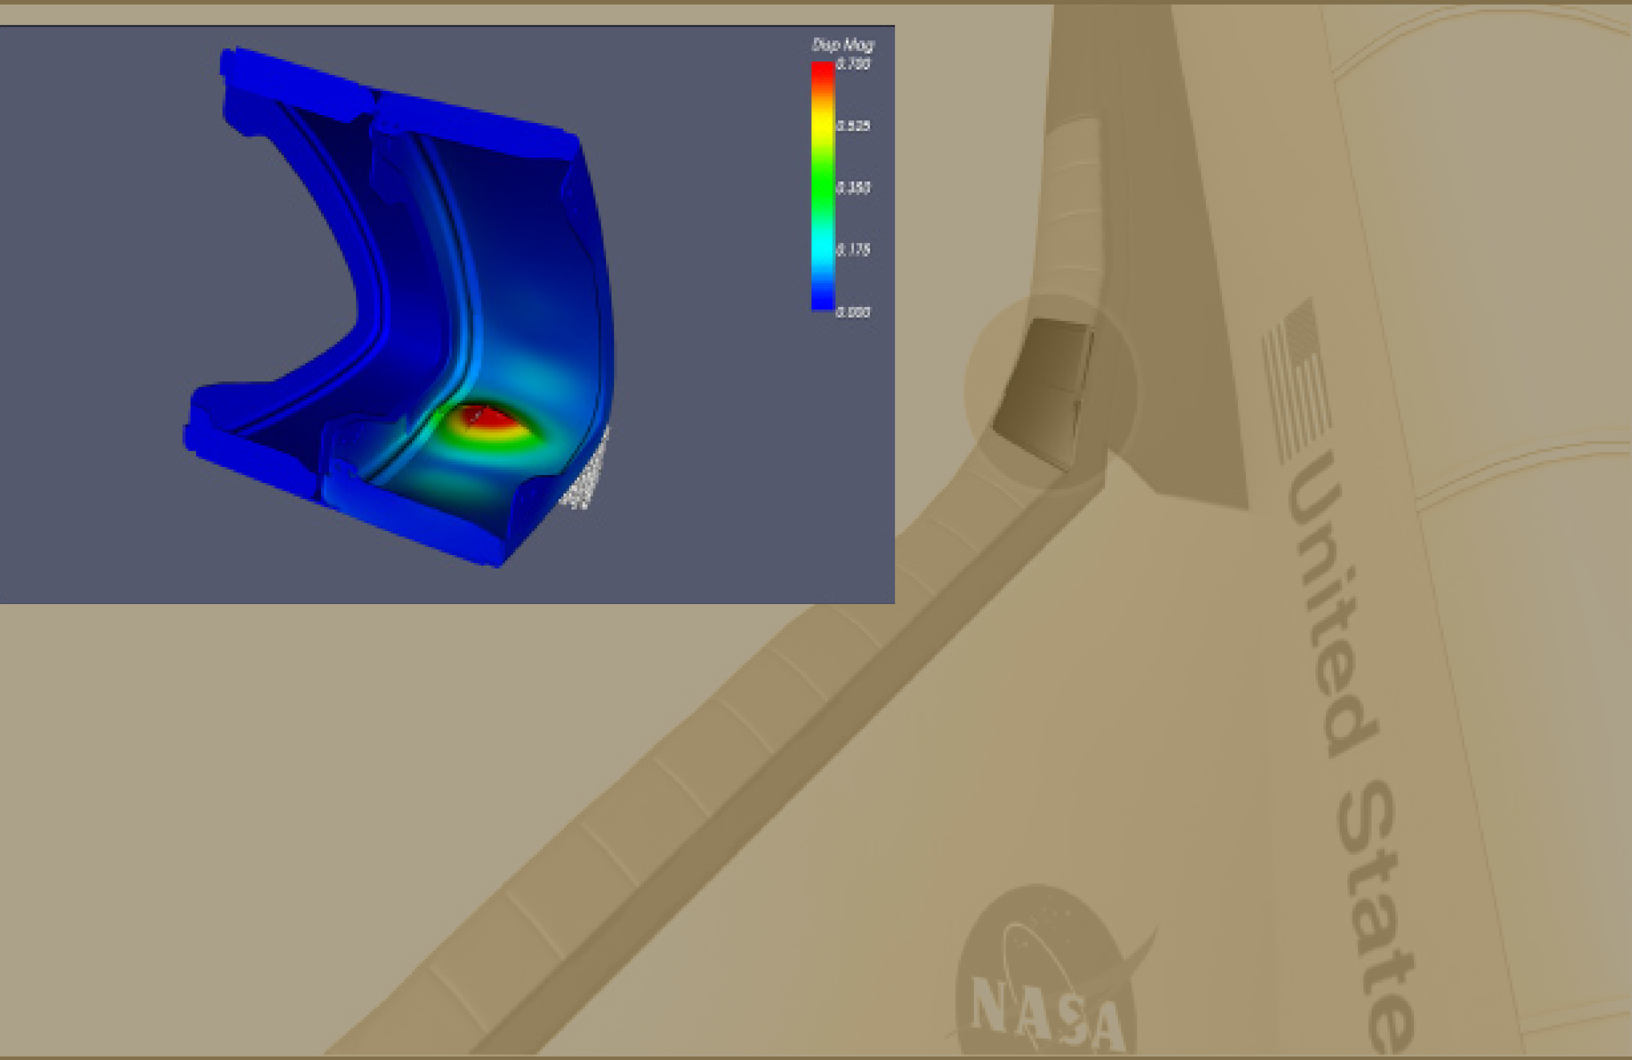

HPC simulation of the leading edge of the wing of the space shuttle Columbia, and the significant panel damage that could occur from foam impacts. Such simulations, which were later confirmed by full-scale impact tests, were critical to verifying the cause of the shuttle accident. Images courtesy of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

CONFERENCE PROCEEDINGS

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

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

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

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

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

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

Discovery and Application Walking Hand-in-Hand

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

1 The Japan Times: April 28, 2005

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

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

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

Big Changes to Big Things

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

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

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

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

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

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

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

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

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

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

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

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

Solving Grand Challenges: The Users Perspective

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

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

Understanding How the Brain Works

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

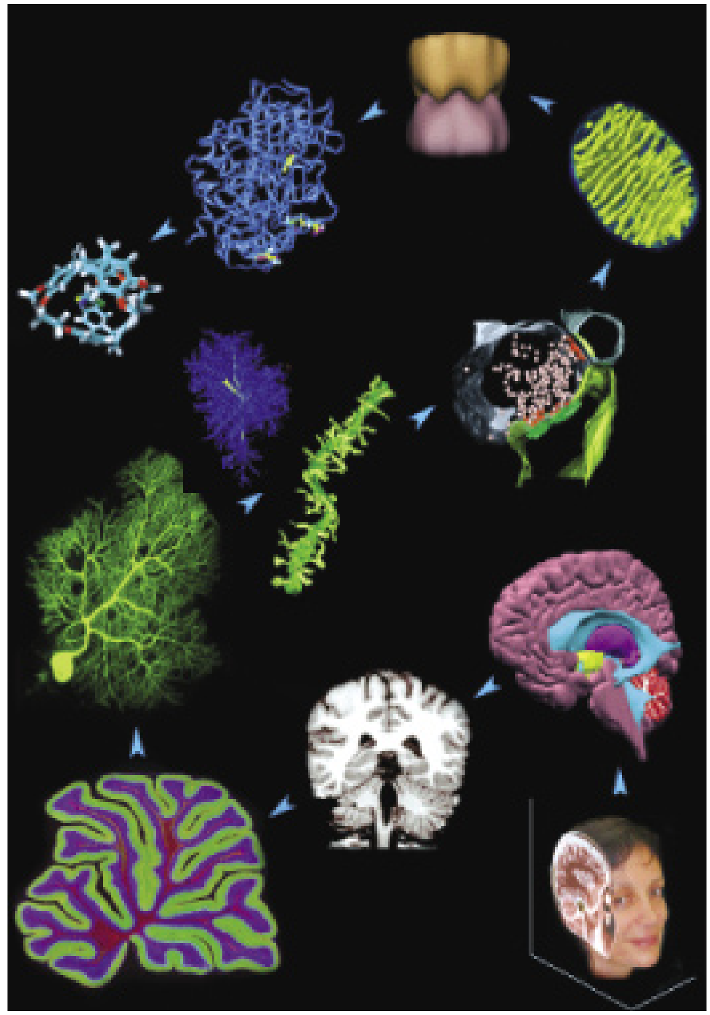

The "tyranny of scale" in brain research results from the massive amounts of data that are generated as researchers delve deeper into the cellular and atomic makeup of the human brain. Image courtesy of Dr. Mark Ellisman and the National Center for Microscopy and Imaging Research, University of California,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

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

In order to advance the study of neural pathways and molecular development, researchers are collaborating with colleagues across the United States, sharing data across distributed networks and from remote databases. Dr. Ellisman described the successful implementation of distributed computing and collaboration programs in brain research at the Biomedical Informatics Research Network (BIRN) at the University of California, San Diego. BIRN provides a robust IT infrastructure for researchers to hasten the understanding and treatment of diseases such as Alzheimer's, depression, schizophrenia, multiple sclerosis, attention deficit disorder, brain cancer, and Parkinson's. "We went with this project of linking HPC and data [because] imaging data with humans or animal models is very large," he explained. "Each one of the data sets is pulled from a different location, and they're brought together so that [the researchers] can then work on them." The project currently involves 26 research sites from 19 universities and hospitals. BIRN is creating a federated data-sharing environment that taps into biological data held at geographically separate sites as if it were a single, unified database. Growing this and other networks and continuing to link data, computational resources, scientific instruments, and people via distributed research networks will require a leap forward in computational power and capacity.

Keeping the Lights On: Managing the Energy Grid

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

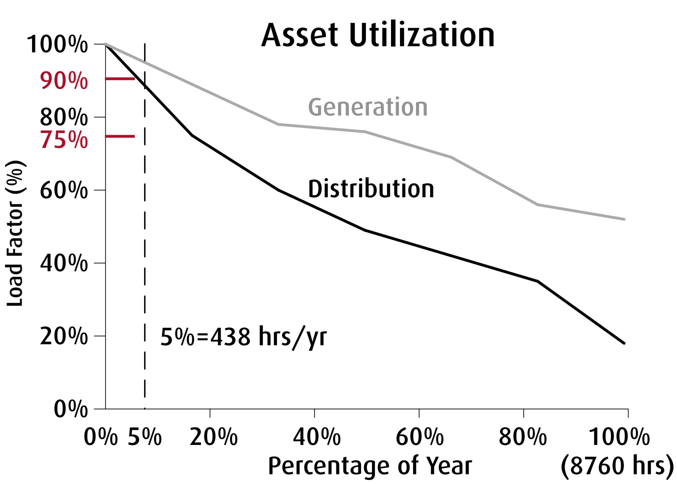

U.S. Power Grid load utilization as a function of full yearly capacity. It is important to note that nationally, the power grid is near peak loading only about 5 percent of the time.

Graphic courtesy of Pacific Northwest National Laboratory.

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

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

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

Expanding Market Share

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

America's Army®

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

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

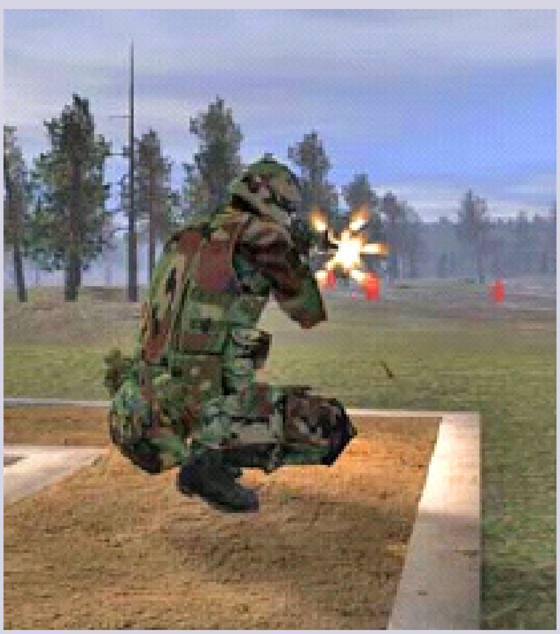

Players hone their skills in rifle range training.
Image courtesy of America's Army®

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

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

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

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

Sustaining Market Competitiveness

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

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

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

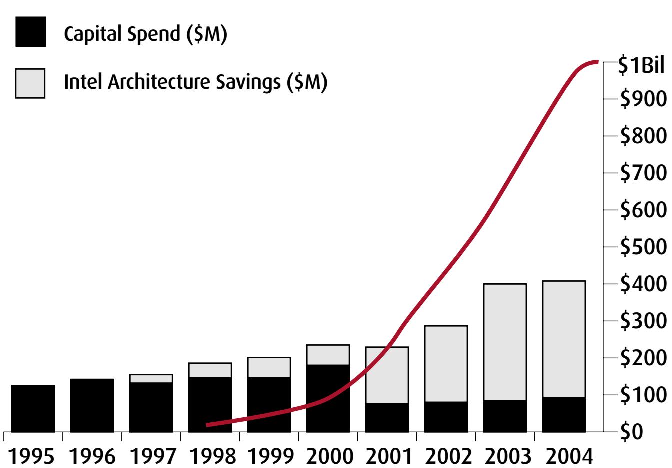

Accumulated savings as a result of Intel architecture migration. Graphic courtesy of Intel Corporation.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

When asked specifically about their plans for:

- Scaling to hundreds of processors: NO PLANS for 37 percent of codes
- Scaling to thousands of processors: NO PLANS for 44 percent of codes
- Scaling to tens of thousands of processors: NO PLANS for 60 percent of codes

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

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

In Their Own Words:

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

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

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

"Show me the business case."

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

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

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

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

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

Challenges in HPC Application Software Development

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Industrial and National Laboratory Software Environments

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

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

Moving Forward

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

Accelerate HPC software development through partnerships among HPC stakeholders.

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

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

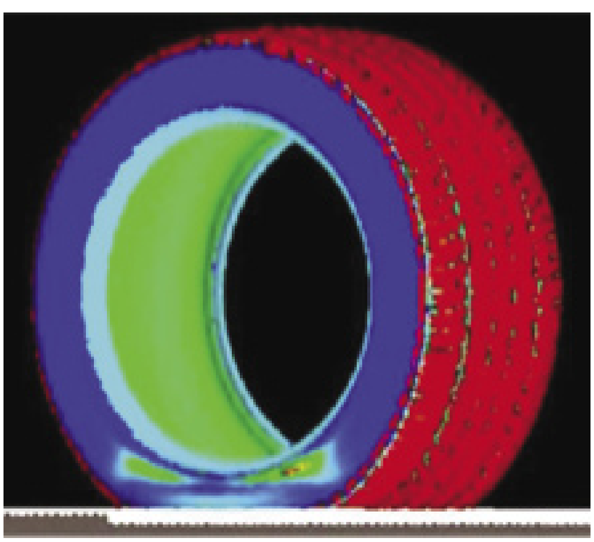

HPC simulation of tire stress points. Image courtesy of The Goodyear Tire & Rubber Company.

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

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

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

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

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

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

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

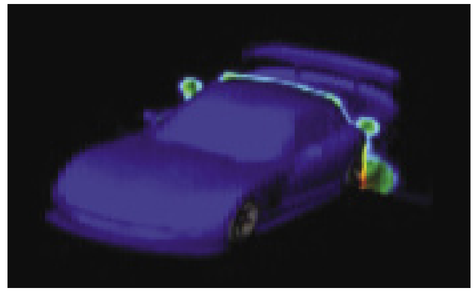

HPC simulation of NASCAR stock car aerodynamics. Image courtesy of DaimlerChrysler.

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

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

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

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

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

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

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

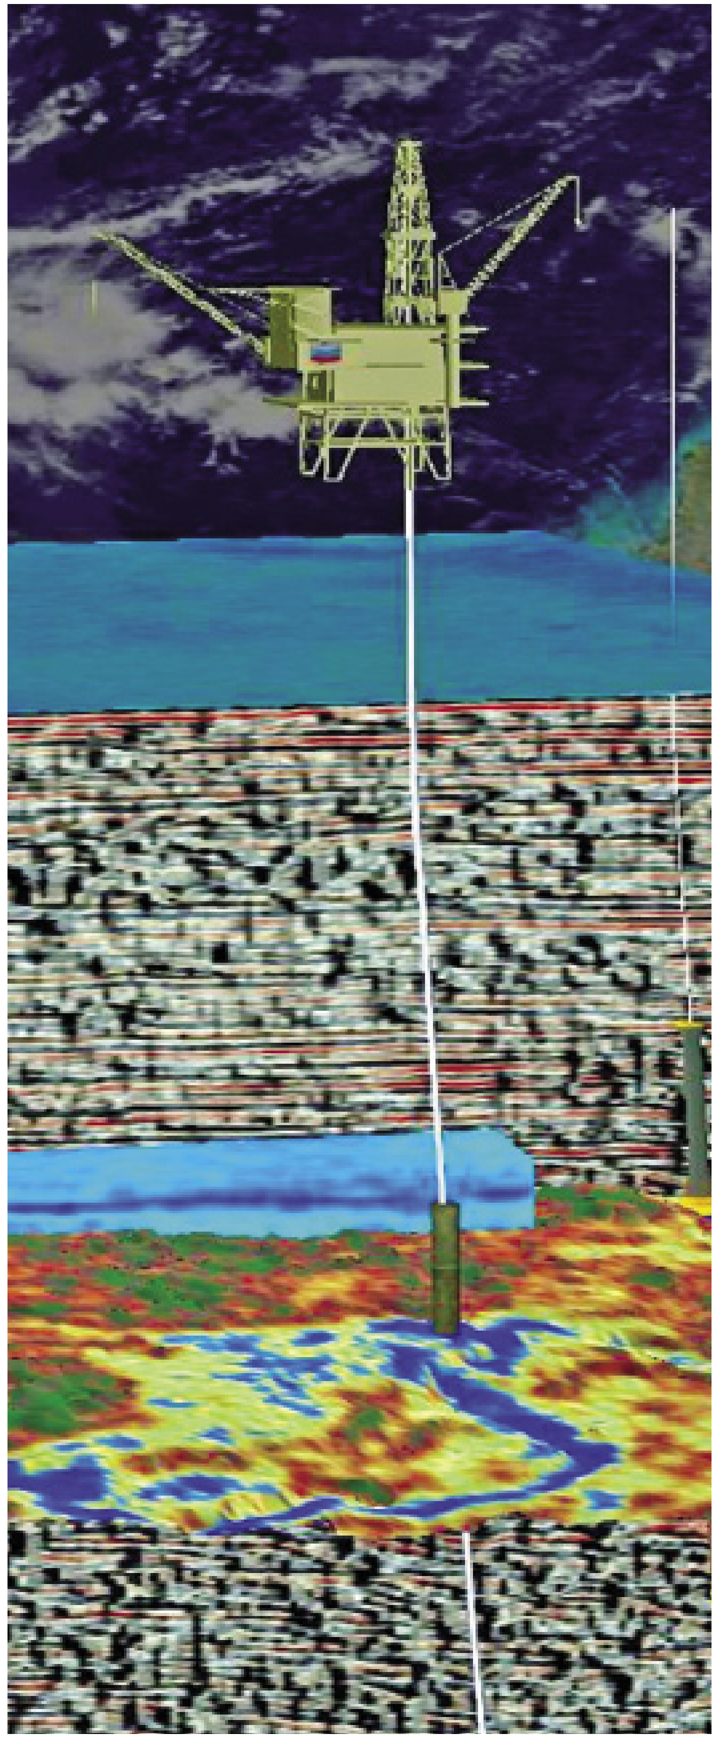

HPC simulation of a well location, tapping into the resources of an oil-filled sandy channel, colored in blue. Image courtesy of Chevron Corporation.

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

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

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

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

Conclusion

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

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

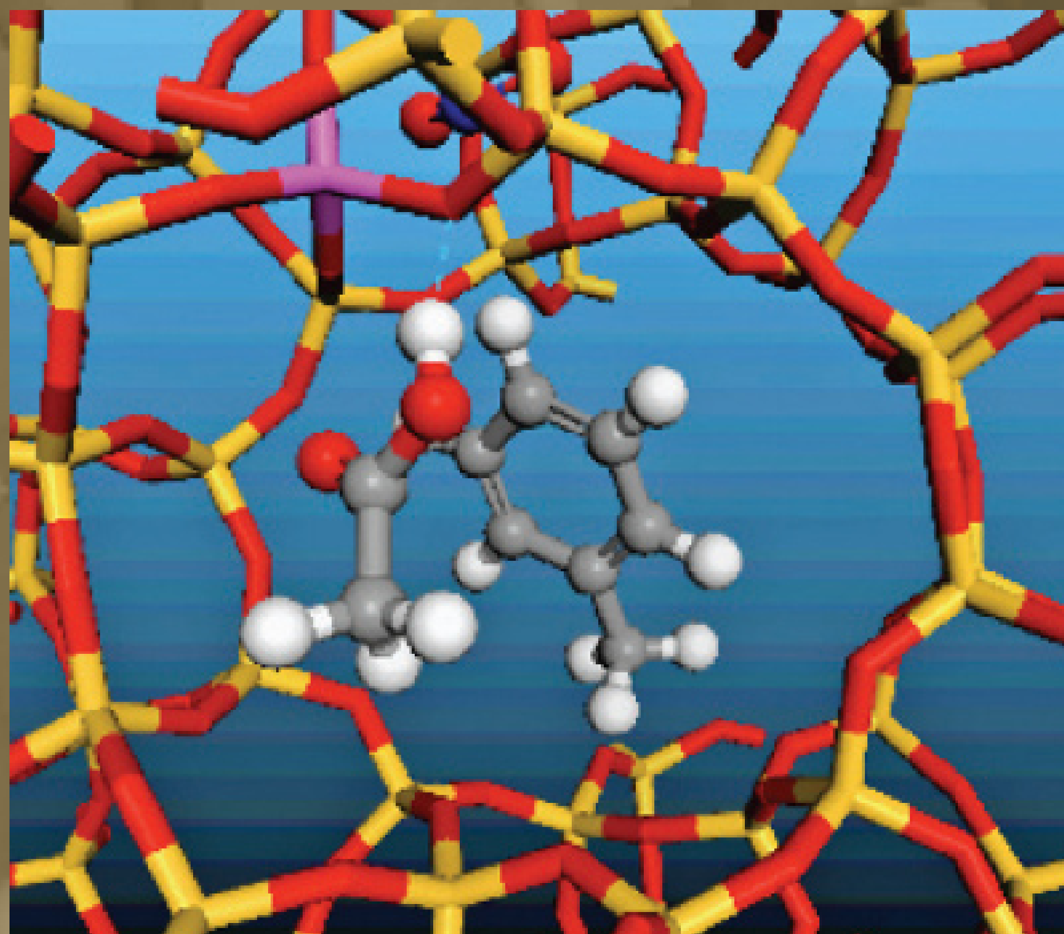

HPC simulation of an “environmentally friendly” chemical synthesis of agrochemicals and pharmaceuticals. Image courtesy of Accelrys.

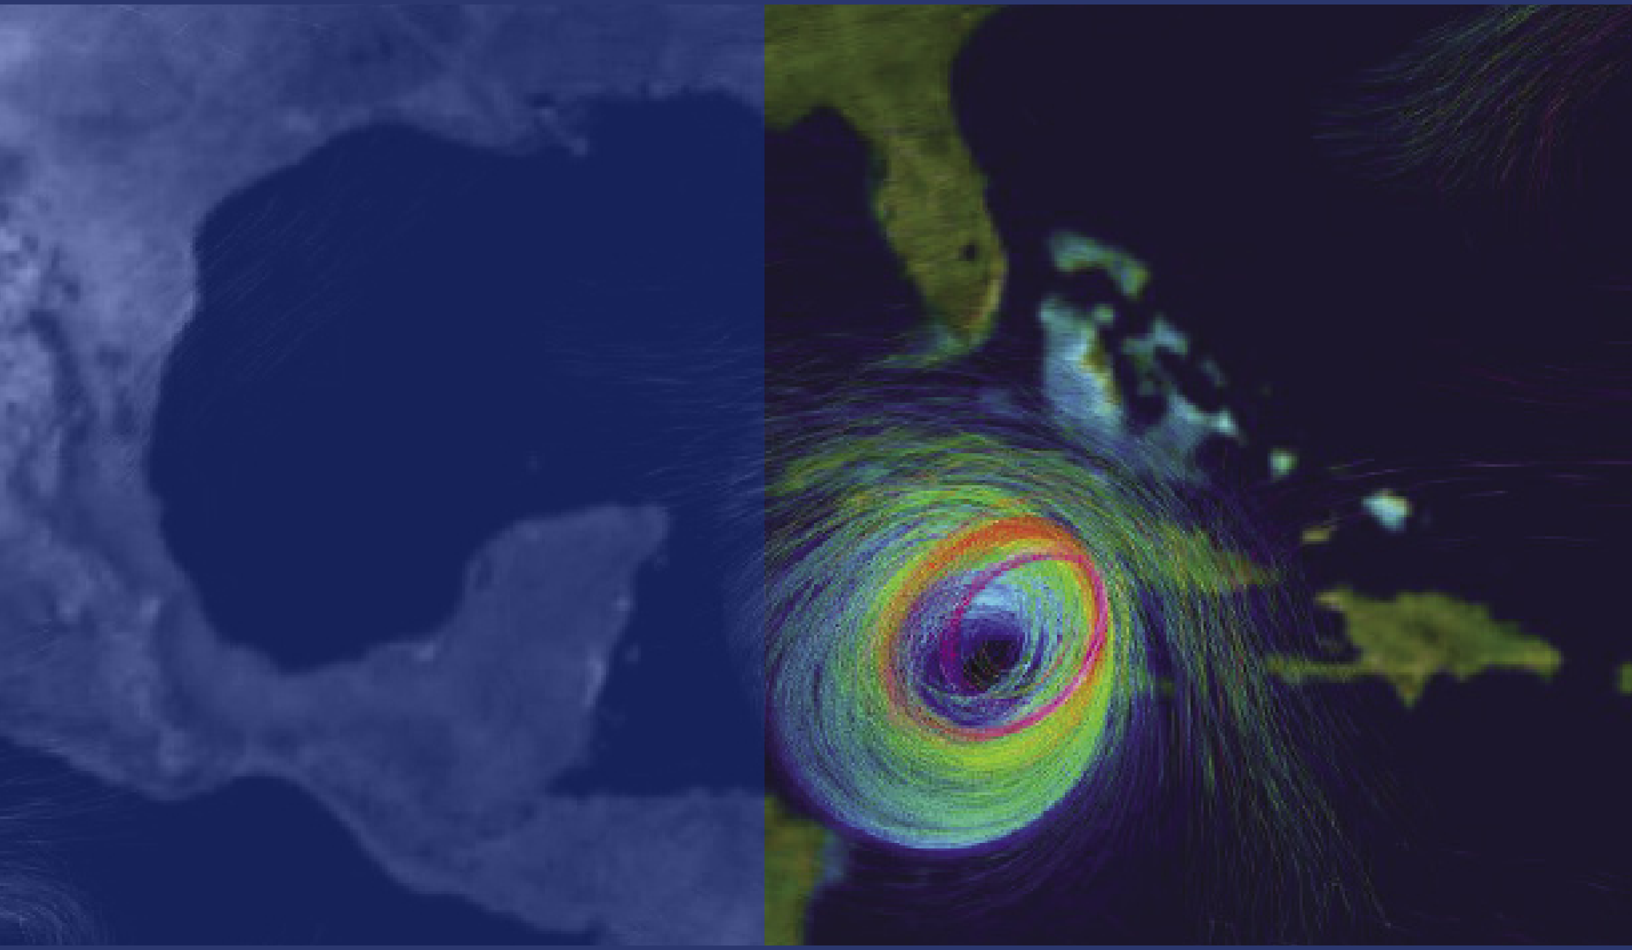

HPC simulation of Hurricane Ivan, 2004. Image courtesy of Intel Corporation.

APPENDICES

APPENDICES

Conference Agenda

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

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

Council on Competitiveness High Performance Computing Advisory Committee

Council on Competitiveness Members, Affiliates and Staff

CONFERENCE AGENDA

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

7:30 a.m. **Breakfast/Registration**

8:15 a.m. **Welcome**
- Deborah L. Wince-Smith, President, Council on Competitiveness
- Dr. Karen A. Holbrook, President, The Ohio State University, and Co-Chair, Council on Competitiveness HPC Advisory Committee

8:30 a.m. **Morning Keynote Address: A Challenge for Our Future**
- Roger Enrico, former CEO, PepsiCo; current Chairman, DreamWorks Animation, SKG, Inc.

9:15 a.m. **Panel 1: High Performance Computing - Driving Solutions to Grand Challenge Problems**
Panelists will discuss cutting edge, "grand challenge" problems that require leaps in software development to successfully address them, and whose successful solution will have a significant impact on society and propel competitiveness and prosperity.
Moderator: Dr. Jeffrey Wadsworth, Director, Oak Ridge National Laboratory
Panelists:
- Guru Bhatia, General Manager and Director, Engineering Computing Information Technology Group, Intel Corporation
- Dr. Mark H. Ellisman, Professor of Neurosciences and Bioengineering, University of California, San Diego, School of Medicine and Director, the Center for Research on Biological Systems, The National Center for Microscopy and Imaging Research and the Biomedical Informatics Research Network (BIRN) Coordinating Center, University of California, San Diego
- Dr. George Michaels, Associate Laboratory Director, Computational and Information Sciences Directorate, Pacific Northwest National Laboratory
- Dr. Michael Zyda, Director, GamePipe Laboratory, Viterbi School of Engineering, University of Southern California Information Sciences Institute

10:45 a.m. **Break**

11:15a.m. **Results of the Study of Independent Software Vendors Serving the HPC Market**
Study results will be released revealing the state of independent software vendors (ISVs) globally, which industries/companies are dependent on which ISVs, user and ISV perspectives of the strengths and weaknesses of the ISV industry, and areas where partnerships could help to advance software development.
- Dr. Earl Joseph, Research Vice President, High Performance Systems Program, International Data Corporation

12:15 p.m. **Luncheon/Keynote:**
- Dr. Arden L. Bement, Jr., Director, National Science Foundation

1:45 p.m. **Panel 2: High Performance Computing: Bridging the Software Gap**
Industry, university and government panelists will discuss the challenges of using, maintaining, and creating application software suitable for a competitive, corporate "production" environment, and the role of universities and national laboratories to help accelerate development of new and/or updated code. The panel format will be a roundtable discussion
Moderator: Dr. Graham B. Spanier, President, The Pennsylvania State University
Panelists:
- Paul Bemis, Vice President Marketing, Fluent Inc.
- Dr. William Camp, Director, Computation, Computers, Information and Mathematics, Sandia National Laboratories
- Dona Crawford, Associate Director Computation, Lawrence Livermore National Laboratory
- Loren Miller, Director, IT Research, Development & Engineering, The Goodyear Tire & Rubber Company
- Dr. Donald Paul, Vice President and Chief Technology Officer, Chevron Corporation
- Dr. Reza Sadeghi, Vice President, Enterprise Computing, MSC.Software Corporation

3:30 p.m. **Next Steps/Building the Strategy**
- Dr. David E. Shaw, Chairman, D.E. Shaw & Co., Inc. and Co-chair, Council on Competitiveness HPC Advisory Committee

4:00 p.m. **Adjourn**

EXECUTIVE SUMMARY: ISV STUDY

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

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

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

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

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

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

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

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

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

EXECUTIVE SUMMARY: ISV STUDY

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

Key Findings

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

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

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

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

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

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

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

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

EXECUTIVE SUMMARY: ISV STUDY

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

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

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

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

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

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

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

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: ISV STUDY

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

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

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

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

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

Original Source of ISV Application

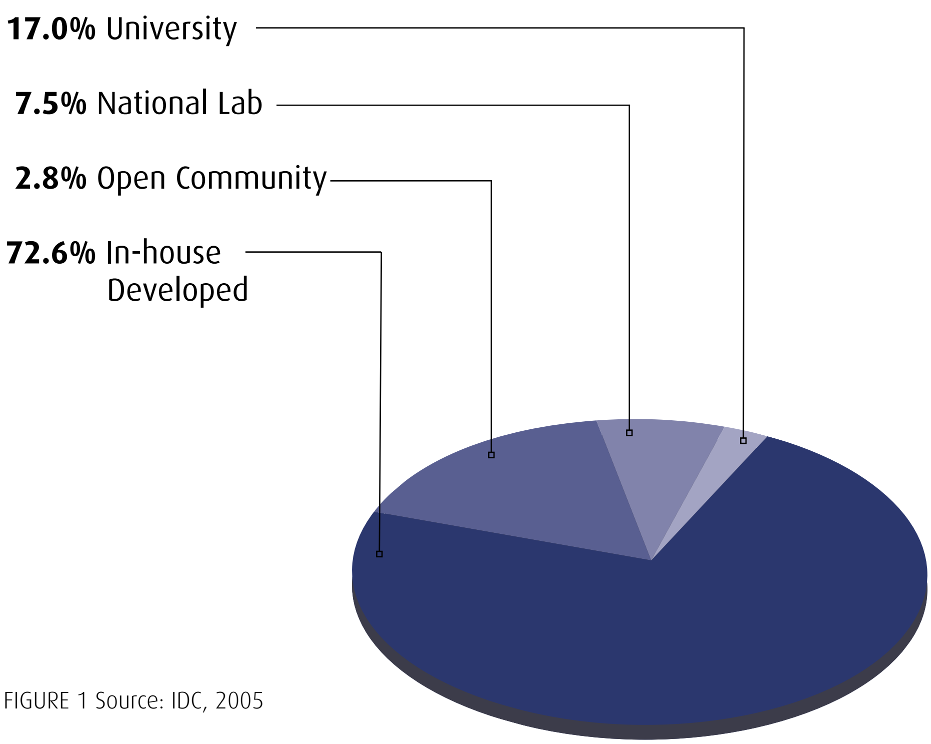

FIGURE 1 Source: IDC, 2005

EXECUTIVE SUMMARY: ISV STUDY

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

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

Percentage of Applications by Type of Ownership

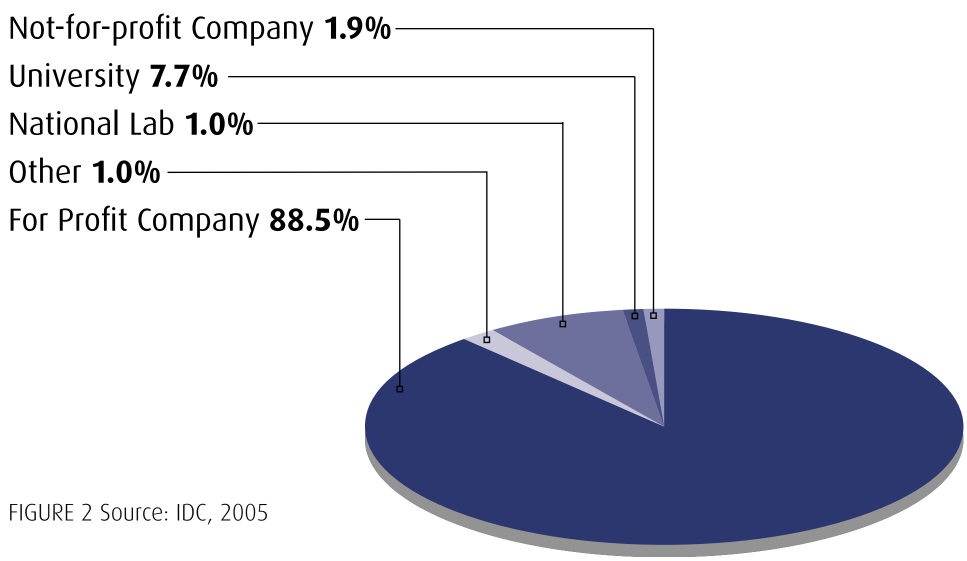

FIGURE 2 Source: IDC, 2005

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

There is a lack of readiness for petascale systems

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

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

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

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

EXECUTIVE SUMMARY: ISV STUDY

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

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

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

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

Number of ISV Applications and Companies by Total Company Revenue

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

TABLE 1 Source: IDC, 2005

EXECUTIVE SUMMARY: ISV STUDY

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

Revenue for fundamental rewrites is generally not available.

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

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

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

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

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

Partners ISVs Selected as Potentially Most Useful, by Application

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

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

EXECUTIVE SUMMARY: ISV STUDY

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

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

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

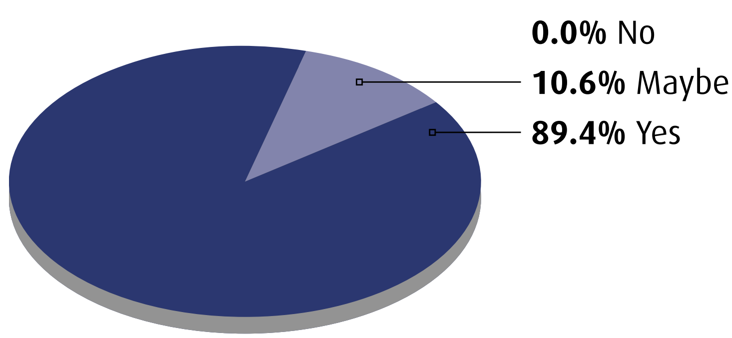

FIGURE 3 Source: IDC, 2005

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

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

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

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

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

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

EXECUTIVE SUMMARY: SOFTWARE WORKSHOP

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

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

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

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

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

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

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

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

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

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

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

EXECUTIVE SUMMARY: SOFTWARE WORKSHOP

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

algorithms that are more comprehensive and accurate.

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

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

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

Recommendations for ISVs:

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

Recommendations for Universities:

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

Recommendations for the Federal Government:

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

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

COUNCIL ON COMPETITIVENESS HIGH PERFORMANCE COMPUTING ADVISORY COMMITTEE

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

University Co-Chair
Karen Holbrook
President
The Ohio State University

PRIVATE SECTOR

Lalitha Subramanian
Director, Contract Research Services
Accelrys

John F. Sparks
Director, Engineering & Technology Programs
Aerojet

J. D. Thompson
Technical Services
ATK Thiokol Propulsion

William Madia
Executive Vice President, Laboratory Operations
Battelle Memorial Institute

Hank Kafka
Vice President,
Architecture and Emerging Technology
BellSouth

Frederick H. Hausheer
Chairman and Chief Executive Officer
BioNumerik Pharmaceuticals

Suresh N. Shukla
HPC Service Manager
The Boeing Company

Donald Paul
Vice President and Chief Technology Officer
Chevron Corporation

Keshav Pingali
Associate Director, Cornell Theory Center
Cornell University

Rich Grzybowski
Director, Systems Engineering
Corning, Inc.

Peter Ungaro
President
Cray Inc.

James Shanks
President, CDW Government
CDW Corporation

John Picklo
Manager,
Mainframes and High Performance Computing
DaimlerChrysler

Ed Leonard
Chief Technology Officer
DreamWorks Animation SKG

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

Paul Bemis
Vice President Marketing
Fluent Inc.

Vincent Scarafino
Manager, Numerically Intensive Computing
Ford Motor Company

Stewart Moore
Principal Engineer
General Dynamics Electric Boat

Sharan Kalwani
CAE HPC Manager
General Motors

Stephen Squires
Chief Science Officer
Hewlett-Packard Company

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

David Turek
Vice President, Deep Computing
IBM

Earl Joseph
Research Vice President,
High Performance Systems
IDC

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

Ben Bennett
Director, HPC Programs Office
Intel Corporation

Eric J. Pitcher
Vice President, Technical Marketing
Linux Networx

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

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

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

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

Reza Sadeghi
Vice President, Enterprise Computing
MSC.Software Corporation

Stan Ahalt
Executive Director
Ohio Supercomputer Center

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

Beverly Clayton
Executive Director
Pittsburgh Supercomputing Center

Greg Brandeau
Vice President of Technology
Pixar Animation Studios

Saadat Syed
Senior Fellow
Pratt & Whitney

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

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

COUNCIL ON COMPETITIVENESS
HIGH PERFORMANCE COMPUTING ADVISORY COMMITTEE

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

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

George Vradenburg, III
President
Vradenburg Foundation

PUBLIC SECTOR

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

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

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

Peter Freeman
Assistant Director for CISE
National Science Foundation

Dona Crawford
Associate Director Computation
Lawrence Livermore National Laboratory

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

Stephen Younger
Senior Fellow, Theoretical Division
Los Alamos National Laboratory

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

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

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

COUNCIL ON COMPETITIVENESS MEMBERS

COUNCIL ON COMPETITIVENESS BOARD

Chairman
Charles O. Holliday, Jr.
Dupont

Chairman Emeritus
F. Duane Ackerman
BellSouth Corporation

Vice Chairman
G. Wayne Clough
Georgia Institute of Technology

President
Deborah L. Wince-Smith

EXECUTIVE COMMITTEE

Robert R. Bishop
Silicon Graphics, Inc.

Molly Corbett Broad
University of North Carolina

William R. Brody
Johns Hopkins University

Jared L. Cohon
Carnegie Mellon University

John J. DeGioia
Georgetown University

Gary T. DiCamillo
TAC Worldwide Companies

Robert C. Dynes
University of California

John A. Edwardson
CDW Corporation

John M. Engler
National Association of Manufacturers

William D. Green
Accenture

Sheryl Handler
Ab Initio

John L. Hennessy
Stanford University

Susan Hockfield
Massachusetts Institute of Technology

Shirley Ann Jackson
Rensselaer Polytechnic Institute

Henry A. McKinnell, Jr.
Pfizer Inc

John P. Morgridge
Cisco Systems, Inc.

Samuel J. Palmisano
IBM Corporation

Michael E. Porter
Harvard University

Luis M. Proenza
University of Akron

Patricia F. Russo
Lucent Technologies Inc.

Kenan Sahin
TIAX LLC

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

Lawrence H. Summers
Harvard University

Lawrence Weber
W2Group

DISTINGUISHED FELLOWS

Erich Bloch
Washington Advisory Group

Daniel S. Goldin
The Intellisis Corporation

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

Thomas Ridge
Former Secretary
Department of Homeland Security

GENERAL MEMBERSHIP

Michael F. Adams
The University of Georgia

Richard A. Anthes
University Corporation for Atmospheric Research

William F. Ballhaus, Jr.
Aerospace Corporation

Steven A. Ballmer
Microsoft Corporation

David Baltimore
California Institute of Technology

Craig R. Barrett
Intel Corporation

Richard A. Bendis
Innovation Philadelphia

Robert J. Birgeneau
University of California, Berkeley

James H. Blanchard
Synovus Financial Corporation

Lee C. Bollinger
Columbia University

David L. Boren
The University of Oklahoma

Richard H. Brodhead
Duke University

Michael J. Burns
Dana Corporation

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

Curtis R. Carlson
SRI International

Albert Carnesale
University of California, Los Angeles

Carol A. Cartwright
Kent State University

John T. Casteen, III
University of Virginia

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

Mary Sue Coleman
University of Michigan

France A. Cordova
University of California, Riverside

Denis A. Cortese
Mayo Clinic

Margaret B. Cozzens
Colorado Institute of Technology

Michael M. Crow
Arizona State University

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

COUNCIL ON COMPETITIVENESS MEMBERS

Denise D. Denton
University of California, Santa Cruz

Peter R. Dolan
Bristol-Myers Squibb Company

Michael V. Drake
University of California, Irvine

Michael T. Duke
Wal-Mart Stores, Inc.

Roger A. Enrico
DreamWorks Animation
SKG Inc.

Marye Anne Fox
University of California, San Diego

Robert M. Gates
Texas A & M University

E. Gordon Gee
Vanderbilt University

Thomas F. George
University of Missouri, St. Louis

James W. Griffith
The Timken Company

Amy Gutmann
University of Pennsylvania

Peter Halpin
World Resources Company

David C. Hardesty, Jr.
West Virginia University

Jack Harding
eSilicon Corporation

Robert Hemenway
The University of Kansas

Adam W. Herbert
Indiana University

Richard Herman
University of Illinois, Urbana Champaign

Karen A. Holbrook
Ohio State University

Jeffery R. Immelt
General Electric Company

Irwin M. Jacobs
QUALCOMM, Inc.

John I. Jenkins
University of Notre Dame

Martin C. Jischke
Purdue University

Carl F. Kohrt
Battelle Memorial Institute

Alan G. Lafley
The Procter & Gamble Company

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

Martin G. McGuinn
Mellon Financial Corporation

W. James McNerney, Jr.
The Boeing Company

Peggy Gordon Miller
South Dakota State University

Richard K. Miller
Franklin W. Olin College of Engineering

James B. Milliken
University of Nebraska

Clayton Daniel Mote, Jr.
University of Maryland

Anne M. Mulcahy
Xerox Corporation

Diana S. Natalicio
University of Texas, El Paso

Mark A. Nordenberg
University of Pittsburgh

Peter O'Donnell, Jr.
O'Donnell Foundation

Constantine Papadakis
Drexel University

Antonio M. Perez
Eastman Kodak Company

Peter G. Peterson
Blackstone Group

James H. Possehl
Republic Financial Corporation

Harold J. Raveché
Stevens Institute of Technology

Steven S Reinemund
Pepsi-Cola Company

Robert L. Reynolds
Fidelity Investments

Michael C. Ruettgers
EMC Corporation

Henry Samueli
Broadcom Corporation

Carl J. Schramm
Ewing Marion Kauffman Foundation

Pamela Sedmak
Ernst Sedmak, LLC

Ivan G. Seidenberg
Verizon Communications

Kevin Sharer
Amgen

Sanford C. Shugart
Valencia Community College

Lou Anna K. Simon
Michigan State University

Albert J. Simone
Rochester Institute of Technology

John B. Simpson
State University of New York at Buffalo

David J. Skorton
University of Iowa

Andrew A. Sorenson
University of South Carolina

Graham B. Spanier
The Pennsylvania State University

Ray Stata
Analog Devices, Inc.

Susan S. Stautberg
Partner Com Corporation

Charles W. Steger
Virginia Polytechnic Institute and State University

Andrew L. Stern
Service Employees International Union

John D. Stobo
University of Texas Medical Branch

COUNCIL ON COMPETITIVENESS MEMBERS, AFFILIATES & STAFF

Matthew J. Szulik
Red Hat, Inc.

Henri A. Termeer
Genzyme Corporation

Lydia Waters Thomas
Mitretek Systems Inc.

Lee T. Todd, Jr.
University of Kentucky

Peter J. Ungaro
Cray Inc.

Steven L. VanAusdle
Walla Walla Community College

Larry N. Vanderhoef
University of California, Davis

George Vradenburg, III
Vradenburg Foundation

G. Richard Wagoner, Jr.
General Motors

Edie Weiner
Weiner, Edrich, Brown, Inc

William C. Weldon
Johnson & Johnson

Jack M. Wilson
The University of Massachusetts

James H. Woodward
University of North Carolina, Charlotte

James Wright
Dartmouth College

Mark S. Wrighton
Washington University

Henry T. Yang
University of California, Santa Barbara

Paul A. Yarossi
HNTB Corporation

Steven G. Zylstra
Pittsburgh Technology Council

NATIONAL AFFILIATES

American Association for the Advancement of Science

American Electronics Association

American Mathematical Society

American Petroleum Institute

American Society for Engineering Education

American Society for Quality

ASME International

Association of American Universities

Association of University Related Research Parks

BITS' Financial Services Roundtable

Community Learning & Information Network

Computer Systems Policy Project

The Conference Board, Inc.

Council for Chemical Research

Council on Governmental Relations

IEEE-USA

Iowa Business Council

The Japan Society

JumpStart, Inc

National Association of Management and Technical Assistance Centers

National Center for Manufacturing Sciences

National Center for Women & Information Technology

NorTech

Oak Ridge Associated Universities

SMC3

*Current as of January 17, 2006

COUNCIL ON COMPETITIVENESS STAFF

Deborah L. Wince-Smith
President

C. William Booher, Jr.
Chief Operating Officer

Debra Van Opstal
Senior Vice President

William C. Bates
Vice President, Congressional Outreach

Merryl Burpoe
Vice President, Global Initiative

Chad Evans
Vice President, National Innovation Initiative

Randall T. Kempner
Vice President, Regional Innovation

Mohamed Khan
Vice President, Information Systems

Betsy Thurston
Vice President for Strategic Development

Suzy Tichenor
Vice President and Director, High Performance Computing Initiative

Jennifer Bond
Senior Advisor, Global Initiative

David Attis
Director, Policy Studies

Melyssa Fratkin
Policy Director, High Performance Computing Initiative

Julie Jordan
Policy Director, International

Gabrielle Trebat
Director, International Policy Studies

Marcy Jones
Executive Assistant

Mary Marchal
Senior Research Associate

Jennifer Mateik
Controller

Lars-eric Rödén
Research Associate,
National Innovation Initiative

Amanda Welch
Manager, Policy Studies

About the Council on Competitiveness

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

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

Third Annual High Performance Computing Users Conference Report

Moving Beyond Islands Of Innovation
September 7, 2006
Washington, D.C.

TABLE OF CONTENTS

Proceedings

MOVING BEYOND ISLANDS OF INNOVATION

The Council's third annual High Performance Computing Users Conference, held during our twentieth anniversary year, marked an important turning point for the Council's—and America's—efforts to make the proven innovation, productivity and competitiveness benefits of high performance computing more pervasive within U.S. industry.

The federal government stepped up its support for High Performance Computing (HPC) through the American Competitiveness Initiative. This initiative, announced in 2006, seeks to double the federal commitment over the next decade to basic research in the physical sciences—including "promising areas such as…supercomputing." In 2006, federal programs aimed at accelerating innovation by giving companies access to some of the nation's most powerful computers also stepped up their investments and activities, including the Department of Energy's INCITE and SciDAC programs and others involving the National Science Foundation (NSF) and the National Nuclear Security Administration (NNSA). Council studies during the year found that the NSF and NNSA public-private partnership programs have been overwhelmingly successful, and the DOE once again chose to announce the INCITE award winners in conjunction with the Council.

Three years of pioneering Council research under the guidance of our HPC Advisory Committee laid the groundwork for announcing the National Innovation Collaboration Ecosystem (NICE), an exciting new partnership program we are leading with the University of Southern California's Information Sciences Institute (ISI). NICE will boost America's national productivity and competitiveness by making HPC systems and expertise more broadly available to companies, entrepreneurs and even solo inventors. The Council-sponsored research identified the benefits of HPC for industry, the barriers preventing wider usage, and successful models of public/private sector collaborations involving HPC. NICE will serve as a catalyst to spread public-private partnerships and help turn today's islands of HPC-based innovation into a national American infrastructure for world-leading productivity and competitiveness.

The Council's HPC activity supports our larger National Innovation Initiative (NII), whose ground-breaking "Innovate America" report clearly established that innovation is the surest path to sustained economic growth and global competitiveness. Other governments are pursuing aggressive strategies to strengthen their innovation capacity and link innovation with economic development. America's business, university, and labor leaders agree that simply doing business the way it has always been done will not be enough to sustain leadership – either for the country or for U.S. companies.

At the same time, the nature of innovation is rapidly changing. Historically, innovation occurred mainly through the efforts of individual investigators working in single, sharply demarcated disciplines. Today, the biggest advances increasingly come from multidisciplinary collaborations. For studying disease pathways through the body, knowledge is needed about physics, chemistry, biology and in some cases also nanotechnology. Building superior cars requires the ability to look concurrently at interdependent factors, including crashworthiness, aerodynamics, fuel-efficiency, cabin noise and vibration, and ride harshness. No technology has shown greater ability to energize multidisciplinary, collaborative innovation and propel competitiveness than HPC. HPC can shrink time to insight, time to market and time to competitive advantage.

HPC is essential for addressing some of the nation's grand challenges: energy independence, protection of critical infrastructure (e.g., power, telecommunications, financial and transportation systems), and scientific leadership. But the real impetus for innovation occurs at the crossroads between businesses, national labs and other research centers, universities and skilled workers throughout the country. Regions have become the critical nodes for innovation-based economic growth, and some U.S. regions are already exploiting HPC in their economic development plans. Much more of this is needed.

The Council’s HPC Advisory Committee will work closely with the Council and ISI in 2007 to prepare NICE for implementation. We want to ensure that NICE responds well to the major needs identified in the Council's research: recognizing the collaborative, multidisciplinary nature of innovation today; learning from successful public-private sector partnerships; supporting experienced and novice HPC users in the private sector; and helping regions to leverage their innovation assets.

The Council's fundamental belief is that U.S. competitiveness and the nation's ability to add high-value economic activity increasingly depend on 21st-century, HPC-based modeling and simulation. We look forward to helping make HPC usage more pervasive, so that U.S. businesses can stay in front of the competition and Americans can enjoy greater prosperity

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

Karen A. Holbrook
President
The Ohio State University

CONFERENCE PROCEEDINGS

Needed: A National HPC Ecosystem For Turbocharging Innovation

Council on Competitiveness President Deborah L. Wince-Smith welcomed attendees to the Third Annual HPC Users Conference on behalf of the Council and its conference partners: the Defense Advanced Research Projects Agency, the National Nuclear Security Administration, the Department of Energy Office of Science, and the National Science Foundation.

Three years ago, under the guidance of its HPC Advisory Committee, the Council launched a coordinated program of original research, conferences and workshops to learn how businesses were using high performance computing in practice, how HPC was contributing to innovation and productivity, and what barriers were preventing wider use of HPC in the commercial sector. The primary finding: HPC is indispensable for the survival of nearly all businesses that have adopted it, including some of America's largest companies. HPC is driving much of the innovation that fuels U.S. competitiveness. It is increasingly true that to out-compete, we must out-compute.

The proven competitiveness benefits of HPC are far from being fully exploited today, however. A relatively small group of highly experienced HPC users in industry are shining examples of this technology's effectiveness and stand out as islands of leading-edge innovation. They are surrounded by a much larger group of entry-level HPC users who have not moved up the performance curve to realize the full benefits of HPC. This lack of advancement gives the HPC market a "missing middle" that represents a substantial productivity loss for our country. An even larger number of American businesses—primarily small and medium-size entrepreneurial firms—have never tapped into the advantages of HPC because they view it as beyond their means. We call this group the "never evers." In addition, many U.S. regions and states do not yet understand that HPC is crucial for economic development, for attracting the best-and-brightest companies and individuals. Council research has shown that even for the highly experienced users of this important technology, inadequate production-quality application software[1] and access to needed talent are barriers to more aggressive exploitation of HPC.

A new 21st-century infrastructure is needed to make HPC more widely available for national, regional and business prosperity, and to increase the return on America's public- and private sector investments in HPC facilities and expertise. The Council has begun an initiative called the National Innovation Collaboration Ecosystem (NICE), under the guidance of its HPC Advisory Committee and in partnership with leading organizations, including the University of Southern California's Information Science Institute. The Council envisions NICE as a powerful resource that will allow businesses of all sizes, as well as entrepreneurs and inventors throughout the nation, to access HPC capabilities and expertise (Figure 1). NICE will be an enabling platform for a new generation of public-private partnerships.

We must aggressively advance this initiative. Failure to take aggressive action will inhibit the competitiveness advantages we need to maintain U.S. leadership in the global economy. "By investing in an HPC-based ecosystem," Ms. Wince-Smith concluded, "America can unleash a new era of innovation-driven growth, create new industries and markets, fuel wealth creation and profits, and generate higher-value, higher-paying jobs."

[1] See the Council's *Study of ISVs Serving the High Performance Computing Market*, Parts A & B, available at www.compete.org/hpc.

The New Frontier in High Performance Computing (Keynote Address)

Through partnerships and multidisciplinary collaborations, the Department of Energy's Argonne National Laboratory has been transforming our understanding of complex systems—even in fields that historically have not exploited HPC, such as supply chains, sociology, history and anthropology, according to Dr. Robert Rosner, director of the laboratory.

Dr. Rosner described four examples of collaborations in which the lab's HPC systems and expertise are spurring innovation and competitive advantage:

Nuclear Power Plant Design: Keeping America in the Forefront

In collaboration with Argonne and other DOE National Laboratories (through the Global Nuclear Energy Partnership), the commercial nuclear power industry aims to apply experimentally validated HPC modeling and simulation for rapid prototyping and safety risk assessments of nuclear power plants. The financial stakes for U.S. industry are high, Dr. Rosner explained: "The French, for example, are among the very likely competitors for American companies in this field world-wide, and CEA, the French equivalent of the U.S Department of Energy, is working on very similar problems. The question is: will U.S. industry be positioned to participate effectively in the world-wide revival of nuclear energy?"

In this industry, costs are heavily driven by regulatory expense. The ability to shorten the design cycle, and to demonstrate the safety of complex reactor systems to the satisfaction of regulators, is an enormous cost advantage. A science-based, experimentally validated approach, made possible by HPC modeling and simulation coupled to thorough code validation (as opposed to a program based primarily on phenomenology), makes it possible to understand how complex systems will behave even when they are perturbed far beyond their design points. Establishing a margin-of-safety purely by experimental means is actually risky in the absence of fundamental understanding, and tends to lead to inappropriate risk assessments. Using experimentally validated HPC, it is also possible to do sensitivity analyses to determine which areas of research and development will produce the best return-on-investment. Argonne, in collaboration with other national laboratories and universities, is developing the HPC simulation tools that can do this, and that can be easily adapted to new HPC systems as they are introduced.

Astrophysics: When Going It Alone Won't Work

The University of Chicago, together with Argonne, is using HPC to transform our understanding of how Type 1a supernovae explode, and in the process has gained critical insights into multidisciplinary HPC collaborations. These complex stars are the classic yardsticks for measuring the size and age of the universe, and for understanding other fundamental questions about its makeup. While developing the specialized codes to understand these explosive processes, we gained insights into the importance of multidisciplinary collaboration in complex problem solving.

At first, the University/Argonne collaboration thought that a group of computational physicists or astrophysicists could write the simulation code on their own, with occasional interaction with computer scientists. That approach failed – the collaboration learned that a multidisciplinary team was needed, and that the team members had to work collaboratively from the start. "If you want to hold a team of physicists, computer scientists and applied mathematicians together, you can't afford to have one group go off, have deep thoughts for a while, and then expect them to come back some time later and still engage with the folks that have been waiting around," Dr. Rosner explained. "That simply will not happen. This was not a matter of the physicist telling the computer scientist, 'Could you just do this for us?' This was really an intimately-linked collaboration, right from the start, between the code builders, computer scientists and applied mathematicians, and the physicists and astrophysicists."

History + Anthropology + Sociology: Mesopotamia Goes Multidisciplinary
The University of Chicago's National Science Foundation-funded collaboration with Argonne used HPC to develop new techniques for multidisciplinary study of history. The team chose to study the evolution of the Tigris/Euphrates area over the past 5,000 years—including population dynamics, weather, land use, and other factors. They successfully constructed a simulation of the physical and sociological environment of that ancient era. And, as Dr. Rosner pointed out, "This is a case where you can actually do validation. You can made predictions taking a certain slice of time, and see whether these are borne out by actually going back and doing the archeology. No one would have thought that archeology would become an HPC application, and here it is."

Dr. Rosner also pointed out that sociology is rarely considered an HPC computational science today, except perhaps for statistical analyses performed on very large databases, e.g., data mining. Having worked on a simulation of the Tigris/Euphrates physical and sociological environment, one can imagine HPC simulations and modeling in diverse areas involving millions to hundreds of millions of 'actors,' exploring the consequences of individual societal interactions driven by processes such as advertising (and mass communication in general), kinship relations, economic forces, geography, and politics. This is an area in which the issue of code validation becomes supreme; and it is one of the great challenges in this field to understand how such simulation efforts will deal with the validation problem.

Challenges for Multidisciplinary Use of HPC
One major barrier to the use of HPC that spans disciplines, and has been mentioned often in Council conferences and workshops, is legacy software codes "that simply will not run on modern high performance computers," Dr. Rosner said. "Codes can live for decades, but the hardware rarely – if ever – lasts much more than a few years. The challenge is to engineer codes so that the changes in underlying hardware architecture become transparent to the higher-level software developers and users, so that one is not always starting from scratch when moving legacy codes between different generations of computing hardware. "Software also must be more "friendly" to the range of potential users. "We really have to distinguish in a much more systematic way, from the user point of view, different levels of expertise: "Argonne is partnering to develop simulation tools from scratch that will be able to keep up with changes in hardware platforms.

Driving Competitiveness through Collaborative Computing (Panel)

Public-private partnerships are evolving to help close the gap between the use of HPC in the public sector and in industry. Dr. Peter Freeman, assistant director for the National Science Foundation's Computer and Information Science and Engineering Directorate (CISE), chaired a panel discussion that explored public-private partnerships from the perspectives of a university HPC center, a national laboratory, and a private-sector business. The panelists provided important lessons on the aspects of collaborations that are working well today, and what must still be done to make the partnerships more successful. Dr. Freeman noted, "It is important to find a collaborative path that allows industry and other less-advanced users to gain early access to HPC technology. It is also important to prevent intellectual property rights issues from remaining a barrier to competitiveness."

The Pittsburgh Supercomputing Center (PSC) received high marks in the Council's recent study of industrial partnerships with National Science Foundation (NSF) centers.[2] PSC also has a strong history of outreach and collaboration at the state and local levels.

---

[2] See the *Council on Competitiveness Study of Industrial Partnerships with the National Science Foundation (NSF), 2006*, available at www.compete.org/hpc

Executive Director Beverly Clayton said that PSC has collaborated with industry since its founding 20 years ago as one of the original NSF-funded university HPC centers. In addition to receiving funding from NSF, PSC has also been funded from its inception by the Pennsylvania Department of Community and Economic Development, with the understanding that PSC would provide HPC resources to Pennsylvania companies. Focusing at first on large firms, PSC today also serves small and midsize companies, reaching them through state economic development programs and private sector organizations such as the Pittsburgh High Technology Council.

In her remarks, Ms. Clayton summarized lessons PSC has learned from two perspectives: collaborating with industrial partners, and working with economic development officials.

University HPC centers and national laboratories must help government officials understand the value of HPC for local, state and regional economic development.

Noting that PSC does not have a "complete recipe for success yet," Ms Clayton stressed that HPC centers must work with executives at the top of state economic development agencies as well as staff "in the trenches" to ensure that they understand the value these centers can deliver. Only then can the officials promote these benefits effectively to companies and help link companies to the centers for regional economic gain. It's difficult to get these messages to stick, however, Ms. Clayton said. There is high turnover in the economic development agencies, and different Pennsylvania governors have shown varying levels of interest. "It's a constant process of education. Every four or eight years, we have to do it again," she said.

If they are going to partner successfully with industry, university and national laboratory HPC centers must be highly flexible in helping companies to define the projects.

University HPC centers such as PSC need to be extremely flexible when collaborating with industry, according to Ms. Clayton. "We try to work with companies based upon their needs. If university HPC centers want to be effective national resources and anchors for regional economic development, they must be able to support a broad spectrum of users, from entry-level to experienced, and assist with a wide range of problems," she said. This is particularly true when companies come to the center without a well-defined project or model. HPC centers must be prepared to help companies define the model and even run it on the Center's HPC system.

Industry-university partnerships will not succeed if companies do not appreciate the value of modeling and simulation.

Ms. Clayton stressed that for public-private HPC partnerships to succeed, companies must understand that HPC modeling and simulation is an important ingredient for business success. It is important, therefore, that companies hire people who have embraced this approach as a way of doing business. Ms. Clayton pointed out that companies won't be able to find appropriate talent if our educational system does not do a better job of teaching modeling and simulation as a third approach to scientific discovery, in addition to theory and physical experimentation. This effort needs to begin at the high school level.

Dr. Ian Foster, director of the Computation Institute at Argonne National Laboratory and the University of Chicago, shared Argonne's perspectives on partnering with industry. Industrial partnerships at Argonne cover a broad range of activities, ranging from formal cooperative research and development projects to industry use of (and sometime collaborative development of) Argonne high-performance computing software. Indeed, Argonne software for parallel computing, numerical modeling, and grid computing is used by tens of thousands.

Some companies approach Argonne for access to one of the fastest HPC systems in the world. However, Dr. Foster pointed out that access to Argonne's expertise was the attraction for many others, and that this expertise encompasses "not only high performance computing in the sense of modeling and simulation, but also large-scale data analysis, data mining, and related computational methods that can also have a dramatic impact on a company's bottom line."

Dr. Foster also communicated some important lessons for collaborations and partnerships involving HPC.

National laboratories and their industry partners can enhance their mutual learning by spending time at each others' sites.

Dr. Foster stressed the importance of treating collaborations as opportunities for mutual learning. For example, companies can enhance their in-house HPC expertise by co-locating staff at the sites of their national laboratory partners. In this way, the company representatives can become familiar with new technologies and techniques that they can then incorporate into their business processes. But the most effective staff exchanges work in both directions. "Having staff from the labs spend time in industry is equally valuable. We've done this at various times in our collaboration with Acxiom, for example," explained Dr. Foster. (Read further for more about the Argonne-Axciom collaboration.) This and other industrial partnerships showed Argonne that the problems industry wrestles with are often similar to those that the laboratory scientists are struggling to resolve. In the case of the parallel virtual file system that Argonne and Acxiom collaborated to develop, "clearly we are able to do more together than either of us could have done alone," Dr. Foster said. No money changed hands in this collaboration, but both sides gained considerable value from it.

Service portals could make software from the national laboratories available to much larger communities, but may require investment in support personnel as well as technologies.

There is a transition occurring in some parts of industry toward accessing software as a service, rather than downloading and installing it on company computers.[3] Dr. Foster expanded on this concept, suggesting that HPC Centers such as those at Argonne and PSC could provide software to a much larger community than they serve today by creating portal interfaces to enable easy access and usage. Argonne is doing this with bioinformatics software, for example, by making available substantial data resources backed by large-scale high performance computing that is not visible to the users. As Deborah Wince-Smith outlined in her opening remarks, the Council sees service portals as an integral component of its National Innovation Collaboration Ecosystem (NICE) initiative. Dr. Foster cautioned, though, that it is expensive to produce a portal with user-friendly interfaces. Even then, additional investment may be needed in the human expertise to support people using the service portals and to help them formulate their problems. "You can only go so far in automating the process," Dr. Foster pointed out.

Dr. Foster also echoed Ms. Clayton's observation that the value of modeling and simulation should be taught in our educational systems. He recommended that students spend time in national laboratories to enhance their HPC skills.

Partnerships and portals can expedite knowledge transfer between the national laboratories and industry, to the benefit of both.

Dr. Terry Talley, chief architect for Acxiom Corporation's Products and Infrastructure Technology Organization offered an industrial perspective on public-private partnerships drawing on his firm's experience working with Argonne National Laboratory. Acxiom is a mid-sized company ($1.3

[3] See the *Council on Competitiveness High Performance Computing Software Workshop Report, 2005*; and *Council on Competitiveness Second Annual High Performance Computing Users Conference Report: Accelerating Innovation for Prosperity, 2005*, available at www.compete.org/hpc

billion in revenue) that builds large data warehouses to provide business intelligence and decision support for Fortune 500 companies. “The whole idea,” said Dr. Talley, “is to apply lots and lots of data, and logically integrate all that data and put it into a decision support system, so that you can make good business decisions based upon that data. Because of the volume of that data, we were very interested in trying to exploit high performance computing.”

According to Dr. Talley, the company began looking for a way to boost its competitiveness by accelerating the process of building data warehouses. But as Acxiom enhanced its computational capabilities through grid computing, it quickly ran into input/output challenges. That's when the company became aware of Argonne National Laboratory's work with the Parallel Virtual File System (PVFS). Acxiom discovered that PVFS was a good match for the types of input/output problems it was tackling, where data had to flow in parallel through a large number of processors. Acxiom entered into a co-development partnership with Argonne to access and enhance this open source software together. Access to Argonne's large HPC systems was also important for the co-development.

Although Acxiom approached Argonne for its unique software expertise, the knowledge flow throughout their collaboration has not been one-way. Although people usually talk about knowledge transfer from research labs to industry, "It goes the other way around as well," Dr. Talley explained. “We spend a lot of time explaining to our partners in the research community what our business problems are, and those are often foreign to them. It's been surprising to me how much time we spend actually doing this."

Like Dr. Foster, he also elaborated on the potential benefits of portals to HPC resources. “I think the idea of having portals is very important because it allows us to identify resources that we could use and adapt to our particular problem.” Dr. Talley was also enthusiastic about the advent of service portals that would allow Acxiom to "demonstrate a proof of concept, so we can make an intelligent investment. Those are all really important things."

HPC collaborations have the potential to transform companies and entire industries.

In summing up Acxiom’s partnership experience, Dr. Talley outlined three tangible competitive benefits the company has received:

“First, Acxiom has been able to reduce the transactional cost for its computing in a significant way by accessing Argonne’s HPC systems. As a mid-sized company, we simply could not afford to purchase commercial machines with that much horsepower. So on the cost side, we're more competitive,” Dr. Talley explained. Second, he said, thanks to the Argonne collaboration, Acxiom now can deal with data volumes and transformation complexity “that our competitors simply can't.” And finally, as a result of its work on the grid and the parallel virtual file systems with Argonne, Acxiom has entered into a relationship with EMC to jointly develop a commercial product.

DreamWorks Animation is also partnering with a national laboratory. Through the Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, the company successfully competed for a large block of time on a high performance computer at Oak Ridge National Laboratory. DreamWorks Animation's goal is to use this resource to transform the way animated films are created.

In the computer graphics (CG) animation industry, the pace of change and the competitive pressures have increased dramatically over the last decade. When Pixar released *Toy Story*, the first CG animated film, in 1995, it created an entirely new art form. Twenty films, approximately two per year, were released over the next decade. Now, the CG animated film industry averages one new release per month. This has created more competitive pressures for every company in this market.

Ed Leonard, chief technology officer for DreamWorks Animation, explained that the most computationally intensive part of creating an animated film is rendering, the important process of applying the light and texture to otherwise finished scenes. This process gives animated films exceptional visual interest. “We have to tell great stories,” explained Mr. Leonard, “but we need to tell them in a way that keeps the audience's interest. It's no longer possible to keep people's attention with a painted background. People want to see trees that move, and leaves that move.”

Today, the final step of rendering has to be done overnight on "batch farms" of computers. Through Department of Energy's INCITE grant, DreamWorks Animation is exploring possibilities for a radical advance that would transform the animated film industry—eliminating the need for overnight batch processing and making rendering part of the interactive, daytime “artist-in-the-loop” process of creating an animated film. Advances are especially important, because rendering consumes ever-increasing amounts of computational power. *Shrek 1* (2001) required five million render hours; *Shrek 2* (2004) took 10 million render hours, and *Over the Hedge* (2006) consumed 15 million render hours.

DreamWorks’ engagement with the INCITE program offered lessons applicable far beyond the animated film industry.

Partnerships with universities and/or national laboratories offer industry the opportunity to explore competitively important problems that are orders of magnitude beyond what they are currently solving in their production environment.

Market pressures demand that DreamWorks utilize its in-house system full time for the production of its films. Resource constraints do not permit the kind of forward-looking experimentation that drives competitive breakthroughs. Through its partnership with Oak Ridge via the INCITE program, DreamWorks Animation will have access to one of the most powerful high performance computers in the world to try out its cutting-edge ideas. "Our motivation was to figure out how we could gain not one order-of-magnitude improvement, but three or four or five, and what that would do to our process,” explained Mr. Leonard. “It became very clear to us that it wasn't about making something you already knew how to do go faster. It was about creating entirely new ways of making film, and what that would do to the creative process."

By partnering with the national laboratories, industry can get a “crystal ball” look at HPC systems several years before they are widely accessible in the commercial market, accelerating its ability to prepare for their use.

National security demands propel the national laboratories to invest in leading-edge HPC systems before they are widely available and affordable in the commercial market. When industrial firms are able to access these systems through partnership programs like INCITE, they gain a “crystal ball” look into their technical future and a head start in preparing their internal processes and programs to embrace it. Early access to cutting-edge HPC resources allows DreamWorks to experiment with new techniques to create imagery and characters, in worlds and environments that “actually separate us from a pack that's crowded,” explained Mr. Leonard. This access also helps the company better understand what kind of HPC systems it will need to maintain market leadership in the future. “Our intention is to learn what's next for us in two or three years, not necessarily what we need today,” Mr. Leonard emphasized. “It's not optional for us to stay on the leading edge—it's essential."

NNSA, NSF Models for Public-Private Sector Partnerships

Suzy Tichenor, vice president and director of the Council's High Performance Computing project, summarized the findings of two pioneering new studies produced by the Council in collaboration with market research firm IDC. The studies document the experiences of companies that have participated in HPC partnerships with university-based centers affiliated with (1) the Department of

Energy's NNSA/ASC Academic Strategic Alliance Program, and (2) the National Science Foundation. Both studies[4,5] are available at http://www.compete.org/hpc.

Industry's partnerships with the NNSA funded University HPC Centers were overwhelming successful, but collaboration tools may still be immature and cultural differences between businesses and universities were an impediment in some cases.

Ms. Tichenor reported that the findings of the NNSA study were based on discussions with 12 aerospace, automotive, energy and software firms engaged in HPC partnerships with the NNSA Alliance Centers at Stanford University, the University of Illinois, and the University of Utah. The companies reported that the collaborations were overwhelmingly successful. A large majority of the firms said the partnerships had met their expectations, and that they would be willing to partner again with the same Center. One-third of the firms said they had achieved an important breakthrough or discovered something totally new. This is very encouraging, given the Council's belief that America needs to become even more innovative to remain at the head of the competitiveness pack in global markets. For some of the companies, learning new problem-solving approaches proved to be transformational. And half of the companies reported solving a specific problem, bringing a product to market faster, achieving a cost reduction, or increasing profitability as a direct result of the partnerships.

The biggest impediment to progress reported by the companies was the slower-than-desired pace of some of the projects. They generally attributed this to cultural differences between businesses and universities. Additionally, the greatest progress occurred when project participants worked together "face-to-face", suggesting that collaboration tools may still be immature.

The NSF industrial partnerships produced many breakthroughs and created more demand for HPC tools, but the NSF centers' HPC resources are already oversubscribed.

This study, according to Ms. Tichenor, queried 38 companies about their partnerships with three NSF funded University HPC Centers: the Pittsburgh Supercomputing Center, the San Diego Supercomputing Center, and the University of Illinois. Many of the firms had prior HPC experience. Three-quarters of the companies engaged in the partnerships to advance strategically important work. Some wanted expertise to help solve an immediate problem.

The commercial firms were highly positive about their experiences working with the experts and the large HPC systems at these facilities. All reported meeting their objectives, including advancing their R&D efforts. An impressive 88% said they had solved the specific problem they set out to conquer. Forty percent of the firms were able to assign a dollar value to the outcome of the partnership, ranging from $100,000 to $57 million. More than half of the companies achieved a breakthough in their existing work, or discovered something entirely new.

As in the NNSA study, the emphasis on face-to-face collaboration raised questions about the adequacy of today's advanced collaboration tools. And cultural differences between the commercial businesses and the university-based Centers were noted as the largest impediment to project progress.

Study participants also agreed that the NSF and the Centers have an opportunity and a responsibility to market the Centers' HPC resources more aggressively, so that more U.S. businesses can take advantage of them. Today, as Dr. Freeman pointed out, the Centers are oversubscribed. The NSF

[4] See *Council on Competitiveness Study of Industrial Partnerships with the National Science Foundation (NSF), 2006* available at www.compete.org/hpc

[5] See *Council on Competitiveness Study of Industrial Partnerships with the U.S. Department of Energy/NNSA, 2006*, available at www.compete.org/hpc

and the Centers will therefore need to make a strategic decision on how much time they can devote of industry partnerships while still fulfilling their mission to support the university research community.

Interestingly, 38% of the firms in the NSF study reported upgrading their HPC systems, or purchasing new ones, following the NSF partnerships. Because of the partnerships, they were able to see the bottom-line value of using larger systems.

The industrial partnerships programs of the NSF and NNSA, as well as the Department of Energy's INCITE program (see below), address two key challenges stalling more aggressive use of HPC by U.S. firms: the need for state-of-the-art HPC expertise, and the need for access to large-scale HPC systems. The Council will build on these successful partnership models in the development of the National Innovation Collaboration Ecosystem (NICE).

Public-Private Partnerships Through the Department of Energy's INCITE and SciDAC Programs

Dr. Raymond Orbach, director of the Department of Energy Office of Science and recently named the nation's first Under Secretary of Science, reinforced earlier conference speakers on the importance of public-private partnerships. At his luncheon keynote address, he praised the Council's initiatives to advance the collaborative use of HPC to boost innovative and competitive gain, and then discussed two Department of Energy programs that enable industry to access some of the nation's most powerful HPC systems.

INCITE: Fostering Innovation through Access to High-End Computers

The Department's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program gives American industry, universities and laboratories access to some of the nation's most powerful high-end computers for scientific and engineering discovery. This successful program, now in its fourth year, grants large blocks of computing time to proposals with great potential for breakthrough discoveries. For fiscal year 2006, 18.2 million hours of computing time were made available on five supercomputers at four DOE national laboratories: Argonne National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory and Pacific Northwest National Laboratory. And for the first time, researchers from four companies were among the 15 awards that passed DOE’s strict peer review process: Boeing, DreamWorks Animation, General Atomics, and Pratt & Whitney. In fiscal year 2007, DOE will dramatically expand the amount of computing time it will make available to 95 million hours.

SciDac2: $60 Million A Year For Collaborations

The Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) program complements the INCITE program by bringing together computer scientists, mathematicians, physicists, chemists, engineers and others in a collaborative framework that can adapt physical problems to powerful high-end computers. To support the winners of the SciDAC2 competition, the Department of Energy will invest nearly $60 million a year in 30 projects involving 70 institutions and hundreds of researchers and students. The department will provide support for three to five years for the projects, which span a wide range of scientific application areas, including fusion energy, turbulence, climate change, chemistry, nucleosynthesis, groundwater transport of contaminants, computational biology and materials science.

Dr. Orbach described three related initiatives within the SciDAC program that are designed to foster multidisciplinary research:

- A unique feature of the SciDAC program will be the integrated set of nine centers for enabling technologies. These centers will specialize in applied mathematics, computer science, distributed computing or visualization, and will be closely tied to specific science applications and meeting the challenges of petascale computing.

- SciDAC will create four new institutes that tap into some of the nation's leading intellectual resources, in order to help educate and train future generations of computational scientists. Through hands-on workshops and tutorials, the institutes will also help researchers to learn from SciDAC teams how to prepare their applications to take advantage of the increasing capabilities at supercomputing centers around the country.

- Finally, the new SciDAC Outreach Center will act as a service portal to provide "one-stop shopping" for support services and outreach to scientists outside of SciDAC who are located in universities, national labs, and industry. This center is a pilot program to leverage existing support services at NERSC, and to gather data about the breadth and specific needs of the emerging petascale community. The National Science Foundation and The Department of Energy's National Nuclear Security Administration (NNSA) have joined the Department's Office of Science in this partnership.

Driving HPC through the Supply Chain (Panel)

Microsoft Corporate Vice President Marshall Phelps moderated a panel discussion of industry executives representing both suppliers and purchasers. They explored the challenges and benefits of using corporate supply chain partnerships to drive HPC to a broader user base, as well as the role of HPC in optimizing a company's supply chain for added profitability and competitive gain.

The panel included Gary Abyad, president of Clopay Plastic Products Company; Tom Lange, director of corporate R&D modeling and simulation for The Procter & Gamble Company; Dr. Jayant S. Sabnis, chief engineer for systems analysis & aerodynamics, Pratt & Whitney; and Nancy Stewart, senior vice president and chief technology officer, Wal-Mart Stores, Inc.

In a nutshell, supply chain management means optimizing the flow of materials, information and money all the way from a company's suppliers to its customers, including what happens within the company itself. Efficient management of supply chains is increasingly important for competitiveness and profitability. There are strong competitive pressures to obtain parts quickly for assembly and to deliver finished goods quickly to outlets. Modeling complex supply chains for maximum efficiency and cost control, particularly within large companies, often presents enormous data management and data analysis challenges. Increasingly, high performance computing is required for the task, and supply chain management is emerging as a critical HPC application. Because of the close relationships that many companies within a supply chain maintain with each other, the supply chain could be a vehicle for expanding modeling and simulation with HPC to companies that are not currently using this technology.

Expanding HPC usage through supply chain partnerships

Panelists expressed differing views on the practicality and suitability of using supply chain relationships to introduce companies to high performance computing. Often the difference in opinion reflected the depth, or lack of depth, in supplier options.

Pratt & Whitney: Actively Helping Suppliers Learn To Use HPC

According to Dr. Sabnis, Pratt & Whitney has long used HPC for the modeling and simulation of jet aircraft engines. As he explained, "whether they recognize it or not, everybody utilizes modeling and simulation. Your models are either analog, also known as prototypes, or you build your models on a computer." Pratt & Whitney builds computational models because the process is faster and lower cost than building physical prototypes. "It's all about taking cost out, and meeting the schedule," Dr. Sabnis emphasized. "And if the supply base does not meet the specifications and if the things they are supposed to deliver don't arrive on time, we have to ultimately bear the cost." Dr. Sabnis also pointed out that in the jet engine business; there is not the depth of suppliers that might be available to other businesses that are designing/delivering commodity products. "I have to ensure that whomever I have picked actually succeeds."

For these reasons, Pratt & Whitney introduces key suppliers to HPC. "They don't always come willingly," Dr. Sabnis pointed out, "because this is something they have not done in the past." And so Pratt & Whitney is careful to first assess the "technology readiness" of a supplier. "We take a lot of time to drill deep, and discover what the capability of the supplier is to understand what questions they should be asking, and what boundary conditions they should be posing on a specific analysis," Dr. Sabnis explained. "And if they don't have the capability, we will actually work with them to make sure they come up to speed on that."

Pratt & Whitney's hands-on approach has been successful. As their suppliers see the benefits that Pratt & Whitney accrues from modeling and simulation with HPC, they are adopting it themselves.

Procter & Gamble: Letting Market Forces Decide

Consumer products giant Procter & Gamble provided insight as both a large supplier and customer. It is one of the largest suppliers of finished product to the demanding retail giant Wal-Mart. And like Pratt & Whitney, P&G relies heavily on HPC for designing and testing its products in order to meet its customers' requirements. "It's part of our innovation equation," explained Mr. Lange. "Every dime we spend on computing and software and hardware, and the people to do it, is a dime we don't spend on prototype molds, prototype equipment, full-scale experiments and physical tests."

However, despite the benefits it derives from using HPC, P&G normally does not introduce its own suppliers to modeling and simulation with this technology as Pratt & Whitney does. "We will respond to a mentoring or partnership requests when asked," Mr. Lange said. But, he pointed out "We don't get asked very much." Instead, P&G relies on market forces to push suppliers to adopt the tools they need to meet P&G requirements. If HPC can help P&G's suppliers "reduce inefficiency, make higher quality product, reduce weight, and reduce waste, we strongly applaud and encourage that," explained Mr. Lange. But "we're going to rely on standard market forces to make that happen, or not."

In Procter & Gamble's world, there are some important barriers inhibiting the use of HPC by suppliers. For example, software licensing restrictions often prevent P&G from allowing a supplier to use its HPC system. Another barrier is a generally lower level of engineering ability among P&G's suppliers than among, for example, auto or aerospace-industry suppliers. Related to this is a shortage of user-friendly middleware. "Not enough people are producing higher-quality middleware to get this scientific capability out into industry for things that are done every day. We need middleware to get some of the routine analysis automated, so people who are not experts in finite element analysis and computational fluid dynamics can run it," Mr. Lange explained. Finally, Mr. Lange pointed out that most of P&G's suppliers simply do not have the scale to make what might appear to be a risky investment in HPC. "It's taken Procter & Gamble, with all of our scale and all of our focus, 20 years to learn how to migrate work that we used to do physically to a virtual environment."

Wal-Mart: Assisting Suppliers on an As-Needed Basis

Ms. Stewart said Wal-Mart, the world's #1 retailer with about 8,000 stores, also does not require its suppliers such as Procter & Gamble to use HPC. But Wal-Mart expects its suppliers to use HPC if that is what is needed to produce the highest quality, lowest cost product. "To the degree that our suppliers, like P&G and others, can demonstrate to us that they are producing high value at low cost, that's what's germane. That's the value of HPC for us and that's the value that it also brings to our customers who are looking for the highest value at the lowest cost."

However, like Pratt & Whitney, Wal-Mart occasionally finds that only one or two suppliers can meet its requirements. "Sometimes we really have only one or two suppliers that can meet the need because of the volumes we require," Ms. Stewart explained. In these instances, Wal-Mart will assist its suppliers as needed with the HPC systems and expertise it uses in-house to manage its global supply chain and network of stores. Ms. Stewart went on to relate a recent incident in which a supplier was unable to fulfill a commitment. Wal-Mart stepped in and used its HPC technology and techniques to make the needed improvements. "We did the analysis," Ms. Stewart said. "We did the work such that our supplier could get the product to market on time and at the prescribed cost metric. But we used our systems to accomplish that. And he was made successful and we were made successful by getting the work done."

Clopay Plastics Products: Not Seeing the Need Yet

Mr. Abyad provided the perspective of a medium-sized company that must meet the stringent demands of Wal-Mart and Procter & Gamble, but does so without using HPC. Despite the fact that his firm, Clopay Plastic Products Company, is a supplier to two committed users of HPC, Clopay isn't convinced yet of the value of this technology. Clopay is therefore unwilling to make what it deems to be a risky investment in HPC.

Echoing Mr. Lange's comments about the importance of "scale," he noted that as a division of a public company (Griffon Corporation), "we don't have the luxury of a ten-year investment in this technology for the sake of a potential payout." Clopay has a fixed resource base and finds that this is "well spent, on a daily basis, trying to come up with the very next iteration of innovation by our traditional (i.e., experimental) means." Siphoning off dollars to invest in HPC "has not made the radar screen." In addition, although Clopay feels the pressure from its customers to reduce innovation cycle time and cost, Mr. Abyad is not aware that any of Clopay's competitors are using HPC and gaining a competitive edge.

What would it take to make HPC flicker on Clopay's radar? Mr. Abyad was candid in his reply. "We don't use HPC because we don't understand the value of it and we don't understand how to use it. We don't have the tools and we don't have the skills in our company. We need to be introduced and led with user-friendly tools that can produce results." And although he thinks that "larger partners can mentor smaller firms like ours in the use of HPC," he acknowledged there are competitive issues. "Using these tools provides a competitive advantage, so P&G would need to decide how this would affect their own competitive position."

Using HPC to model supply chains for optimal efficiency and cost

In addition to discussing the use of supply chain relationships as a venue to expand HPC usage, Ms. Stewart pointed out an emerging and increasingly important industrial HPC application: using HPC to manage and optimize the supply chain process itself for maximum efficiency and cost control. "The value we derive from using HPC for cost reduction is what drives us," Ms. Stewart explained.

On a typical day, Wal-Mart sees 27 million customers and processes 500,000 transactions worth $2 billion of revenue. (On the day after Thanksgiving, Wal-Mart's busiest day of the year, the company processes about a billion transactions worth $20 billion in revenue). The company ships about 740,000 items daily and Wal-Mart Supercenters stock about 500,000 products each. Suppliers compete fiercely for shelf space.

Every day between 3:00 and 5:00 a.m., the company runs very large models on high performance computers to determine what is selling well in each store. That information is then sent to Wal-Mart suppliers so that they know what to stock in each store. It is also sent to every store so that shelves can then be "reformated" to meet customer needs more appropriately. "If you really understand your information or your data, then you can do your predictive analysis," Ms. Stewart explained.

In addition to using HPC for shelf space determinations, store planning and resource planning, the company also uses HPC for operational "ergonomics." From its headquarters in Bentonville, Arkansas, Wal-Mart services all of its stores worldwide, right down to turning on the lights in the stores. "Within a day, I basically process a petabyte of data," Ms. Stewart stressed.

Ms. Stewart emphasized that Wal-Mart "couldn't do these kinds of things without this level of technology value. That's why we made the investment in HPC. We can see the return and it's helped make us that much more efficient."

HPC Reveals a Major New Oil Trend For Chevron and Its Partners

After spotting Chevron Chief Technology Officer Dr. Donald Paul in the audience, Conference co-MC Dr. David Shaw asked him to come forward and comment on the recent, highly publicized discovery, by

Chevron and two of its partners, of a new field in Gulf of Mexico deepwater that could yield 3-15 billion barrels of oil—boosting U.S. reserves by up to half.

Dr. Paul said HPC was a crucial feature, “not just an add-on,” for enabling this important discovery. HPC has been used for seismic processing for many years, but Chevron’s "Jack-2" reservoir was at the very edge of current seismic imaging capability. Imaging at the scale of this project was unprecedented, with data sets up to a quadrillion (10^15) points. Processing such vast data sets was impossible until the past few years brought advances in HPC capabilities and visualization technologies.

The features of the newly discovered reservoir were invisible until recently, because of a huge canopy of salt that is sometimes miles thick, and geologists were skeptical about the amount of potential oil in that region. But with high performance computing, what was invisible became clear. "The HPC systems become significantly faster, so you can see more, adjust the algorithms, and finally image what you're looking for. This opens up an enormous exploratory area 300 miles long and 100 miles wide," said Dr. Paul.

Once HPC permitted Chevron to “see” the possibilities, the company had the confidence to proceed with the enormously expensive process of drilling a test well. HPC was used again for the even larger challenge of modeling in real time what the drilling process might be like. Specialized ships costing up to $1 billion each were needed to drill through 7,000 feet of water and 20,000 feet of underlying rock. The steel drillstrings were five miles long (8 kilometers).

The next stage is to use HPC to model these reservoirs so that decisions can be made about how best to develop them. This will involve simulations with tens of millions of cells, eventually some of the largest models ever. Again, the modeling will not be done in the lab, but "on the front line of production work.”

“It would not have been possible to have had this exploration success five years ago,” Dr. Paul summarized. “We just didn’t have the horsepower to do the computations to apply in practice this kind of imaging, drilling, and reservoir modeling technology. HPC was absolutely critical.”

Bridging The Gaps With A Collaborative Ecosystem

Dr. Shaw concluded that exploiting the full potential of HPC to advance U.S. competitiveness is a significant challenge, but "we're making enormous progress already," thanks to the Council's work and a pattern of increasing participation and collaboration. Referring to the Council's HPC Initiative, Dr. Shaw said, "This isn't just an independent research project. It's something we wanted to be driven by the various stakeholders, and that has really happened. The Council’s HPC Initiative is making a very significant contribution to high performance computing, on the one hand, and to U.S. competitiveness, on the other."

While the Council continues to do HPC-related research and evaluation, it will also move forward on important efforts that have grown out of the HPC Initiative, including the National Innovation Collaborative Ecosystem. "This is exactly what we need," Dr. Shaw said. "It's a way to bridge the gaps and get people working together. NICE is going to be a very exciting program, with high risk and very high potential returns." He told conference attendees that the Council's HPC Advisory Committee "will continue to rely on your help and your insights as we move ahead."

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Council on Competitiveness Study of U.S. Industrial HPC Users

Sponsored by: Defense Advanced Research Projects Agency

Earl Joseph, Ph.D. Addison Snell
Christopher G. Willard, Ph.D.
July 2004

IDC OPINION AND OBSERVATIONS

The U.S. industrial market for high-performance computing (HPC) technical servers has been through a number of changes, evolutions, and revolutions over the past few years. In many ways, it is like going through a physical phase change from a liquid to a solid state, and things have not yet fully solidified. There are many clear trends in the U.S. industrial market today:

- Price and price/performance have both provided new possibilities and driven buyers to acquire different types of HPC servers.
- Clusters have proven themselves as capable servers to handle a sizable portion of the HPC workload.
- Industrial users are still very engaged and excited about applying HPC to help their organizations become more competitive and ultimately more successful.
- Industrial users are less interested in investing in programmers and computer scientists and have refocused their investments and people toward work directly in their business areas.
- Higher-performance computers are desired, but most sites cannot afford to purchase the fastest computers today.

U.S. industrial users/buyers really want and need faster computers that fit their budgets and that don't require specialized programming skills.

The U.S. industrial sites interviewed clearly see HPC as fundamental to the business, as the following quote illustrates: "Our business model could not exist without HPC." These sites are still strategically using HPC and are investigating new ways to apply HPC.

Recent Market Events

2003 was a good year for HPC technical servers, as the technology reversed the economic downturn from the previous two years with a healthy 12% growth in sales revenues. The expansion in the market is taking place at lower price points due to a combination of factors, including the tight economy placing pressure on budgets, the processor improvements provided by Moore's law, the usability improvements in clusters, and the lack of strong high-end product offerings. Hence the high-end capability market segment actually declined by 24%, while the lowest-end departmental segment grew by 35%. We expect to see the movement to lower-priced platforms continue over the next five years.

Definitions

HPC

In this document, the term *HPC* is used in the same way as the terms *HEC*, *HPTC*, and *high-end computing*. We are referring to computer servers used to solve computational or highly data-intensive problems. Our definition requires large scientific/engineering/economic problems.

Industrial Sector

This is a study of nongovernment and nonuniversity sites/organizations that use HPC computers in their businesses. The terms *industrial sites*, *commercial sites*, and *business sites* are all used in the same manner to represent the overall industrial sector.

Capability-Class and Capacity-Class HPC Computers

IDC defines capability-class computers as systems purchased primarily to tackle the largest, most complex single problems. Capability-class HPC systems are generally priced at $2 million to $4 million or more, with costs occasionally approaching or even exceeding $100 million. Traditional symmetric multiprocessor (SMP) technical servers and supercomputers of sufficient size fall into this category, but large-scale clusters also qualify as capability-class systems if they are purchased primarily to address large single problems.

Capacity-class systems are purchased primarily to solve many small and medium-sized problems. Capacity-class HPC systems may also be priced at more than $1 million and may include any category of HPC computer. IDC further divides capacity-class computers/servers by price band:

- **Enterprise:** $1 million or higher
- **Division:** between $250,000 and $1 million
- **Department:** below $250,000

The primary purchasing rationale differentiates capability- and capacity-class systems.

Clusters

Clusters are considered capability systems when they are used for the most challenging problems (e.g., when used for "traditional" capability-type problems). These clusters are typically very large with an average size of more than 300 nodes.

EXECUTIVE SUMMARY

This study was commissioned by the Council on Competitiveness (COC) and sponsored by the Defense Advanced Research Projects Agency to explore the usage and impact of high-performance computing (HPC) resources in industry and other business sectors — including currently available HPC computers and potential future computers assumed to be dramatically faster and easier to use. The study asked about both capacity-class computers, purchased primarily to address many small and medium-sized problems, and capability-class computers, purchased mainly to tackle the largest, most daunting individual problems.

The 33 participants in this study are seasoned private-sector chief technology officers (CTOs), chief information officers (CIOs), and production and research managers representing a wide range of business segments that employ HPC today — from leading aerospace, automotive, petroleum, electronics, pharmaceutical, life sciences, and software companies to financial services, transportation logistics, and entertainment firms.

Major Findings

High-Performance Computing Is Essential to Business Survival

High-performance computing is not only a key tool to increasing competitiveness, it is also a tool that is essential to business survival. Nearly 100% of the respondents indicated that HPC tools are indispensable, stating that they would not exist as a viable business without them or that they simply could not compete effectively. A majority (70%) of the respondents indicated that HPC is so important that their organizations could not function without it.

Nearly 100% of the respondents indicated that HPC tools are indispensable.

Typical comments include:

"There is no other way for us to complete our work. We would not exist."

"There is no other way for us to complete our work. We would not exist."

"The time to market would prohibit our business from existing."

"We would not be able to stay technologically ahead of other competing nations."

The number 1 reason given for purchasing high-end computers is their unique ability to run very large and very complex computational problems that companies must successfully address to maintain their competitive advantage. In addition to running these large-scale problems, the majority of respondents are also able to harness the computer power to run a larger number of smaller-scale, important problems than they were able to run in the past.

Companies Are Realizing a Range of Financial and Business Benefits from Using HPC

Companies described a range of impressive competitiveness benefits realized from using high-performance computing. Approximately one-quarter of the respondents were able to quantify the ROI to their organizations, in some cases in the millions of dollars. Strategic competitive benefits included gains such as shortened product development cycles and faster time to market (in some cases more than 50% faster), not to mention the resultant reduced costs, all of which can improve a company's bottom line.

"It has been a continuous stream of revenue to our bottom line, giving us the ability to look into other development areas."

"It drives innovation, R&D effectiveness, and productivity."

"It drives innovation, R&D effectiveness, and productivity."

Companies Are Failing to Use HPC as Aggressively as Possible

Despite the acknowledged importance of high-performance computing to business competitiveness, a majority of respondents acknowledged that they are not using HPC as aggressively as possible. Two-thirds of the respondents indicated that they have important problems that they simply can't solve today. The remaining third said that they need more powerful systems to achieve more effective solutions. Reasons for both vary. In some cases, systems with the needed capability are on the market but companies face obstacles in owning or accessing them or in using them to their fullest capability. These barriers are discussed below. In other cases, the systems required simply don't exist.

Examples of current unsolved problems include modeling block engine assembly in full detail, simulating vehicle rollover, real-time processing of data from remote sensors, protein folding, and coordinating databases across tens of thousands of servers.

Business and Technical Barriers Are Inhibiting the Use of Supercomputing

The largest single factor preventing more aggressive use of HPC is the lack of computational scientists.

Respondents noted a range of reasons that HPC is not used more aggressively. The largest single factor is the lack of computational scientists — human experts (internal or external) who can apply HPC tools to the problems in question — and the budget to hire them. In most cases, the concern was the lack of resources to hire people, but in a few cases, it was the lack of available talent in the marketplace. Closely related is the ease-of-use issue; most industrial sites require software compatibility in their HPC servers and the cost to change or rewrite software is frequently seen as prohibitive.

Despite the often proven returns from using high-performance computing, respondents noted that upper management often does not appreciate the value of HPC hardware and software tools. As a result, HPC is often viewed as a cost instead of an investment, and many sites find it difficult to obtain internal funding to acquire additional HPC resources. More than half of the respondents expect their budgets for all HPC tools will decline (43%) or remain the same (17%) over the next two years.

Companies Don't Have the HPC Tools They Want and Need

When asked if there are currently available HPC tools they would like to own or access, a majority of the respondents answered in the affirmative. Relatively even numbers of respondents pointed to currently available software and hardware tools they would like to own or access. However, 31.6% stated that there are either hardware or software tools missing in the market today, and 21% said that they need hardware systems that are more powerful than any available on the market today.

Most Companies Do Not Rely on Remote Access to HPC

When respondents were questioned about their methods of accessing HPC resources, most responded that they use onsite purchased or leased HPC systems instead of accessing them remotely at partner or external provider sites. And most do not expect to outsource their most complex (and therefore most competitively sensitive) problems in the future. Security is an important inhibiting factor for some companies.

Dramatically More Powerful and Easier-to-Use-Computers Would Deliver Strategic, Competitive Benefits

When respondents were asked what they could accomplish with systems 100 times more powerful and/or 10 times easier to use, their replies again reflected the strategic importance of HPC to competitiveness. They saw opportunities to simulate larger, more accurate models and tackle completely new problems that they cannot address today, resulting in the ability to produce higher quality products, achieve faster time to market, and improve their financial performance.

When asked what could be accomplished if the "ease-of-use" barrier were addressed with systems that are 10 times easier to program, respondents overwhelmingly indicated that they could develop more powerful applications and fundamentally rewrite their current codes. Not surprisingly, they also indicated that they could shorten design cycles and time to market, a natural by-product of better applications. In addition, more easily programmable systems would enable a wider universe of researchers, scientists, inventors, designers, manufacturers, and mathematicians to use high-performance computing to solve their problems, extending the benefits of these systems more broadly across the private sector for increased industrial and national competitiveness.

"We could test two-generations-out models that we are researching today."

"It would increase revenues for the company and market share."

"We would look to rewrite the entire science underlying the current technology and methodology we are using."

"It would make these tools available to a much wider array of scientists who have good ideas but may not have programming skills."

"It would make these tools available to a much wider array of scientists who have good ideas but may not have programming skills."

Dramatically More Powerful and Easier-to-Use-Computers Could Add Billions to the Bottom Line

Although not all respondents were able to quantify the potential benefits from access to more powerful and easier-to-use systems, those who could suggested bottom-line improvements from tens of millions to billions of dollars, an enormous increase over the positive financial benefits users are already achieving today.

"We save $1 billion from a faster product cycle."

"We save $1 billion from a faster product cycle."

"I can't release [the amount], but it is in the billions a year."

Methodology

This study was undertaken on behalf of the Council on Competitiveness and sponsored by the Defense Advanced Research Projects Agency to shed light on the factors that encourage or inhibit the use of HPC by users in the private sector to solve their current research and business challenges as well as the future applications these users foresee. The study is based on extensive interviews with 33 private sector CTOs, CIOs, and production and research managers representing a wide range of business segments that employ HPC. Appendix 1 contains a description and profiles of the sites interviewed. Appendix 2 contains the set of questions used in the interviews.

Respondents needed to be the primary buyers of HPC, or at least be actively involved in the selection process, and have:

- An excellent understanding of high-performance computing technology and its current application within their organization
- A vision for the problems that the organization should be solving in order to remain competitive but an inability to solve these problems today due to insufficient computing capability
- The ability to describe the criteria used and process followed by the organization when considering an investment in high-performance computing
- An active role in developing and presenting the bottom-line justification for this substantial investment
- The capacity to identify the factors that encourage or inhibit the use of high-performance computing tools within the organization or across the sector

Study Limitations

While IDC always aims to provide an accurate, comprehensive view of the subject being studied, certain limitations inevitably affect the results. We believe that the group of 33 private sector officials we surveyed is large and diverse enough to represent important market conditions and trends, but it would be presumptuous to claim that there are not others whose situations differ in certain respects from any in this group. Also, consistent with its purpose, this study is deliberately United States–centric and does not claim to fully mirror the tendencies of users in other parts of the world. Finally, with a group size of 33 respondents, some less popular options for responding to questions are thinly represented, occasionally with only one or two responses. IDC has tried to exercise extreme caution in generalizing from such results and cautions readers to do the same.

SITUATION OVERVIEW

Criteria Used to Justify Investment in HPC

Q4. What are the primary reasons why your organization uses HPC computers?

The primary reasons given for using high-performance computers (see Table 1) mirror the business sectors of the respondents. Independent software vendors (ISVs), for example, are in the business of software development and testing, while automotive companies employ HPC computers mainly to help with the design, manufacturing, and testing of new vehicle platforms. Simulation and visualization, on the other hand, are activities common to virtually all HPC usage, whether in industry, government, or academic settings.

TABLE 1

Primary Reasons Why Organizations Use HPC Computers

Category	Number of Responses	Overall Percentage
Software development and testing (to be sold to others)	7	21
Simulation (high end)	6	18
Automotive design, manufacturing, and testing	4	12
Oil exploration/seismic processing	3	9
Microprocessor/circuit design	3	9
Visualization (high end)	2	6
Pharmaceutical research/drug discovery	2	6
Electronic data automation	2	6
Web site development	1	3
(Trucking) fleet management	1	3
Business problem-solving	1	3
Price/performance and scalability	1	3
Total	33	100

n = 33

Note: More than one response per site was allowed.

Source: IDC, 2004

Actual responses to this survey question underscore the importance of HPC to the industries and businesses in question, as the following examples demonstrate:

"We will launch 25 new [automotive] products in the next three years, and they all require major use of HPC."

"Designing and verifying a world-class microprocessor requires significant CPU power."

"Oil exploration is an extremely compute-intensive process for reservoir simulation, seismic simulation, and seismic processing."

"The modeling problems are intractable for smaller systems. We have many users who share this system, so this provides a cost-effective solution to a large, diverse R&D community."

"HPC has been critical to our business model."

"Designing and verifying a world-class microprocessor requires significant CPU power."

"The modeling problems are intractable for smaller systems. We have many users who share this system, so this provides a cost-effective solution to a large, diverse R&D community."

Q5: How does HPC impact your organization's primary goals and/or mission?

As Figure 1 illustrates, a decisive majority (70%) of the respondents indicated that HPC is so important that their organizations could not function without it. Another 27% said that HPC contributes significantly to what they do. Only one respondent rated HPC as merely "moderately important."

70% of the respondents indicated that HPC is so important that their organizations could not function without it.

FIGURE 1

Impact of HPC on Organizations' Primary Goals and/or Mission

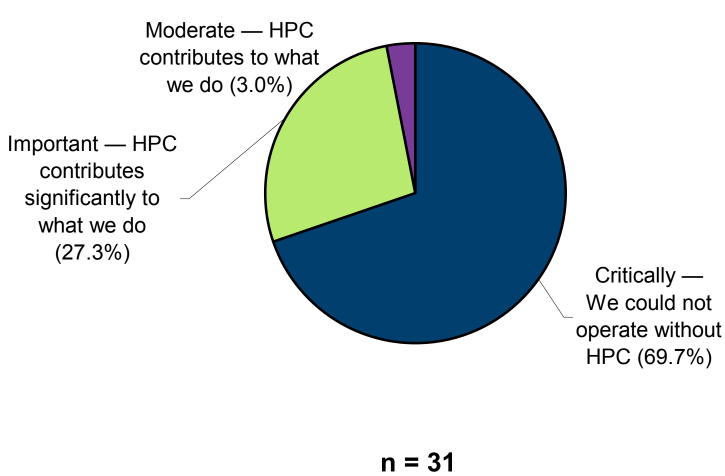

Source: IDC, 2004

"HPC is essential to today's computing needs in industry."

"We can't build airplanes without HPC."

"It is crucial for the design of our [automobile] products."

"It is a key component of our R&D. It helps us calculate properties before engaging in experiments. It has on occasion been a vital part of rescuing product development projects that were not making expected progress through experimental means."

"This is our core business, and the [petroleum] industry could not exist without it."

"We as a company could not produce our [automotive] products, and we would not be in business."

"It allows for faster development cycles and more economical research."

"This is our core business, and the [petroleum] industry could not exist without it."

Q7. How does your organization acquire access to these computing resources? (More than one response is possible.)

Most of the organizations (54%) purchase HPC systems and install them in their own facilities, and another 16% install systems in their facilities under leasing arrangements (see Table 2). Hence, 70% of the organizations access HPC systems on site — further evidence of the crucial importance of these systems.

54% of organizations purchase HPC systems and install them in their own facilities.

TABLE 2

Methods of Accessing HPC Resources

Method of Accessing HPC Resources	Number of Responses	Overall Percentage
Purchase and install them in our facilities	30	54
Use systems installed in partner facilities (e.g., vendors, universities, labs)	11	20
Lease and install them in our facilities	9	16
Use resources over a grid or from an Internet provider	1	2
Other	5	9
Total	56	100

n = 33

Note: More than one response per site was allowed.

Source: IDC, 2004

One in five of the respondents accesses HPC systems at partner facilities. This is a common practice, for instance, among independent software vendors. Many ISVs are small businesses (fewer than 100 employees) that must test their software on relatively large HPC systems from multiple vendors without having the financial wherewithal to purchase the systems. Instead, they frequently arrange to use systems at universities, government laboratories, or the sites of the vendors that manufacture the systems.

Only one of the organizations accesses HPC resources via a grid or Internet provider. In recent years, grid computing has evolved well beyond its origins in projects such as SETI Online, where many thousands of Internet-connected PCs contributed available cycles to a common, massive computing task — in effect becoming a free-of-charge throughput supercomputer. Although grid computing is actively employed today by only a minority of HPC users, a growing array of standards, software, and partnerships promises to expand its use.

Only one of the organizations accesses HPC resources via a grid or Internet provider.

Q8. If you access your HPC tools via an external provider or if you supplement your needs with access to a university or national lab, etc., who makes the decision to do it externally and what criteria are used to justify the decision? (Please list only their job title and/or organization level.)

Even among the minority of respondents who said they use external HPC resources (external providers or partners), on average only 7.5% of their total high-performance computing needs are currently being met in this way. Security concerns are an important inhibiting factor for some organizations. The decision to use outside resources is typically not assigned to a single official. Instead, these decisions may be made by various officials within an organization, depending on the momentary need for additional resources. The implication is that, at least in most cases, the use of external HPC resources is not part of a planned approach today and is not yet considered an ongoing business necessity. Because the use of outside resources is still relatively new, respondents who employ these resources also find it hard to describe their level of satisfaction when asked to do so.

Q9a. Who makes the purchase decisions for HPC tools/computers in your organization — on the technical side?

Decisions about which HPC computer an organization will purchase (or lease) can be complicated affairs. This is not surprising, given the often mission-critical importance of these resources and their substantial price tags. The extent of the purchasing process — and how high up it reaches into an organization — is often a function of the organization's size and the computer's price.

In virtually all cases, the process involves a technical decision maker. In about 75% of cases, a financial officer is also involved in making the decision.

In virtually all cases, the process involves a technical decision maker. In about 75% of cases, a financial officer is also involved in making the decision.

The primary technical decision maker (see Figure 2) most often (55%) is either the vice president of product development — sometimes titled vice president of engineering/R&D — or the manager of the technical datacenter. Especially in large industrial organizations, such as automotive and petroleum companies, the decision-making process frequently involves one of these "technical champions" gaining consensus among the technical staff and then advocating their proposal to financial and other nontechnical executives.

FIGURE 2

Primary Technical Decision Maker

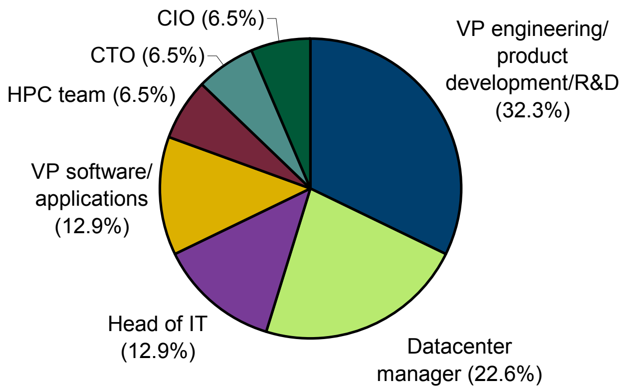

n = 31

Source: IDC, 2004

In smaller industrial organizations and ISV firms, the CTO typically has far fewer steps to follow and, in some cases, may be authorized to make the final purchase decision. In commercial (nonindustrial) organizations (e.g., in the financial or transportation logistics sectors), the CTO's role more often is CIO or head of IT.

Q9b. Who makes the purchase decisions for HPC tools/computers in your organization — on the financial side?

For the majority of purchasing processes that also require a financial decision maker, this party's title can vary dramatically — from an IT officer up to the company's CEO or even board of directors (see Figure 3).

FIGURE 3

Primary Financial Decision Maker

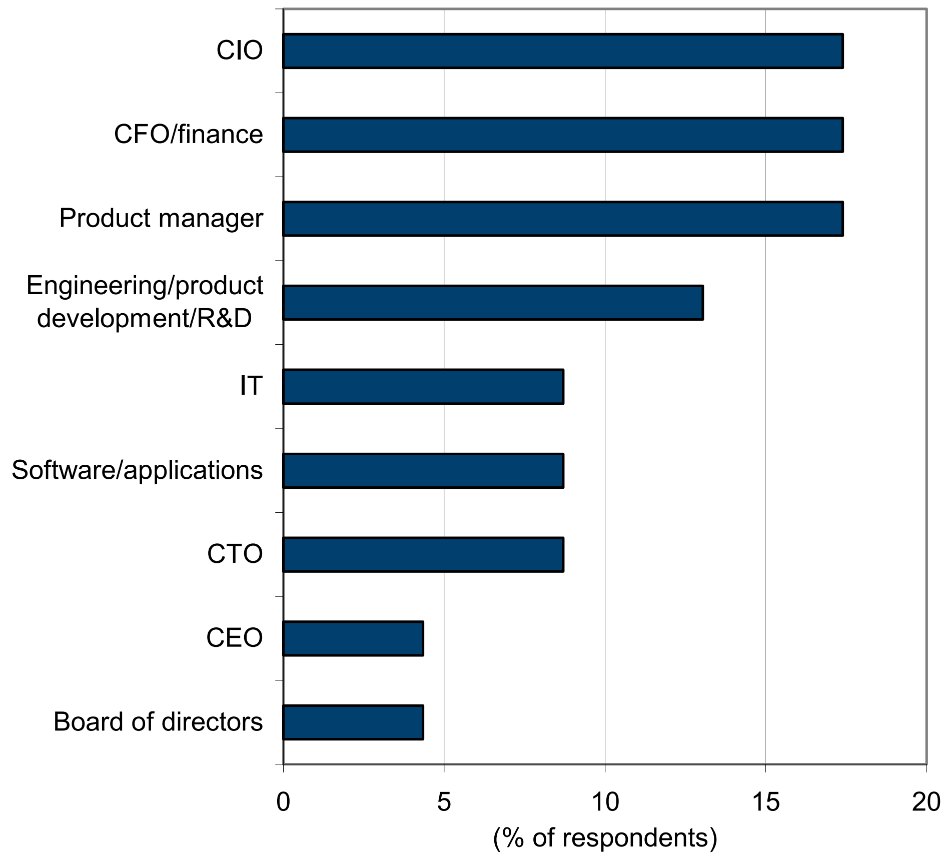

n = 23

Source: IDC, 2004

Q10. Regarding the type of computational problems running at your site, what percentage is capability and what percentage is capacity?

As Figures 4 and 5 indicate, across the sites surveyed, slightly more than one-quarter (29.4%) of the average workload of the organizations consisted of capability-class problems, with capacity-class problems making up the remaining nearly three-quarters (70.6%) of the typical workload. However, 63.6% of the organizations run at least some capability problems. Most sites (55%) run a mix of capability and capacity problems.

29.4% of the average workload of the organizations consisted of capability-class problems.

FIGURE 4

Capability Versus Capacity Problems as Percentage of Workloads

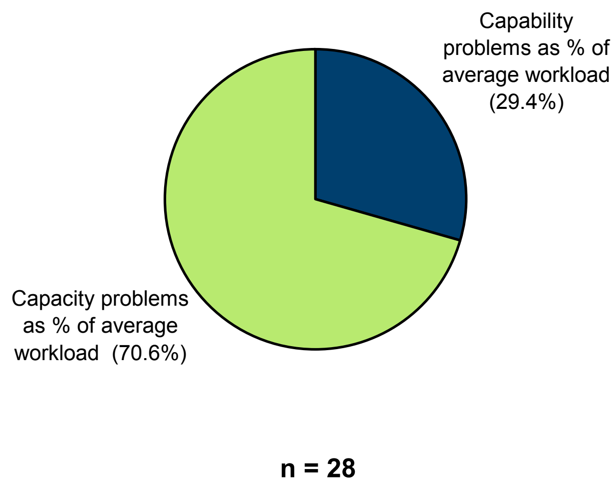

Source: IDC, 2004

FIGURE 5

Sites Running Capability and/or Capacity Problems

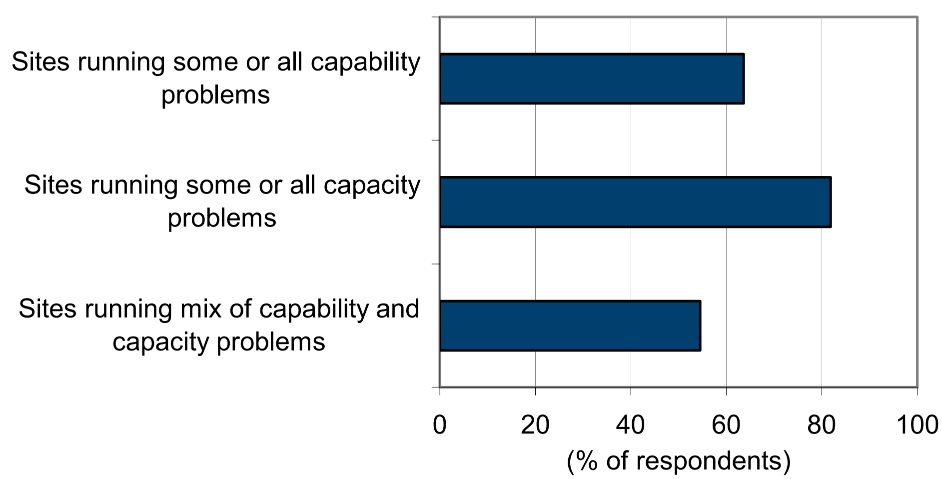

n = 33

Source: IDC, 2004

Q11a. Please rate the following criteria as used by your organization to make your HPC purchase decisions for CAPABILITY-class computers (use a scale from 1 to 10, with 10 = most important and 1 = not used).

The top five criteria for purchasing capability-class computers are closely grouped in popularity, but "performance on our applications" stands out most prominently, with a rating of 8.9 on a scale from 1 to 10 (see Table 3). Solving problems — often by running the same application repeatedly to close in on an optimal solution — can be far more time-critical for industry than for government and university organizations pursuing scientific research. Industrial and other business firms are driven by external competition in a never-ending race to be first to market with the best products. In these battles for global market supremacy, faster application performance often means faster time to market. In the race for market supremacy, the ability of a particular HPC computer to run an organization's applications faster than competing computers can outweigh other considerations, including price/performance (8.0 rating) and price (7.7 rating). Cost considerations are nevertheless very important, as these ratings indicate.

In the race for market supremacy, the ability of a particular HPC computer to run an organization's applications faster than competing computers can outweigh other considerations.

Respondents explained that not all capability-class computers can run every software application. For applications to run at all on the systems, they have to be explicitly "ported" (i.e., adjusted to operate compatibly on the system). (To run as well as possible on the computer in question, the applications also need to be optimized — modified to take advantage of the system's design features.) Accordingly, the ability to run specific applications that are important to an organization figures just below application speed and cost considerations in importance (7.6 rating).

Less critical, yet still important, are the computer's quality/reliability, the reputation of the vendor, and the prospects for running next-generation applications on the system.

TABLE 3

Ratings of Criteria for Purchasing Capability-Class Computers

Criterion	Rating (10 = Most Important)
Performance on our applications	8.9
Price/performance ratios	8.0
Price or budget level	7.7
Ability to run certain software and/or applications	7.6
Quality/reliability	7.1
Vendor reputation	5.9
Future application requirements	5.5
Other	7.3

n = 19

Note: This question applies only to sites with capability-class computers.

Source: IDC, 2004

Q11b. Please rate the following criteria as used by your organization to make your HPC purchase decisions for CAPACITY-class computers (use a scale from 1 to 10, with 10 = most important and 1 = not used).

The same prioritization of purchasing criteria applies to capacity-class HPC computers (see Table 4), although application performance does not stand out quite as distinctly from the other highly rated factors. The prioritization match between capability and capacity computers is not surprising. In both cases, industrial and other business organizations are employing the computers in the same battle for market supremacy.

TABLE 4

Ratings of Criteria for Purchasing Capacity-Class Computers

Criterion	Rating (10 = Most Important)
Performance on our applications	8.8
Price/performance ratios	8.4
Price or budget level	8.1
Ability to run certain software and/or applications	7.6
Quality/reliability	6.8
Vendor reputation	5.9
Future application requirements	5.6
Other	8.3

n = 22

Note: This question applies only to sites with capacity-class computers.

Source: IDC, 2004

Q12. What percentage of your CAPABILITY-class problems are you interested in outsourcing today? In the future?

When we asked organizations about their use of external HPC resources in general (refer back to Table 2), only a small percentage said that they outsource today. The same pattern emerges when the organizations are more specifically asked what portion of their capability-class problems are being outsourced. Only about 2% of these problems are outsourced today, and that figure is expected to climb only to about 7% in the future. The highest figure reported for outsourcing today was 25%, and one organization expects to outsource as much as 75% of its capability workload in the future.

Looked at another way, 81% of the organizations said that they do not outsource any capability-class problems today. 77% have the same expectation for the future.

TABLE 5

Percentage of Capability-Class Problems Being Outsourced

	Average Percentage
Portion outsourced today	2.3
Portion to outsource in the future	6.8
% responding "zero outsourced" — today	81.0
% responding "zero outsourced" — in the future	77.0
Highest % — today	25.0
Highest % — in the future	75.0

n = 31

Source: IDC, 2004

Q13. What organizational or competitive risks do you have if you DON'T have access to high-end HPC computer systems/tools?

97% of the industrial and business organizations consider access to high-end HPC computer systems/tools indispensable.

When asked to name the most important organizational and competitive risks they would face if they did not have access to HPC computers (see Figure 6), nearly half (47%) of the organizations echoed responses to earlier questions by stating that they could not exist as businesses without HPC. An additional 50% replied that they would be unable to compete, emphasizing either time-to-market and related cost issues or product testing and quality issues. Hence, 97% of the industrial and business organizations consider access to high-end HPC computer systems/tools indispensable. Only one organization indicated that it could still exist and compete without the use of HPC.

F I G U R E 6

Organizational/Competitive Risks from Not Having Access to HPC Computers

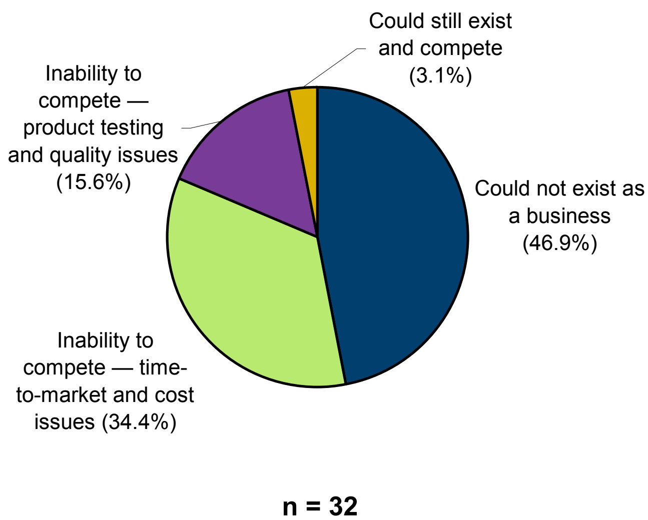

Source: IDC, 2004

"There is no other way for us to complete our work. We would not exist."

"There would be a great lack in quality."

"We would not be able to stay technologically ahead of other competing nations."

"In the semiconductor business, it is critical to continue to push the processor node to remain competitive. Lack of HPC tools similar to those of our foreign competitors could reduce our ability to compete."

"We can't build cars without them."

"The time to market would prohibit our business from existing."

"We cannot support customers without access to HPC machines. This means we would have to tell customers to run in 'at your own risk' mode, which is generally not acceptable."

"There is no other way for us to complete our work. We would not exist."

"In the semiconductor business, it is critical to continue to push the processor node to remain competitive. Lack of HPC tools similar to those of our foreign competitors could reduce our ability to compete."

Q14. Does your purchase process consider future application requirements?

Earlier in this study, some respondents listed future application requirements among the important criteria they consider when purchasing capability- and capacity-class computers. When respondents were asked whether they consider future application requirements to any extent in purchase decisions (see Figure 7), 91% replied yes and only 9% said no. For those replying yes, the average future time frame under consideration was 2.5 years.

FIGURE 7

Are Future Application Requirements Considered in the Purchasing Process?

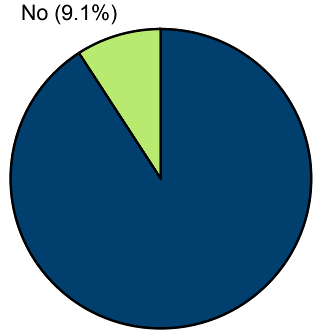

Note: The average number of years respondents looked is 2.5 years.

Source: IDC, 2004

Benefits from Using HPC

Q15a. What has been the benefit of HPC on your organization? Impact on bottom line — can you quantify?

Most sites have difficulty proving the direct benefits of HPC. Some of the major success stories are due to a combination of strong HPC computers, great scientists and engineers, and market conditions; therefore, sites find it hard to pinpoint and quantify which part provided the success. As Table 6 shows, about one in four (23%) respondents was able to quantify the bottom-line benefit of HPC to the organization, and most others (65%) could describe the realized benefit, ranging up to millions of dollars.

About one in four (23%) organizations was able to quantify the bottom-line benefit of HPC to the organization, ranging up to millions of dollars.

For those respondents who were able to quantify, financial benefits were typically substantial, ranging up to millions of dollars and impressive returns on investment in the HPC computers. Nonfinancial benefits included, in one case, an impressive reduction in time to market from five years to two years.

Nonquantified realized benefits from HPC computers fell into the major categories of business criticality (27%) and increased revenue (23%), with increased quality and productivity constituting less frequent responses.

TABLE 6

Bottom-Line Benefit of HPC on Organizations

Impact to Bottom Line	Typical Comment	Number of Responses	Percentage
Able to quantify			23
	50% return on capital employed (ROCE)	1	
	Greater than 5x return on investment (ROI)	1	
	Time to market dropped from five to two years	1	
	More than $1 million	1	
	Several million dollars	1	
	ROI returned within one year of purchase	1	
Can't quantify but can describe			65
	Critical to our business	7	
	Increased revenue	6	
	Reduced costs	2	
	Increased quality	1	
	Increased productivity	1	
Unable to quantify or describe			12
	Unable to quantify	3	

n = 26

Source: IDC, 2004

One respondent indicated that the use of HPC and modeling and simulation tools saved a product development effort that was on the verge of being cancelled. That product is now on its way to market.

"It has been a continuous stream of revenue to our bottom line, giving us the ability to look into other development areas."

"It drives innovation, R&D effectiveness, and productivity."

"The ROI is returned within one year of the purchase."

"The ROI is returned within one year of the purchase."

Q15b. What has been the benefit of HPC on your organization? Increased competitiveness — describe how.

Respondents were clear about HPC's benefits to their organizational competitiveness (see Figure 8). Many confirmed that HPC is a competitive necessity (41%) and that it provides faster time to market (41%). Others noted that HPC has enabled them to extend their market reach and improve product quality.

Respondents were clear about HPC's benefits to their organizational competitiveness.

FIGURE 8

Benefits of HPC to Organizational Competitiveness

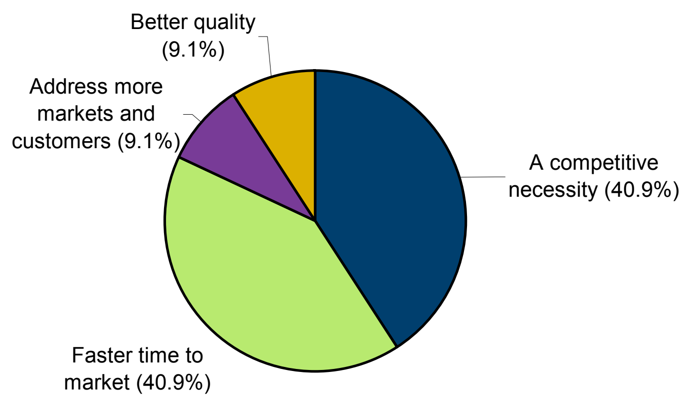

n = 22

Source: IDC, 2004

"It is the only way we can keep up."

"It is the only way we can keep up."

"We can develop drugs faster."

"Can get patent quicker."

"Can get patent quicker."

"Our ability to support our software on HPC gives us an entry into many sectors, which would be difficult if we could not develop or deploy our tools on these platforms."

"These resources maintain the U.S. lead in several critical asset areas."

"These resources maintain the U.S. lead in several critical asset areas."

Q15c. What has been the benefit of HPC on your organization? Increased productivity — in what way?

Respondents also underscored the benefits to organizational productivity of using HPC (see Figure 9). The leading productivity benefit — supporting the all-important time-to-market goal — is shortening product development cycles (45%). The ability to tackle more and larger problems also ranks high as a productivity benefit. Presumably, some of these competitively important problems would have been unaddressed without sufficiently powerful HPC computers.

FIGURE 9

Benefits of HPC on Organizational Productivity

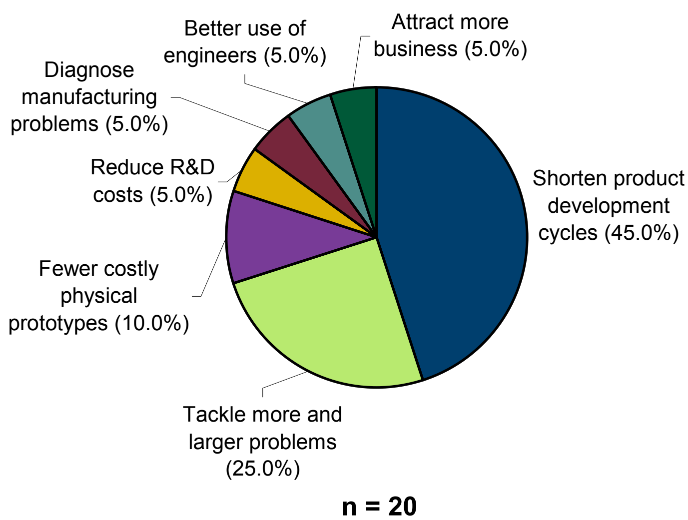

Source: IDC, 2004

"We can get much more done and experiment in many more ways."

"We have the ability to run more jobs and take more samples before investing millions of dollars in the physical tools."

"Better engineering analysis delivered faster allows faster [semiconductor] yield enhancement."

"Productivity is addressed on two levels: first, solving problems and doing R&D that otherwise would not be done; and second, providing more timely answers to problems."

Are Companies Using HPC as Aggressively as Possible?

Q16. Is your organization using HPC tools as aggressively as it could?

A majority of the organizations (58%) indicated that they are not using HPC as aggressively as they could (see Figure 10). The main reasons provided were budget limitations, insatiable demand from end users, time needed to train everyone, and need for better HPC computers in the market.

A majority of the organizations (58%) indicated that they are not using HPC as aggressively as they could.

FIGURE 10

Are Organizations Using HPC as Aggressively as Possible?

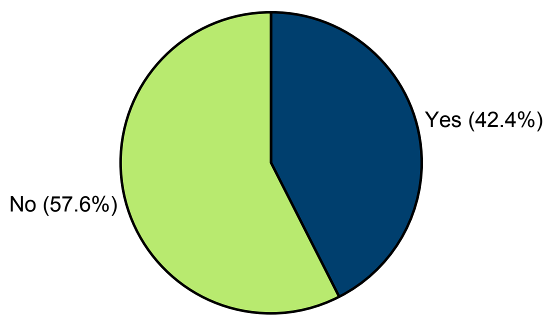

n = 33

Source: IDC, 2004

Yes, Using HPC as Aggressively as Possible

"We are bringing all available tools to bear on our problems."

"We have large compute resources that are designed to run lots of jobs. Our workload was up 10x last year."

"We have access to the very latest systems."

No, Not Using HPC as Aggressively as Possible

"Despite good success, we are under continued pressure to cut the headcount of qualified computational scientists."

"Despite good success, we are under continued pressure to cut the headcount of qualified computational scientists."

"There can never be enough HPC for us. Benchmarking and testing require infinite resources."

"Budget is always an issue, but the key point here is the cost of simulation software. It dwarfs the cost of the hardware."

"Software licensing issues are causing a downsizing in the application seats we purchase. This is a sole source application."

"Without enough human resources, adding additional compute power does not matter."

"Software licensing issues are causing a downsizing in the application seats we purchase."

Q17. What important computational problems do you have today that you can't solve today?

As Figure 11 illustrates, two-thirds of the organizations have important problems that they can't solve today. The remaining third need faster HPC computers to solve current problems more effectively.

FIGURE 11

Do Organizations Have Important Computational Problems They Can't Solve Today?

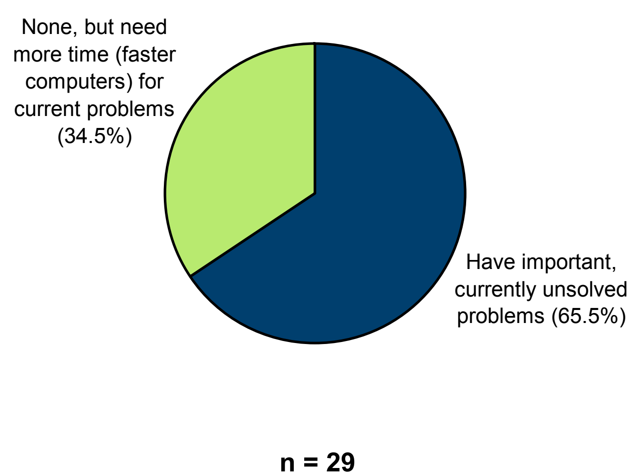

Source: IDC, 2004

Table 7 lists the problems respondents identified as currently unsolved in their organizations or industries.

TABLE 7

Current Unsolved Computational Problems

Industry Sector	Current Unsolved Problem
Automotive/aerospace	Model engine block assembly in full detail (no "submodeling")
	Problems 10x larger than the largest problems today
	Air design and testing
	Crash tests with better body models
	3D system simulation and optimization within 24–48 hours
	Vehicle rollover
	Aero-acoustics/wind noise
	Combustion
	Manufacturing
Petroleum	Real-time processing of data from sensors in remote locations
Pharmaceutical	Protein folding
	Ab initio molecular mechanics — more than 100x current speeds
Semiconductor	2D and 3D modeling of device physics, semiconductor electronics
General (non-specific)	Highly distributed database — coordinate across our 21,000 servers
	Increase mesh size for increased resolution from our data
	Increased breadth and depth of analysis
	Simulate full system
	Compute models two generations ahead
	Problems requiring very large memory

n = 29

Note: Some of the unsolved problems listed were given by more than one site.

Source: IDC, 2004

Barriers Inhibiting Use of Supercomputing

Q18. If you are not using HPC tools as much as you think your organization should, what is holding you back? (More than one response is possible.)

Table 8 shows that the largest single factor (16%) preventing organizations from using HPC tools more aggressively today is a shortage of human experts — internal or external — able to apply the HPC tools to the problems in question. Clearly related to the shortage of expertise, ease of use of hardware and software tools ranks second as an inhibitor (11%). Some other important constraints are attributed to limited management vision: the difficulty in getting approval to make future-oriented investments (10%) and the difficulty higher-ups have in grasping the important contributions of HPC (8%). The response to question 21 indicates that computers that are easier to use and program could help reduce the top two barriers.

The largest single factor (16%) preventing organizations from using HPC tools more aggressively today is a shortage of human experts.

TABLE 8

Factors Holding Back Organizations from Using HPC Tools More Aggressively

Reason for Not Using HPC Tools as Much as Possible	Number of Responses	Overall Percentage
Availability of internal or external people to apply the tools to our problems	12	16.4
Ease of use (hardware and software)	8	11.0
Easier to get decision on investment that reduces costs now versus future	7	9.6
Cost of HPC tools (hardware, software) versus other business investments required	7	9.6
Decision makers do not grasp HPC impact versus other business pressures	6	8.2
Scalability of commercial ISV software	6	8.2
Cost of developing in-house software	5	6.8
Ease of accessing outside resources	5	6.8
Don't have the workload to justify the expense	4	5.5
Technology is changing too quickly to keep up	4	5.5
Hesitant to run company-sensitive problems on outside resources	2	2.7
Availability of appropriate commercial software or applications	1	1.4
Ability to charge against a government contract	1	1.4
Other	5	6.8
Total	73	100.0

n = 29

Note: More than one response per site was allowed.

Source: IDC, 2004

"Time to change people and their reluctance to try new things."

"Internal human resources."

"ECAD is very expensive today, and there are few sources for the types of tools we need."

"The corporate structure has been changing, so less focus has been on technology to run the business."

Q19a. Can you explain why your organization uses different tools for specific applications and the limitations of your current tools on these applications? Why did you choose these particular computational tools for your applications?

When respondents were asked to give their reasons for selecting one HPC computer over others (see Table 9), the predominant reason (31%) was superior performance on the organization's in-house application codes and related requirements. The next two reasons, which are closely related — "only choice available that meets our needs" and "perform best on external (ISV) codes we use" — elevate this overall response category to 62%. This finding is consistent with the primary importance assigned to organization-specific application performance presented earlier. For many (not all) industrial and business organizations, HPC computers are more than mere productivity tools — they are enabling tools without which, as we also saw earlier, these organizations would not be able to compete effectively and survive.

TABLE 9

Reason for Choosing HPC Computational Tools

Reason for Choosing HPC Computational Tools	Number of Responses	Overall Percentage
Perform best on our in-house codes and requirements	9	31.0
Only choice available that meets our needs	6	20.7
Perform best on external (ISV) codes we use	3	10.3
Price/performance	3	10.3
Trusted vendor relationships	2	6.9
Need to use same tools our customers use	2	6.9
Need to meet customer requirements	1	3.4
Have always used the same tools	1	3.4
Follow market trends	1	3.4
Efficiency	1	3.4
Total	29	100.0

n = 29

Source: IDC, 2004

Q19b. Can you explain why your organization uses different tools for specific applications and the limitations of your current tools on these applications? Are they adequate for your current needs? If not, why not?

Most organizations (72%) believe that HPC hardware and software tools currently available in the market are adequate for their needs, whereas 28% do not. For the latter group, the reasons for the inadequacy are almost as various as the respondents (see Table 10). Insufficient computing power is one repeated theme, however.

FIGURE 12

Are Currently Available HPC Hardware Tools Adequate for the Organizations' Needs?

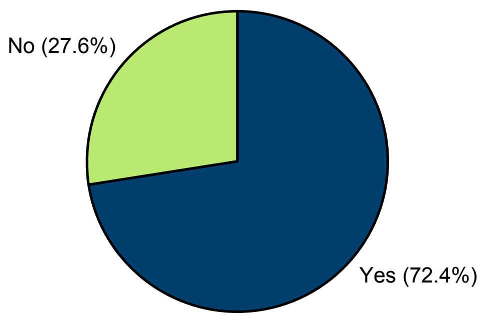

n = 29

Source: IDC, 2004

TABLE 10

Reasons Why Current Tools Are Inadequate

Reason	Number of Responses
Do not provide high enough problem resolution	1
We always need the next generation	1
We lack high-end systems to do real science	1
Development tools on newer (cluster) systems lag traditional SMPs	1
Need better file system than NFS	1
Third-party software is cost-prohibitive	1
We are still growing and must balance the costs	1
We need more linked systems to provide the best information to customers	1
Total	8

n = 8

Source: IDC, 2004

HPC Tools That Industry Would Like to Have

Q19c. Can you explain why your organization uses different tools for specific applications and the limitations of your current tools on these applications? Are there any other HPC tools (hardware or software) on the market today that you would like to own or have access to? (Please list them).

Most organizations are more limited by their budgets and not the computers available in the market. Relatively even numbers of respondents pointed to currently available software (42%) and hardware (37%) tools that they would like to own or access. The most desired currently available tools are cluster management software and clusters based on AMD Opteron processors. 21% of the respondents said that the higher-performance computer hardware systems that they need are not available today. Figure 13 shows that 31.6% feel that adequate hardware and/or software tools are not currently available in the market today.

TABLE 11

HPC Tools Available Today That Organizations Would Like to Own or Access

Tool Desired	Specific Comment	Responses	Percentage
Software			42.0%
	Cluster management software	3	
	Grid management software	1	
	Cluster- and grid-enabled modeling tools	1	
	Distributed memory tool kit	1	
	Affordable Linux file system	1	
	Affordable back-end design software	1	
Hardware			37.0%
	Opteron cluster	3	
	SMP system	1	
	Rendering gear	1	
	Alternative architectures	1	
	More powerful system	1	
Tools we need are unavailable			21.0%
	More powerful computers than today	4	

n = 19

Source: IDC, 2004

"It would be useful to have access to 32–64-way SMP machines with 250–500GB of RAM and approximately 5TB of disk."

"The area where we could use additional refinement is cluster management."

"The area where we could use additional refinement is cluster management."

"The AMD [Opteron] systems we use are the fastest in the world. If we were to use products based on competitive processors, we'd be over budget and late to market."

"The market lacks good products."

"Better computers are needed."

Q19d. Can you explain why your organization uses different tools for specific applications and the limitations of your current tools on these applications? What is stopping you from owning/accessing them?

Costs/budget limitations are the primary reason (42%) that organizations cannot acquire the HPC tools they currently do not have and want, but the unavailability of the desired tools in today's market is not far behind (32%) as a stated reason (see Figure 13). A number of organizations (26%) simply lack the time to determine whether the desired HPC tools truly live up to their promise. (Note: Figure 13 applies to all aspects of the system including hardware, software, networking, etc.)

FIGURE 13

Reasons for Not Owning/Accessing Desired Tools

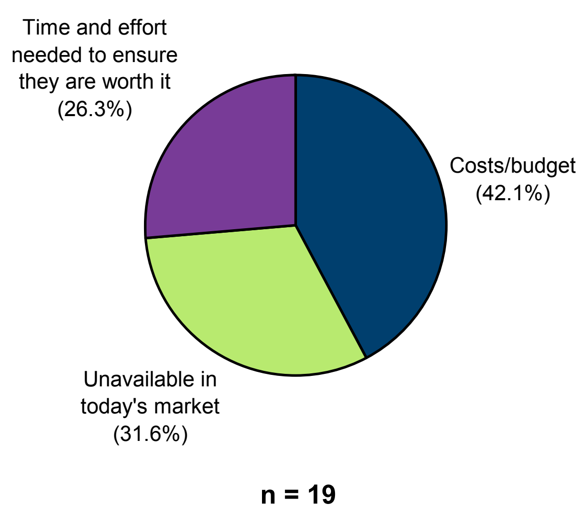

Source: IDC, 2004

"Cost is primary. Buying such machines is outside our budget. Finding such machines that we can use extensively is difficult."

"Budget is too tight."

"We have budget. The tools are just not there yet."

"They do not exist. We need a really good global shared file system. Current ones are immature."

"Before we implement a new technology, we have to be sure it meets our business needs, which can be a very time-consuming endeavor."

Impact of 100x Performance and 10x Ease of Use

Q20a. If you had access to a computer with 100x performance, how would you use it and what new computational science issue could you address?

Naturally enough, HPC users look forward to the arrival of dramatically faster computers. When asked about the impact of a computer 100 times faster than those available today, respondents cautioned that this would need to mean 100 times faster on their applications and then described the effects they foresaw (see Figure 14). Chief among these effects (31%) is the ability to produce higher-quality products based on more accurate simulation. Faster time to market and the ability to tackle larger problems (21% each) also were popular choices, closely followed by the ability to address next-generation problems and make breakthrough advances (17%).

Chief among these effects (31%) is the ability to produce higher-quality products based on more accurate simulation.

FIGURE 14

Overall Impact of Computer with 100x Current Performance

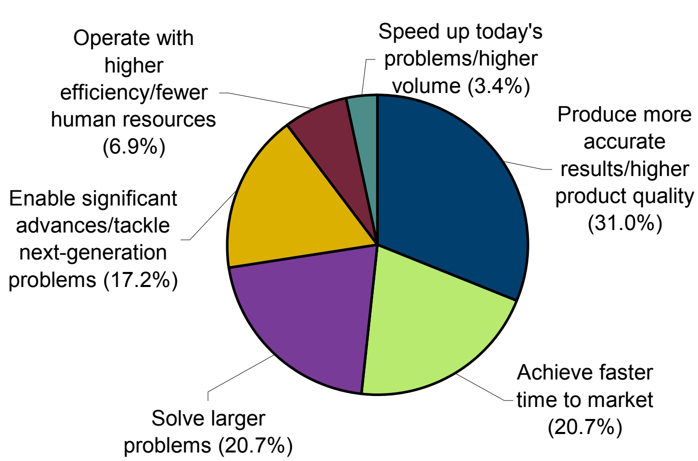

n = 29

Source: IDC, 2004

"100x faster on our codes would be very helpful against our large competitors, in both commercial and military aircraft designs."

"If it is really 100x faster on our applications, we could accomplish some real advances in our science. Better products, better financials."

"We can always find ways to use more power. The first would be to shorten the design cycle while designing better cars."

"We could substantially improve our time to market and chip design variance studies."

"We could test two-generations-out models that we are researching today."

"It could collapse some multiday processes into a few hours and reduce the overall design and manufacturing cycle."

"If it is really 100x faster on our applications, we could accomplish some real advances in our science. Better products, better financials."

Q20b. If you had access to a computer with 100x performance, how would it change your organization?

Respondents were less certain about the organizational impact of a 100-fold faster HPC computer (see Figure 15). A fair number (25%) thought the main impact would be to reduce staffing levels, while others (13%) assumed that more staff would be needed. One in five (21%) respondents believe that a dramatically faster computer would mean little or no change, or else they were not sure what the impact might be.

FIGURE 15

Organizational Impact of Computer with 100x Current Performance

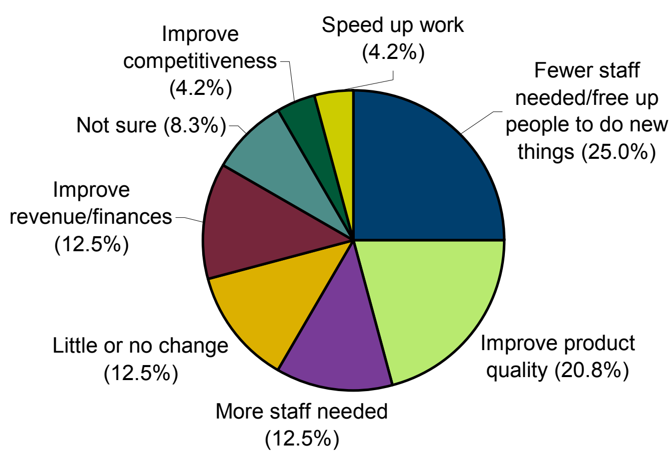

n = 24

Source: IDC, 2004

"It would reduce the time we spend on current tasks, and we could do the other research and program development we want to do."

"We would have to hire a lot more scientists to look at the additional data."

"We would have to hire a lot more scientists to look at the additional data."

"It wouldn't change the organization."

"It would increase revenues for the company and market share."

"Not much [change] at first until we understood the system."

Q21. What could you accomplish if computers were 10x easier to program?

There was greater certainty about the impact of computers that are 10 times easier to program than today's HPC products (see Figure 16). The two most popular responses are not surprising: the ability to develop more powerful applications and fundamentally rewrite current codes (31%) and the ability to shorten design cycles for faster time to market (17%). Less expected was the third-ranking response (14%): the ability to make HPC available to an expanded universe of researchers who don't understand programming — and, with easier-to-program computers, they would not need to train scientists in computer programming areas.

Computers that are easier to use and program could help reduce the top two barriers cited in question 18.

FIGURE 16

Impact of Computers That Are 10x Easier to Program

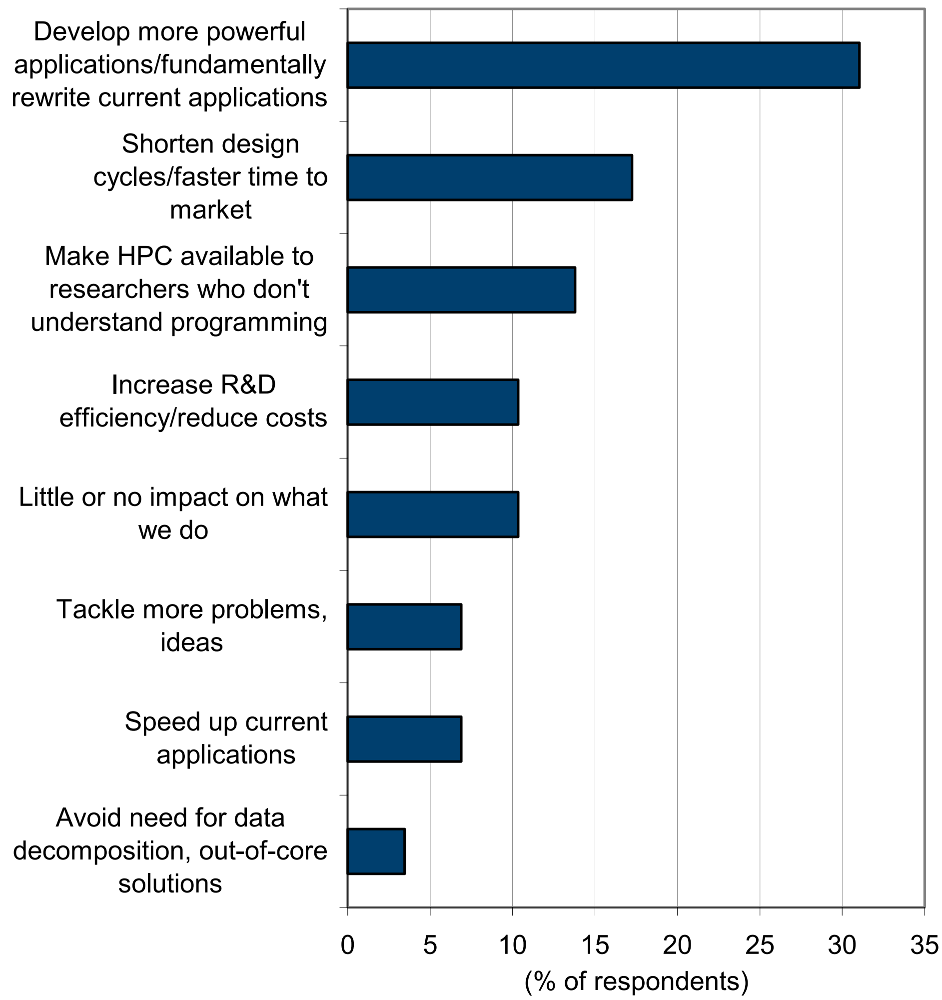

n = 29

Source: IDC, 2004

"We would look to rewrite the entire science underlying the current technology and methodology we are using."

"We would look to rewrite the entire science underlying the current technology and methodology we are using."

"We would be able to develop more powerful and reliable codes."

"It would make these tools available to a much wider array of scientists who have good ideas but may not have programming skills."

"We could try a lot more ideas and significantly reduce the number of assumptions we make."

"If ISV codes are ported, then we could accomplish a lot more — shorten design cycles, faster time to market."

"Not applicable. We mostly use off-the-shelf tools."

Q22. What could you do that you cannot do today if you had these tools?

Figure 17 summarizes the new advances that respondents believe would become possible with computers 100 times more powerful and 10 times easier to program than today's HPC computers. The most important advance (32%) would be in the ability to simulate larger, more accurate (detailed) models. Tackling important new problems (23%), improving product quality while lowering cost (18%), and boosting R&D efficiency (14%) were also mentioned frequently. All of these advances would increase organizational competitiveness.

FIGURE 17

New Advances Made Possible If Organizations Had the Tools They Desire

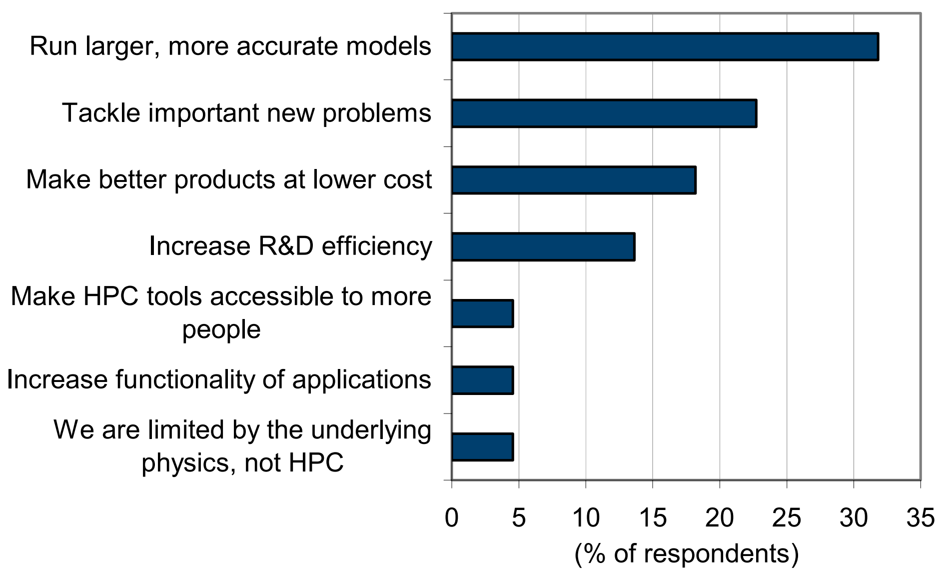

n = 22

Source: IDC, 2004

"We could achieve unprecedented accuracy in the models and significantly reduce the number of assumptions we currently make."

"We could try more ideas with many more techniques."

"We could rewrite all of our underlying science with the new thought processes."

"We could build better aircraft at lower cost."

"We could build lower-cost cars with many additional types."

"We could achieve unprecedented accuracy in the models and significantly reduce the number of assumptions we currently make."

Q23. What features or capabilities would you MOST like to see in future HPC computers, looking over the next five to 10 years? (Please rate on a scale from 1 to 10, with 10 = most important and 1 = not important.) (More than one response is possible.)

Greater processor speed (highest point total) and better price/performance (highest average rating) were popular choices for features organizations would most like to see in HPC computers five to 10 years from now. Greater processor speed refers to greater *delivered* performance. As Table 12 illustrates, there are a number of site-specific desired features/capabilities that were highly rated.

TABLE 12

Desired Features/Capabilities in Future Systems

Desired Feature/Capability in Future Systems	Number of Responses	Average Rating
Individual, site-specific features/capabilities	4	8.5
Better price/performance	27	8.1
Ability to run larger problems	29	7.9
Greater processor speed (delivered performance)	30	7.8
Better system price	25	7.6
Ability to do new science	27	7.3
Easier-to-program computers	23	7.0
Increased reliability	28	6.6
Improved life-cycle cost	27	6.4
Easer-to-use computers	26	6.4

n = 33

Source: IDC, 2004

Comments from the "Individual site-specific features/capabilities" category:

"We need better memory access speeds [mostly bandwidth]. Greater processor speed without better memory bandwidth is making less sense."

"We want to see more memory bandwidth. It seems CPU speed is alright."

"We need a big step change in the data storage — disk write capability."

"We need additional fault tolerance capabilities."

"We need better memory access speeds [mostly bandwidth]. Greater processor speed without better memory bandwidth is making less sense."

Q24a. Given the availability of dramatically better and easier-to-use HPC tools, what would be the impact to your bottom line if you had more capable computers as described above (and they were able to address the problems that you listed above)?

Asked about the bottom-line impact of more capable computers (see Figure 18), the largest percentage of respondents (31%) pointed to anticipated financial improvements. Productivity advances (19%) and better simulations/products (12%) were also noted. Nearly one-quarter (23%) of the organizations found the impact hard to predict.

FIGURE 18

Bottom-Line Impact of More Capable Computers

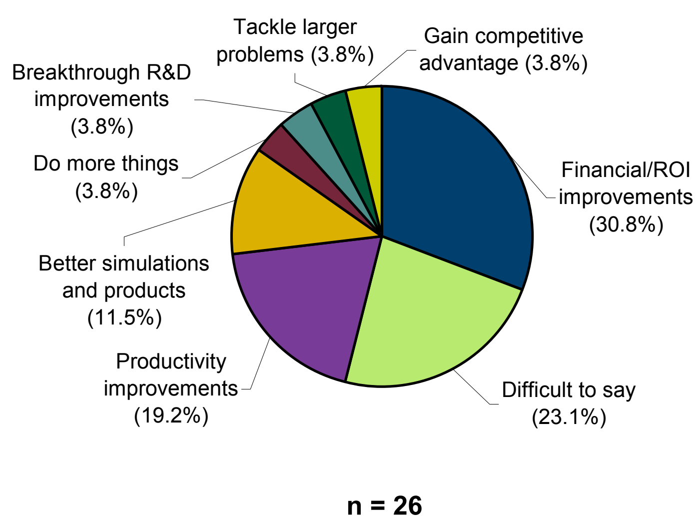

Source: IDC, 2004

Q24b. Given the availability of dramatically better and easier-to-use HPC tools, what would be the impact to your bottom line if you had more capable computers as described above (and they were able to address the problems that you listed above)? Can you quantify the value of solving these problems?

When asked more specifically to quantify the value of solving currently intractable problems, only 17% of respondents were able to do so (see Table 13). For those who did quantify, amounts were large and ranged from $10 million to several billion dollars. Of the remaining 83%, some expect unspecified financial benefits and faster development of improved products, and two respondents said that they are not permitted to discuss anticipated benefits.

TABLE 13

Can You Quantify the Value of Solving Currently Intractable Problems?

Response	Specific Comment	Number of Responses	Percentage
Unable to quantify but can categorize			83.0
	No specifics given	7	
	Financial/ROI benefits	6	
	Better products, faster	2	
	Not permitted to disclose	2	
	New research areas possible	1	
	More efficient use of resources	1	
Able to quantify			17.0
	$10 million	1	
	5% overall gain	1	
	$1 billion	1	
	Multiple billions of dollars	1	

n = 23

Source: IDC, 2004

"There would be significant savings to the entire healthcare system in the areas of surgeon effectiveness, recovery time, and better diagnosis."

"With faster turnaround, we could attract more clients."

"We save $1 billion from a faster product cycle."

"I can't release [the amount], but it is in the billions a year."

"We save $1 billion from a faster product cycle."

Q24c. Given the availability of dramatically better and easier-to-use HPC tools, what would be the impact to your bottom line if you had more capable computers as described above (and they were able to address the problems that you listed above)? What would be the impact to your competitiveness?

91% of the organizations (see Table 14) expect that dramatically faster, easier-to-use HPC computers would boost their competitiveness, most prominently by enabling them to bring better products to market faster or by increasing their financial strength. The remaining organizations argued that the impact of the improved computers would be competitively neutral because "we are already the leader in our business."

TABLE 14

Competitive Impact of Dramatically Better, Easier-to-Use HPC Computers

Competitive Impact	Specific Comment	Number of Responses	Percentage
Positive impact			91.0
	Better products, faster	9	
	Increased financial strength	4	
	Ability to keep pace with competition	3	
	First to market	2	
	Increased customer confidence	1	
	Expand current leadership	1	
	No specifics	1	
Neutral impact			9.0
	Already the leader in our business	2	

n = 23

Source: IDC, 2004

"We would be even more competitive and use the natural resources [oil and gas] to a better degree, creating more financial freedom."

"It would shorten our design cycles and make our cars safer, more reliable, and more interesting so that customers buy them more often."

"It would help us compete against the European manufacturers."

"It would help us compete against the European manufacturers."

"We are staying ahead of the world in specific technical problems. This would just expand that lead."

"We have plenty of market share currently. We just don't want it to decrease."

Q24d. Given the availability of dramatically better and easier-to-use HPC tools, what would be the impact to your bottom line if you had more capable computers as described above (and they were able to address the problems that you listed above)? Other effects?

Other positive effects of dramatically better HPC computers cited by respondents include more powerful pharmaceutical drugs and faster disease cures, more environmentally friendly manufacturing, reduced litigation expense, and more entertaining animated films.

"It would allow our clients to accomplish more research, thus leading to better drugs in a smaller time frame."

"Quicker research means we could get to a cure much faster and reduce long-term healthcare costs."

"Better *in silico* research can have a dramatic environmental impact."

"Lower lawsuit costs."

"Shorter design cycles mean higher margins."

"People would get better animated movies to watch."

Q25. If HPC computers were easier to use and therefore required less dependence on HPC programming experts:

Would that increase your use of HPC computers in general?

Many organizations (56%) said that easier-to-use HPC computers would increase their use of these systems.

Many organizations (56%) said that easier-to-use HPC computers would increase their use of these systems (see Figure 19). A smaller but still substantial percentage (39%) said the same thing vis-à-vis capability-class computers (see Figure 20).

FIGURE 19

Would Easier-to-Use HPC Computers Increase Your Use of These Systems?

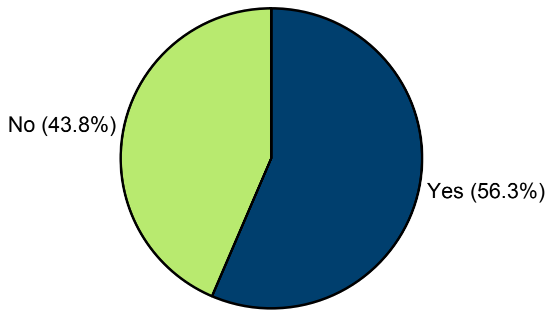

n = 32

Source: IDC, 2004

Would it increase your use of CAPABILITY-class computers?

FIGURE 20

Would Easier-to-Use Capability-Class Computers Increase Your Use of These Systems?

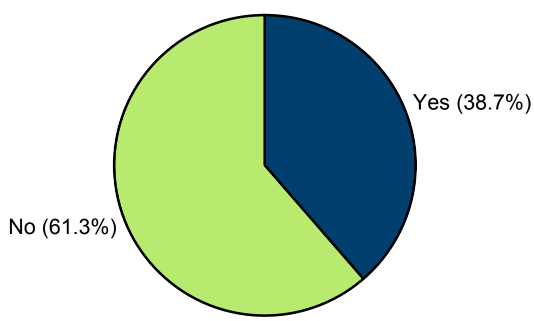

n = 31

Source: IDC, 2004

Budgets and Budget Growth

Q28. Using the following broad categories, what was your company's revenue last year?

As Figure 21 illustrates, more than half (55%) of the organizations' annual revenues exceeded $1 billion in 2003, with another 26% of the organizations reporting revenues in the $1 million to $9.9 million range. This same bifurcated distribution has been evident in prior IDC studies of the U.S. industrial HPC market.

FIGURE 21

2003 Revenue of Respondents' Organizations

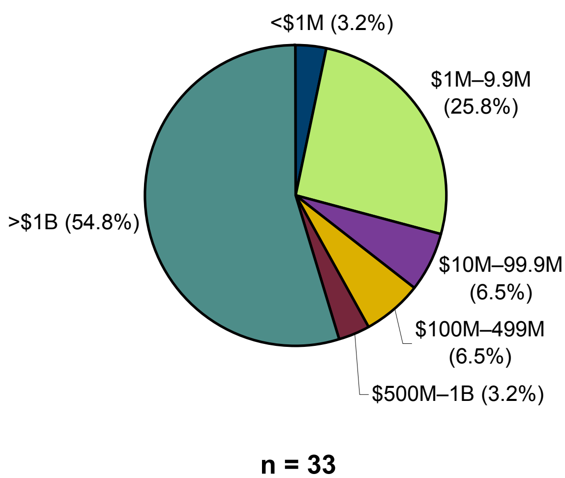

Source: IDC, 2004

Q26. Did your budget over the last year grow, shrink, or stay the same?

Nearly three-quarters (74%) of the budgets (for all HPC solutions including hardware, software, networking, and services) either grew or remained stable during the past year, whereas one-quarter declined (see Figure 22).

FIGURE 22

Budget Change During Past Year

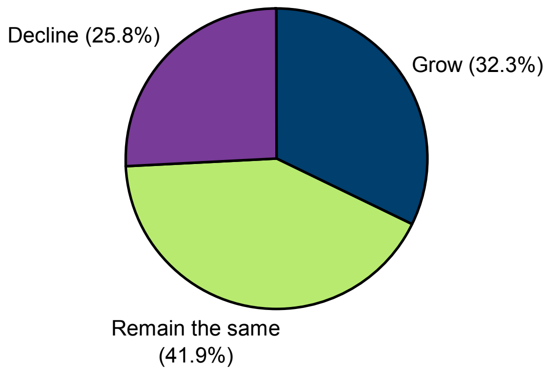

n = 31

Source: IDC, 2004

Q27. Will your budget over the next two years for <u>all HPC</u> spending grow, shrink, or stay the same?

For the next two years, 57% of respondents expect budgets to increase or remain the same, with 40% of the sites interviewed stating that they expect their budgets for HPC to increase. Meanwhile, over the next two years, 43% of respondents expect their budgets to decline and 17% of respondents expect their budgets to stay the same (see Figure 23).

FIGURE 23

Budget Change over Next Two Years

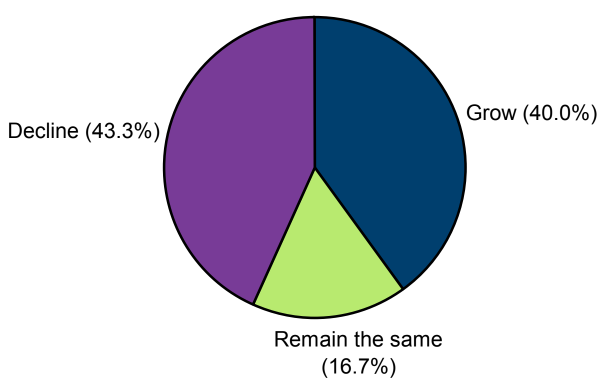

n = 30

Source: IDC, 2004

CONCLUSION

HPC in the United States is an important tool for fostering commercial competitiveness. Buyers have changed the type of systems they acquire and how they apply them to their problems, but they still see HPC as a critical technology for their organizations.

The U.S. industrial market for HPC technical servers has been through a number of changes, evolutions, and revolutions over the past few years. In many ways, it is like going through a physical phase change from a liquid to a solid state, and things have not yet fully solidified.

Major findings in this study include:

- High-performance computing is essential to business survival.
- Companies are realizing a range of financial and business benefits from using HPC.
- Companies are failing to use HPC as aggressively as possible.
- Business and technical barriers are inhibiting the use of supercomputing.
- Companies don't have the HPC tools they want and need.
- Most companies do not rely on remote access to HPC.
- Dramatically more powerful and easier-to-use computers would deliver strategic, competitive benefits.
- Dramatically more powerful and easier-to-use computers could add billions to the bottom line.

Companies are realizing a range of financial and business benefits from using HPC.

Dramatically more powerful and easier-to-use computers could add billions to the bottom line.

Additional observations from recent IDC studies of U.S. industrial HPC end users include:

- Industrial users are still very engaged and excited about applying HPC to help their organizations become more competitive and ultimately more successful
- Price and price/performance have both provided new possibilities and driven buyers to acquire different types of HPC servers.
- Clusters have proven themselves as capable servers to handle a sizable portion of the HPC workload.
- Higher-performance computers are desired based on actual delivered results on end-users computational problems, but most sites cannot afford to purchase the fastest computers available in the market today.

The U.S. industrial sites interviewed clearly see HPC as fundamental to the business, as the following quote illustrates: "Our business model could not exist without HPC."

"Our business model could not exist without HPC."

These sites are strategically using HPC and are investigating new ways to apply HPC.

According to one respondent, "… the use of HPC saved a product development effort that was on the verge of being cancelled. Within a few days, a compound was modeled that would meet all the design criteria and exceed some of them. That product is now on its way to market. The difference between canceling the project and going to market with an improved product could be worth several million dollars."

When respondents were asked if there are currently available HPC tools that they would like to have, 79% answered yes. The inability of so many organizations to access desired current tools due to financial or budget constraints may explain why only a minority indicated that today's tools are inadequate.

For those who were able to quantify dramatically better HPC computers, amounts ranged from $10 million to several billion dollars for their organization alone.

HPC users look forward to the arrival of dramatically faster computers. When asked about the impact of a computer 100 times faster than those available today, respondents cautioned that this would need to mean 100 times faster on their applications. The chief benefit they foresee is the ability to produce higher-quality products. One benefit of computers that are 10 times easier to program than today's HPC products is surprising: the ability to make HPC available to an expanded universe of researchers who don't understand programming. For those who were able to quantify dramatically better HPC computers, amounts ranged from $10 million to several billion dollars for their organization alone.

Other positive effects of better HPC computers include more powerful pharmaceutical drugs and faster disease cures, more environmentally friendly manufacturing, reduced litigation expense, and more entertaining animated films.

APPENDIX

Appendix 1: Site Background and Profiles

Q1. What type of business or industry is your company/department primarily in?

The participants in this study substantially reflect the range of industries and business sectors that employ high-performance computing tools today (see Table 15). In addition, the number of respondents from each industry or business sector is roughly proportionate to the prevalence of HPC tools in that industry or sector. For example, HPC tools are more commonly used in the aerospace, automotive, and petroleum industries today than in the telecommunications or transportation logistics sectors.

HPC tools are more commonly used in the aerospace, automotive, and petroleum industries today than in the telecommunications or transportation logistics sectors.

TABLE 15

Primary Business or Industry of Respondents' Organizations

Industry	Number of Responses
IT and electronics	6
Energy (petroleum, oil, and gas)	5
Chemical	5
Pharmaceutical, biological, life sciences, biomedical	5
Aerospace	4
Automotive	4
Software company	2
Entertainment	2
Telecommunications	1
Transportation and logistics	1
Financial services and economic modeling	1
Other	4
Total	40

n = 33

Note: More than one response per site was allowed.

Source: IDC, 2004

Q29. How many employees are employed at your company?

As Figure 24 illustrates, organizations employing HPC tools tend to be either relatively small (fewer than 100 employees) or relatively large (500 to thousands of employees). The smaller organizations include independent software vendors (ISVs), which may employ as few as a dozen people to develop, test, and maintain crucial third-party software applications. The largest organizations include well-known multinational corporations.

It was interesting to find that the smaller companies and large companies had similar concerns and needs.

FIGURE 24

Number of People Employed at Respondents' Organizations

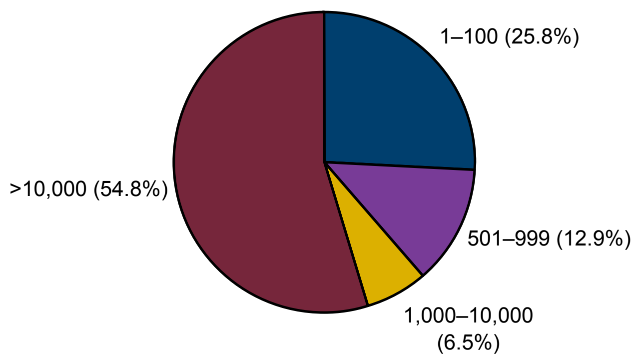

Source: IDC, 2004

Q3. What categories of HPC computers are currently installed or being used by your company? (More than one choice is possible.)

It is not uncommon for organizations that employ HPC to own more than one HPC computer.

For example, the organizations represented in this study have 2.4 HPC computers installed on average. Some organizations acquire multiple HPC computers because each is best-suited for certain problems. Other organizations have multiple HPC systems because they use both current-generation and prior-generation systems. The installed computers may vary greatly in size and purchase price.

The organizations represented in this study have 2.4 HPC computers installed on average.

The preponderance of symmetric multiprocessor (SMP) servers and supercomputers (including vector computers) at the respondents' organizations is not surprising (see Table 16). These systems have been the traditional workhorses of industrial high-performance computing, dedicated to the largest and most daunting computational tasks. Clustered HPC systems as a whole account for approximately half the installed systems; whether purchased from HPC vendors or assembled in-house by user organizations, they offer lower initial prices and better price/performance compared with the more traditional systems. Largely because of these cost benefits, clusters are used for an increasing number of problems in industry and business, although the most challenging problems are generally still reserved for the traditional systems.

TABLE 16

Types of HPC Computers Currently Installed at Respondents' Organizations

Computer Type	Number of Responses	Overall Percentage
SMP technical servers or supercomputers	23	29
Clusters purchased from an HPC vendor	16	20
Beowulf workstation or PC clusters assembled in-house	15	19
Grid computers	14	18
Use of compute cycles from an external party	7	9
Vector computers	4	5
Total	79	100

n = 33

Note: More than one response per site was allowed.

Source: IDC, 2004

Q6. How long (years) has your organization been using HPC technical servers?

The survey respondents as a group are industry veterans who have been using HPC systems for an average of 15 years (see Table 17).

TABLE 17

Number of Years Organizations Have Been Using HPC Technical Servers

Average	Median	High	Low
14.9	15	30	3

n = 33

Source: IDC, 2004

Appendix 2: Survey Questionnaire

Council on Competitiveness Study by IDC
The Industrial Use of HPC Technical Compute Servers

Name: ____
Title: ____
Company: ____
Phone: ____
Fax: ____
Email: ____

****
Qualifiers

I. Do you currently have any technical servers or supercomputers installed at your site or have access to which are used to directly support R&D, engineering and scientific efforts such as simulation, modeling, research, computer-aided design, etc. or computers used to solve problems that are computationally intensive?
____Yes (→ go to question III.)
____No (→ go to question II.)

II. Did your organization previously use HPC technical servers?
____Yes (→ please explain in detail why your organization no longer uses HPC, then end the interview)
____No (→ end the interview)

III. Have you evaluated or been directly involved in the purchase evaluation of a HPC computer system?
____Yes
____No

IF NO -- Is there any other person or department in your organization that may be involved in purchasing system used for technical applications?
Referral Name: ______________________
Phone Number:______________________

****
About IDC

IDC delivers dependable, high-impact insights and advice on the future of all computing markets in the information technology arena such as the high performance computing market.

IDC divides the technical server market into four competitive market segments:
Technical Capability: Systems configured and purchased to solve the largest most demanding problems
Technical Enterprise: Systems purchased to support technical applications in throughput environments selling for $1 million or more
Technical Divisional: Systems purchased for throughput environments selling from $250,000 to $999,000
Technical Departmental: Systems purchased for throughput environments selling for less than $250,000

Additional information about IDC and the reports we produce can be found at www.idc.com and www.idc.com/hpc.
****
Site Overview Questions
Note: For all questions the use of the term HPC computers is used the same way as technical servers, supercomputers, clusters, etc. including Grids, HPC software and storage and is meant to apply to all technical computers used for scientific and engineering problems and applies to computers used to solve problems that are computationally intensive.

Q1. What type of business or industry is your company/department primarily in?

a. ___ Petroleum, oil and gas
b. ___ Chemical
c. ___ Pharmaceutical, biological, life sciences & biomedical
d. ___ Financial services and economic modeling
e. ___ Aerospace
f. ___ Automotive
g. ___ Telecommunications
h. ___ IT and electronics
i. ___ Transportation and logistics
j. ___ Entertainment
k. ___ Other: ________________________

Q2. What type of major high performance compute servers do you currently have installed or have access to? Please list the major systems in use by vendor name, model, number of processors and the top applications and/or problems that you use the computers to solve:

Vendor	System Model	Number of CPUs	Top Applications Or Problems Being Solved	Installed or Have External Access
A				
B				
C				
D				
E				

Q3. What categories of HPC computers are currently installed or being used by your company (more than one choice is possible).

a. ____ SMP technical servers or supercomputers
b. ____ Vector computers
c. ____ Clusters purchased from a HPC vendor
d. ____ Workstation Clusters or PC Clusters Beowulf assemble in-house
e. ____ Grid Computers
f. ____ Use of compute cycles from an external party – Where? (E.g., university, national, lab, service bureau)

****

Next We Would Like To Explore Why You Acquire and Use HPC Technical Servers

Q4. What are the primary reasons why your organization uses HPC computers?

____

____

____

Q5. How does HPC impact your organization's primary goals and/or mission?

____

____

____

____

Q6. How long has your organization been using HPC technical servers?: ___ (# of years)

Q7. How does your organization acquire access to these computing resources?

___ Purchase them and install them in our facilities
___ Lease them and install them in our facilities
___ Use systems installed in partner facilities (including universities & labs)
___ Use resources over a Grid or from an Internet provider
___ Other (please explain: __________________________________)

Q8. If you access your HPC tools via an external provider or if you supplement your needs with access to a university or national lab, etc. -- Who makes the decision to do it **externally** and what criteria is used to justify the decision? (**Please** only list their job title and/or organization level)
______________/______________ (Title and organization level)
Decision justification/criteria: ____
____
____

c. What percentage of your computing are you using external resources today?
____ % of HPC computing done externally
d. What is your level of satisfaction with external sources
(10 = Great, 1 = very low): ___
If low, why: ____

e. What could be done to improve your satisfaction:
____
____

f. If you had access to external HPC resources along with experts to guide you, would you make use of additional external HPC computers resources?: ____ Yes/No

Q9. Who makes the **purchase decision** for HPC tools/computers in your organization (**Please** only list their job title and/or organization level)
Technical decisions are made by: __________________/________________________ (Title and organization level)
Who makes the overall HPC financial decisions (or sets the budgets):
___________________/______________________ (Title and organization level)

Q10. Regarding the type of computational problems running at your site:
Are you currently running CAPABILITY class problems?: ___ Yes/No,
if Yes, what is the percentage _____ %
Are you currently running CAPACITY class problems?: ___ Yes/No,
if Yes, what is the percentage _____ %

Q11.a. Please rate the following criteria as used by your organization to make your HPC purchase decisions for **CAPABILITY CLASS** computers (use a scale from 1 to 10, with 10 = most important and 1 = not used):
______ Price or budget level
______ Performance on our applications
______ Price/performance ratios
______ Quality/reliability
______ Ability to run certain software and/or applications
______ Vendor reputation
______ Future application requirements
______ Other (Please explain: ______________________________)

Q11.b. Please rate the following criteria as used by your organization to make your HPC purchase decisions for **CAPACITY CLASS** computers (use a scale from 1 to 10, with 10 = most important and 1 = not used):
______ Price or budget level
______ Performance on our applications
______ Price/performance ratios
______ Quality/reliability
______ Ability to run certain software and/or applications
______ Vendor reputation
______ Future application requirements
______ Other (Please explain: ______________________________)

Q12. What percentage of your **CAPABILITY** class problems are you interested in outsourcing:
a. _____ % Today
b. _____ % In the future

Q13. What organization or competitive risks do you have if you DON'T have access to high end HPC computers systems/tools?

____
____
____
____

Q14. Does your purchase process consider future application requirements?
___ Yes -- How far into the future to look and/or consider: ____ (years)
___ No

****
Next We Would Like To Explore Your Use of HPC

Q15. What has been the benefit of HPC on your organization?
___ Impact on bottom line—can you quantify:
____
___ Increased competitiveness—describe how:
____
___ Increased productivity—in what way:
____
___ Other: _____________________________

Q16. Is your organization using HPC tools as aggressively as it could?: _____ Yes/No
Please explain:

____
____
____
____

Q17. What important computational problems do you have today that you can't solve today?

____
____
____
____
____
____
____

Q18. If you are not using HPC tools as much as you think the organization should, what is holding you back?
___ Technology is changing too quickly to keep up
___ Final decision-makers do not understand sufficiently the potential impact of using high performance computing to adequately evaluate this investment decision against other business requirements/pressures.
→ If this is the case, where is the internal selling process breaking down:
____
____
____
___ Easier to get a decision on an investment that reduces costs now versus an investment that has future payoffs
___ Availability of internal or external people to apply the tools to our problems
___ Availability of appropriate commercial software or applications
___ Don't have the workload to justify the expense
___ Cost of HPC tools (HW and SW) versus other business investments required
___ Ability to charge against a government contract
___ Ease of use (HW and SW)
___ Scalability of commercial ISV software
___ Cost of developing in-house software
___ Ease of accessing outside resources
___ Hesitancy in running company-sensitive problems on outside HPC resources

___ Other: ________________________________

Q19. Can you explain why your organization uses of different tools for specific applications, and the limitations of your current tools on these applications:

a. Why did you choose these particular computational tools for your applications?
____
____
____

Are they adequate for your current needs? If not, why not:
____
____
____

b. Are there other HPC tools (hardware or software) on the market today that you would like to own or have access to? (Please list them):
____
____
____

What is stopping you for owning/accessing them?
____
____
____

****

Next We Would Like To Explore What Would Be Possible If You Could Obtain Significantly More Powerful HPC Computers

Q20. If you had access to a computer with 100x performance, how would you use it and what new computational science issues could you address?
____
____
____

… And how would it change your organization?:
____
____
____

Q21. What could you accomplish if computers were 10X easier to program?
____
____
____

Q22. What could you do that you cannot do today if you had these tools?
(PLEASE EXPLAIN IN DETAIL)
____
____
____
____
____

Q23. What features or capabilities would you MOST like to see in future HPC computers looking out over the next 5 to 10 years (please rate with 10 = the most important and 1 = not important):

a. ___ Better System Price
b. ___ Improved Life cycle cost
c. ___ Better Price/Performance
d. ___ Increased Reliability

e. ___ Greater Processor speed
f. ___ Ability to run larger problems
g. ___ Ability to do new science
h. ___ Easier to use computers
i. ___ Easier to program computers

___ Other: ____

Q24. Given the availability of dramatically better and easier-to-use HPC tools:

a. What would be the impact to your bottom line if you had more capable computers as described above (and they were able to address the problems that you listed above)?

____
____
____

b. Can you quantify the value of solving these problems?

e.g. reduced costs through more simulation and fewer physical experiments or prototypes; increased revenue from faster introduction of new products. For example, if you could shave off X% of your development costs, this would equate to $Y in cost savings. Or if you could introduce a new product one year early, that could equate to $Y in additional revenue:

Please explain: ____
____
____
____

c. What would be the impact to your competitiveness?
e.g. your ability to increase market share though the development of more new products and/or faster time to market with new products, able to respond more quickly to changing customer needs/demands:
Please explain: ____
____
____
____
____

d. Other effects?
e.g. safer products, reduced environmental hazards from more efficient manufacturing process, faster development of products that affect health/well-being of society etc., reduced prices for customers due to development cost savings, more efficient products/products with longer life-cycle for customers, reduced reliance on foreign manufacturer for important product:
Please explain: ____
____
____
____
____

Q25. If HPC computers were easier to use, and therefore required less dependence on HPC programming experts:
a. Would that increase your use of HPC computers in general?: ______ Yes/No
b. Would it increase your use of **CAPABILITY** class computers?: ______ Yes/No

****

Next We Would Like To Explore How Your HPC Budgets Are Changing

Q26. Did your budget over the last year?:
____ Grow → By what percentage ______%
____ Shrink → By what percentage ______%
____ Remained the same

Q27. Is your budget over the next two years for **all HPC** spending?:
____ Growing → By what percentage ______%
____ Shrinking → By what percentage ______%
____ Remaining the same

****

Additional Company Demographics

Q28. Using the following broad categories, what was your company's revenue last year?
a. ___ < $1 Million
b. ___ $1M - 9.9M
c. ___ $10M - 99.9M
d. ___ $100M - $499M
e. ___ $500M - $1B
f. ___ >$1B

Q29. How many employees are employed at your company?
a. ___ 1 to 100 employees
b. ___ 101 to 500 employees
c. ___ 501-999 employees
d. ___ 1,000-10,000 employees
e. ___ Over 10,000 employees

Thank you for your assistance with this research.

IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Council on Competitiveness Study of ISVs Serving the High Performance Computing Market: Part A - Current Market Dynamics

Earl Joseph, Ph. D.
Christopher G. Willard, Ph.D.
Dolores Shaffer, STA/DARPA
Addison Snell
Suzy Tichenor, Council On Competitiveness
Steve Conway, Conway Communications

July 2005

Commissioned by the Council on Competitiveness
Sponsored by the Defense Advanced Research Projects Agency

EXECUTIVE SUMMARY

This is Part A of a two-part study of the ISVs serving the high performance computing (HPC) market. The complete study is the first independent, assessment of the landscape and market dynamics surrounding Independent Software Vendors (ISVs) that serve HPC users. This first part reflects the opinions and insights of over 100 independent software vendors. See Part B for end user perspectives on these issues.

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[1]. 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 also the *Council on Competitiveness Study of ISVs Serving the High Performance Computing Market: Part B -- End-User Perspectives*, available at www.compete.org/hpc

Key Findings

1) 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 applications 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.

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

"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."

2) 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 Independent software vendor (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.[2]

Few Independent software vendor (ISV) applications today can take advantage of more than 128 processors.

[2] See also the *Council on Competitiveness Study of ISVs Serving the High Performance Computing Market: Part B -- End-User Perspectives*, available at www.compete.org/hpc.

"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."

3) 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 re-write, 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.

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 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."

4) 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% 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% of the applications are based on open source software. (See Figure 1.)

FIGURE 1

Original Source of ISV Application

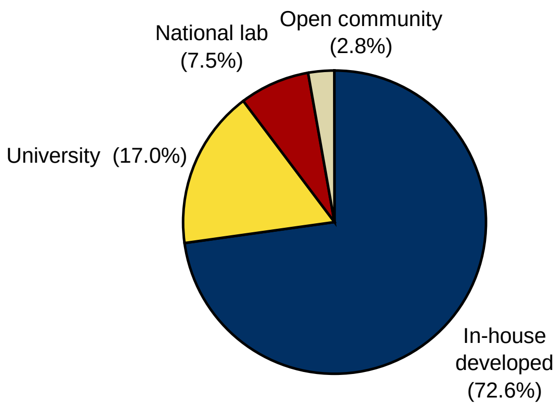

Source: IDC, 2005

The majority of the ISV applications (88%) are supplied by for-profit businesses. By contrast, only 7-8% 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.

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

FIGURE 2

Percentage of Applications by Type of Ownership

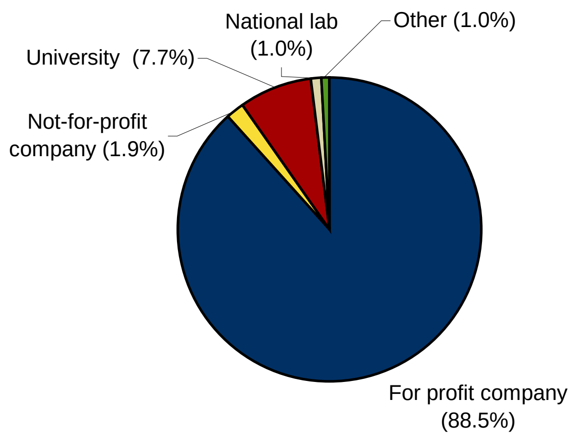

Source: IDC, 2005

5) 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."

6) 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 improve the scalability of ISV software that is needed for improving the competitiveness of U.S. businesses.[3]

Overall annual sales revenues (all products and services) 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.

TABLE 1

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	100.0%	110	100.0%

Source: IDC, 2005

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% 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.

[3] See also the *Council on Competitiveness Study of ISVs Serving the High Performance Computing Market: Part B -- End-User Perspectives*, available at: www.compete.org/hpc.

"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."

7) 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). HPC end-users also were willing to partner with other organizations. For further discussion of partnerships, see Part B of this study, *End-User Perspectives*, available at www.compete.org/hpc.

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	100.0%

Note: Multiple responses permitted.

Source: IDC, 2005

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.)

FIGURE 3

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

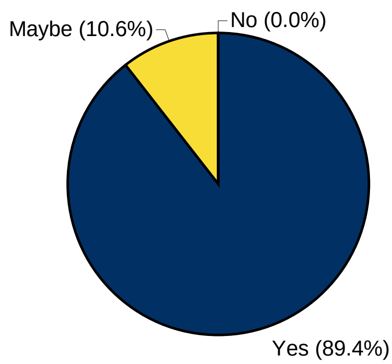

Source: IDC, 2005

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

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

DEFINITIONS AND TERMINOLOGY

Application Software (or Application Software Package)

This term, also called an application program, an end-user program or simply an application or code, refers to a program that performs a specific type of function directly for the user. This is in contrast to system software, such as the operating system, and middleware, such as compilers, libraries, optimization tools and debuggers, which exist to support application software. This study investigates application software used for technical computing. Technical computing application software is used for a wide range of scientific and engineering tasks, ranging from automotive and aerospace design to drug discovery, oil exploration, weather prediction and climate modeling, process engineering, fundamental scientific research, national security, visualization and advanced 3D animation.

Capability-Class and Capacity-Class HPC Computer Servers

IDC defines capability-class computer servers as systems purchased primarily to tackle the largest, most complex single problems. Capability-class HPC systems are generally priced at $2 to $4 million or more, with costs occasionally exceeding $100 million. Traditional symmetric multiprocessor (SMP) technical servers and supercomputers of sufficient size fall into this category, and large-scale clusters also qualify as capability class systems if they are purchased primarily to address large problems.

Capacity-class systems are purchased primarily to solve many small and medium-sized problems. Capacity-class HPC systems may also be priced at more than $1 million and may include any category of HPC computer server. IDC further divides capacity-class computer servers by price band:

- Enterprise: $1 million or higher
- Division: between $250,000 and $1 million
- Department: between $50,000 and $250,000
- Workgroup: below $50,000

HPC

In this document, the term HPC (high performance computing) is used synonymously with the terms HPTC (high performance technical computing) and HEC (high end computing). IDC uses these terms to refer to all technical computer servers used to solve problems that are computationally intensive or data intensive, and also to refer to the market for these servers and the activities within this market. It includes both capability and capacity computers, but excludes single user desktop workstations and PCs.

Clusters

IDC defines clusters used in technical markets as a set of independent computers combined into a unified system through systems software and networking technologies. Thus, clusters are not based on new architectural concepts so much as new system integration strategies. Clusters are considered capability systems when they are used for the most challenging problems (e.g., when used for "traditional" capability-class problems). In the case of capability computing the majority of the cluster's resources (i.e. processors, memory, etc.) will be devoted for a time to solve a single problem. Most clusters are sold as capacity class computers.

Heterogeneous Problem

A heterogeneous problem, also called a multi-physics or multidisciplinary problem, is one that involves multiple scientific disciplines—for example, studying the complex interaction between the structure of an automobile and the fluid dynamics of air flow around it. HPC users are increasingly interested in solving heterogeneous problems, but the software and current hardware systems available are very limited in their ability to address the complexity of this type of problem.

Highly Scalable Systems

The term highly scalable systems is used to refer to HPC computer servers with many—typically hundreds or thousand of—processors. Clusters and massively parallel processing (MPP) computers are two types of highly scalable systems. In the future it is expected that the most capable computers will be configured with hundreds of thousands of processors. As the industry adopts and applies petascale computers to technical problems the issues related to scaling applications to these large sizes is a key concern.

ISV (Independent Software Vendor)

This study uses the term ISV (independent software vendor) to refer to an organization that develops, maintains and makes available application software that is used for technical computing for computer servers. HPC usage typically represents less than five percent of the revenues for many of the application ISVs represented in this study, and in some cases the figure is less than one percent. ISVs may be for-profit, private-sector businesses or public-sector organizations in university or government settings. Although the vast majority (89%) of the ISV's represented in this study are for-profit businesses, the study uses the term "ISV organizations" because some respondents are public-sector entities. ISV's may offer application software, middleware or other software solutions. In the body of this report, only application software is represented; information about middleware appears in the Appendix section.

Last Technology Update

The last time an ISV software code was enhanced, without being substantially re-written. Often ISV's will add new features and functionality to their software on a regular basis without changing the underlying algorithms used in the program. When they invest in a major technology refresh to the underlying algorithms they usually bring the new version to market as a different application package.

Legacy Code

The term legacy code in this study means ISV software that has existed for at least five years, often considerably longer, without being fundamentally updated through a major rewrite. Many of the most used technical application programs are over 20 years old, and were typically designed to run on a single processor.

Major Rewrite

We use this terms to refer to the fundamental rewriting of an ISV application software, typically preceded by a rethinking of the approach to the underlying approach to the problem addressed by the application software. It includes changing the underlying algorithms used by the application program.

Middleware

Middleware refers to a software program that additional functionality over an above that provided by the operating system. Middleware software handles specific tasks like network management, high-level job scheduling, keeping track of files and where they are located, etc. As its name implies middleware sits between the operating system and the application, and may act as "glue" between the two. Middleware may also be used to connect two applications, or two sides of a single application.

Open Source Software

Open source software, also called open community software, refers in this study to ISV application software that is provided to the user community at no or minimal costs. The intellectual property rights are often retained by the ISV organization. It is generally designed to run on open source operating systems, primarily Linux. Open source license agreements typically provide mechanisms for users and other developers to view and modify the original program, or "source code" (thus the term "open source"). Modifications or extensions are generally provided to community as a whole as part of the license agreement. Open source software may or may not be available free of charge. Most of the available open source software for HPC is middleware, rather than end-user applications.

Petascale Computer

A petascale computer is a computer able to operate at petascale performance levels, which is one million billion calculations per second. The DARPA High Productivity

Computing Systems program is currently researching the development of petascale systems for the end of this decade, 2010. There is a broad concern that these systems will require a new level or type of software to be able to extract the full value from these systems. In many cases application software will need to be redesigned and in many cases different types of advanced applications will need to be created, e.g. combining several applications into a single heterogeneous application package in order to take advantage of the capability provided by petascale computers.

Scalability

As used in this study, scalability means the ability of application software to effectively exploit a large number of processors of an HPC computer server, often hundreds or even a few thousands of processors today, growing to tens of thousands in the near future. Many frequently used applications in industry today only scale to 1 to 4 processors in practice, while some may scale to 16 or 32 processors.

Technical Computing and Commercial Computing

The term technical computing, also called scientific and technical computing, refers to the body of computing methods used for scientific, engineering and related computationally intensive tasks. Technical computing activities can be found in industry, government and academia. Industrial activities include: automotive and aerospace product development, oil and gas exploration, drug discovery, weather prediction, complex financial modeling and advanced 3D animation. Scientific researchers in academia and government organizations also use technical computing methods. Technical computing is in contrast to commercial computing as used for business operations such as accounting, payroll, sales, customer relations, transaction processing, human resources and purchasing.

STUDY BACKGROUND

This study provides the first extensive, independent assessment of the landscape and market dynamics surrounding ISVs that serve HPC users. 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). This earlier study and the Council's annual high-performance computing (HPC) users confererence found that many U.S. businesses could not exist, or could not compete effectively, without the use of HPC. But HPC users also indicated that application software challenges were preventing them from using HPC more aggressively.

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. However, 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 computers with hundreds, thousands or (soon) tens of thousands of processors for competitive advantage. Increasingly, these leading-edge users want to solve problems that involve multiple scientific disciplines—for example, studying the complex interaction between the structure of an automobile and the fluid dynamics of air moving around it, or how to extract valuable oil supplies through a porous "mudrock" reservoir. Current ISV applications rarely incorporate multi-disciplinary ("multi-physics," "heterogeneous") capabilities, and few ISV applications today "scale" beyond 100 processors, while many of the most-used ones scale to only a few processors in practice.

Users want to solve problems that involve multiple scientific disciplines, yet current ISV applications rarely incorporate multi-disciplinary capabilities.

This study – also undertaken on behalf of the Council on Competitiveness and sponsored by DARPA – was launched to better understand the causes and extent of this gap. It assess the current capabilities of ISV applications software, the business models and financial resources standing behind this software, and the willingness of ISV organizations to collaborate with outside parties to accelerate progress.

Methodology

This study is based on a broad survey of ISV providers and their applications software packages. We began with a list of 471 software applications and middleware solutions that users and computer vendors pointed to in IDC studies of HPC over the past five years. Though intensive phone interviews and research, IDC gathered current information on 54 of the most important ISV organizations and 110 applications software packages, as well as 20 key suppliers of 64 middleware software solutions. The study included interviews with a total of 104 respondents. Some of the respondents worked for organizations that provide more than one ISV application and gave answers for more than one ISV application, and some were middleware providers.

As described in the "Definitions" section, middleware plays an important role in supporting applications. Even though middleware is seen as a support tool for developing in-house applications and for supporting other ISV applications, they were excluded from the study.

It is important to realize that throughout the study, the consistent unit of reference is the ISV application package and not the company or organization. Data refers to the ISV applications—rather than to the ISV organizations that offer them. We began with a list of ISV applications that were identified to IDC by end-users as their top 3 HPC applications and, for each one, asked a series of questions represented in this part of the study. We were interested in understanding not only the age, condition and scalability of the applications, but the financial strength and human resources that stood behind them. Had we started instead with a list of ISV organizations, we would have reach equally legitimate conclusions but would not have shed as much light on the applications themselves. Since most of the technical computing ISV organizations are relatively small and have more than one application package per organization, the financial funds available to any single package is less than the total amount for each organization.

We assigned ISV software to industries based on the primary usage of the software. The "other" category serves as a catch basin primarily for general science codes and applications software used by only a small number of companies in the telecommunications, transportation/ logistics or entertainment industry. We selected the top five ISV vendors for each major industry based on a combination of revenues, number of customers, and number of licenses. We investigated revenue growth, but did not attempt to assess the profitability of ISV organizations in this study, in large part because past experiences taught us that ISV organizations frequently refuse to provide this information even when assurances of anonymity are given.

Although the vast majority (89%) of the ISVs represented in this study are for-profit businesses, the study uses the term "ISV organizations" because some respondents are public-sector entities.

Because there was no accurate, up-to-date information source covering this important ISV community, DARPA and the Council on Competitiveness asked IDC to create a directory of this information in conjunction with gathering information for this study. (Middleware data that was excluded from analysis is also included.) The directory—a first of its kind—is available through the Council's website (www.compete.org). DARPA, the Council on Competitiveness and IDC hope it will prove useful well beyond the scope of this study.

Study Limitations

While IDC aims to provide an accurate, comprehensive view of the subject being studied, certain limitations inevitably affect the results. We believe that the 54 ISV organizations and 110 applications software packages covered here represent the vast majority of those fitting the parameters established for this study, but there are likely others we have missed. Because our primary focus was on technical computing ISV applications software that all sectors, including industry, may access and use, software that is used only by government or academia is less well represented. For similar reasons, we surveyed codes that are used in the United States or on a worldwide basis and excluded codes used only outside of the U.S., in many cases only in a single country-of-origin. In most cases we ignored minor ISV codes with only a few users. Figures for revenues, market share, customer counts and other business data are as reported by the ISV respondents themselves. Few ISVs are public companies that are required to disclose this information on a broad basis.

STUDY RESULTS

ISV Technical Software Demographics

Independent software vendors serving the HPC market come in many sizes and from many geographic locations. They are a mix of private and public sector organizations. Their software applications target a wide range of industries and disciplines.

Q: What is the primary industry for this application?

IDC asked the primary industry of the respondents in the survey. (See Table 3.)

TABLE 3

Primary Industry of ISV Applications

Industry	Number of Applications	Percent
Auto/Aero	46	41.8%
Bio/Pharm	41	37.3%
Oil/Gas	8	7.3%
Other	15	13.6%
Total:	110	100.0%

Source: IDC, 2005

Two industries—the automotive/aerospace sector and the bio-pharmaceutical sector—together are the primary targets for four-fifths (78%) of the 110 ISV applications covered in this study (see Table 1.a). Each of these industries accounts for about 40% of the total, with the oil and gas industry running a distant third as a primary target for about 7% of the ISV applications. Automotive and aerospace firms were among the first industrial users of HPC, starting in the 1970s. They rely heavily on many of the same ISV software applications in their product design cycles—hence the common practice of grouping automotive and aerospace firms together in an HPC context. It is also important to note that the automotive and aerospace industries are among the largest and most visible sectors in a larger general manufacturing category. Others sectors within this category also use many of the same structural analysis, fluid flow and other applications to design and manufacture products. Although firms in the bio-pharmaceutical industry have used certain ISV applications for years, this number has grown substantially in the "post-genomic" period following the sequencing of the human genome. While the oil and gas industry has long-time users of large capability systems, they typically do more internal application development than other industries, relying less on the ISVs.

Two industries—the automotive/aerospace sector and the bio-pharmaceutical sector—together are the primary targets for four-fifths (78%) of the 110 ISV applications covered in this study.

Respondents from the "Other" category represented industries including Chemical, Electric/Utilities, Entertainment, Finance, General Sciences, Information Technology, Telecommunications, and Transportation. None of these had enough responses to warrant breaking them out separately.

Q: What type of organization controls the application?

The vast majority of the ISV applications (88%) are supplied by for-profit businesses. (See Table 4.) By contrast, only 7% are from universities, and an even smaller number (3%) are maintained primarily by the open-source community. This preponderance of for-profit applications means that most ISVs need to pursue profitable growth and cannot afford investments of time or money that are unlikely to contribute to this goal.

TABLE 4

Types of Organizations Supporting Each Application

Organization Type	Number of Applications	Percent
Company	97	88.2%
University	8	7.3%
National Laboratory	1	0.9%
Open-Source Community	3	2.7%
Other	1	0.9%
Total:	110	100.0%

Source: IDC, 2005

The open-source community is unlikely to become an important source for HPC application software in the future. IDC research shows that the open source model is difficult for business users. Although the software is usually available either free or at a nominal charge, the lack of long-term funding can limit the open source organization's ability to provide formal certification for the product, long-term support, and upgrades for the code. (Even for open source operating systems, users are finding it necessary to look to for-profit companies and/or computer vendors to provide for system reliability and long-term support.) IDC research shows that the open source model is difficult for businesses because there is no formal software certification and validation process, little ability to modify the software for your own environment, and no single responsible party when things go wrong.

The open source model is difficult for businesses because there is no formal software certification and validation process, little ability to modify the software for your own environment, and no single responsible party when things go wrong.

Q: Where is the organization that is directly responsible for this application located?

Eight out of nine applications (88%) are offered by organizations based in the United States, with Europe taking up most of the remainder.

TABLE 5

Organization's Geographical Region, by Application

Region	Number of Applications	Percent
U.S.	97	88.2%
Europe	10	9.1%
Canada	2	1.8%
Other	1	0.9%
Total:	110	100.0%

Source: IDC, 2005

We would not expect the exclusion from this study of ISV applications used only outside of the U.S. to alter the results in any fundamental way. Although non-American ISVs, especially in Europe, provide a number of crucial, widely used applications, the global HPC "ecosystem" and its ISV community continue to be concentrated heavily in the United States and to use English as their common language.

Q: Where is the parent organization that is associated with this application located?

In some cases, the organization directly responsible for the application is a subsidiary of a larger, parent organization. In other cases there is no higher parent; i.e., the organization directly responsible for the application is also the parent. When the question is framed in this way, U.S. dominance declines slightly and the role of Europe becomes more visible, with 20% of the total share of applications (see Table 6). Most of the shift from the U.S. to Europe comes from two companies with multiple applications in the study.

TABLE 6

Parent Company's Geographical Region, by Application

Region	Number of Applications	Percent
U.S.	84	76.4%

TABLE 6

Parent Company's Geographical Region, by Application

Region	Number of Applications	Percent
Europe	22	20.0%
Canada	1	0.9%
Other	3	2.7%
Total:	110	100.0%

Source: IDC, 2005

Business Models of ISV Organizations

Q: What type of ownership does your organization have?

The types of ownership (See Table 7) correlated well with the types of organizations in Table 4. It is worth repeating at this point that the great majority of ISVs are for-profit entities that need to pursue profitable growth and can not afford investments of time or money that are unlikely to contribute to this goal.

TABLE 7

Ownership Model of Organizations Supporting Each Application

Type of Ownership	Number of Applications	Percent
For profit company	92	88.5%
Not-for-profit company	2	1.9%
University	8	7.7%
National lab	1	1.0%
Other	1	1.0%
Total:	104	100.0%

Source: IDC, 2005

Q: What pricing model or models does your organization use?

As Table 8 illustrates, multiple pricing models exist among ISV organizations. The well-established models of charging by number of users (35%), by number of processors the application might be run on (27%) and by issuing site licenses for unrestricted use (12%) together constitute three-quarters of all responses. It will be interesting—and important—to see how ISV organizations grapple with current developments in HPC hardware systems. Will those pricing by the number of computers charge the same amount for a 10,000-processor server as for a 100-processor server? How will those pricing by the number of processors count the emerging wave of multi-core processors?

TABLE 8

Pricing Models of ISV Applications

Pricing Model	Number of Applications with that Pricing Model as an Option	Percent of Overall Responses
Charge by user count	37	34.9%
Charge by CPU	29	27.4%
Site license	13	12.3%
Charge per computer	7	6.6%
Free	7	6.6%
Share within community	3	2.8%
Other arrangement	10	9.4%
Total:	106	100.0%

Note: Multiple responses permitted.

Source: IDC, 2005

Although it is not the most commonly cited model, more than one in four responses cited per-CPU pricing as an option. As dual-core and multi-core chips enter the market, these ISVs will need to examine their pricing structures and decided whether their licenses should be assigned *per chip* (all cores on one chip are covered by the per-chip license) or *per core* (each core requiring its own license).

Figure 4 shows the industry breakdown of the licensing schemes from Table 8.

FIGURE 4

Application Pricing Models by Industry

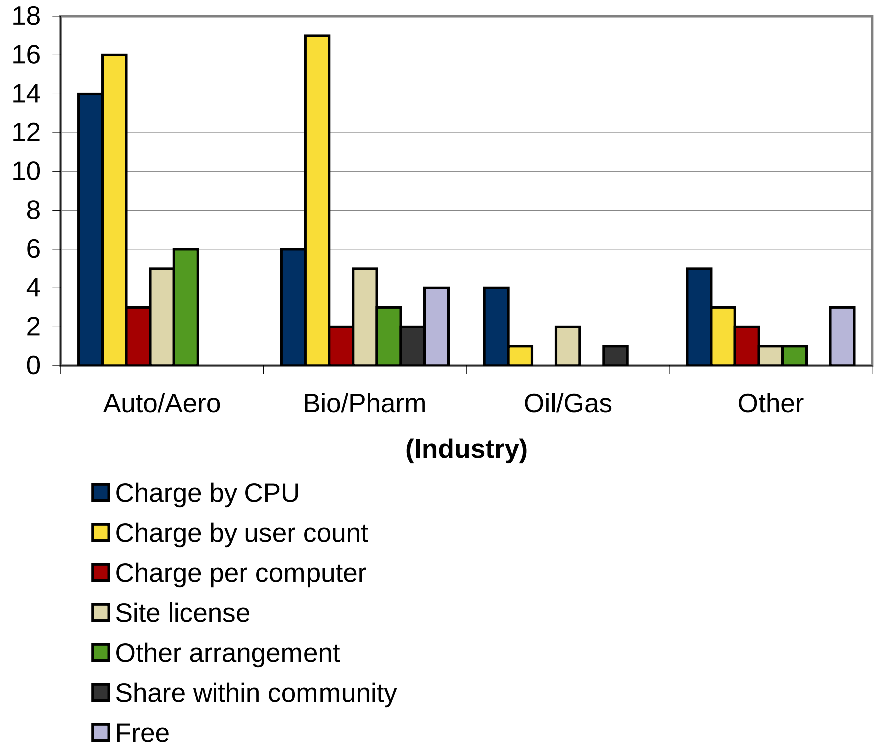

Note: Multiple responses permitted.

Source: IDC, 2005

Figure 4 illustrates that the automotive/aerospace and oil/gas industries are particularly reliant on application codes with per-CPU licensing schemes. Conversely, bio/pharmaceutical has a relatively low proportion. You can also see that the "share with community" and "free" responses are clustered in bio/pharm, possibly related to the preponderance of university customers in that sector.

Financial Condition of ISV Organizations

As for-profit companies (in most cases), ISV organizations 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 show that the HPC portion of the technical server market often represents less than five percent of their overall technical computing revenues, and in some cases this figure is less than one percent. Software development is expensive and labor-intensive. Even when business in their mainstream markets is doing well, ISVs typically cannot afford to spend the time and expense that would be needed to rewrite their applications software to meet the more-demanding requirements of HPC users. For business reasons, the needs of HPC users are often an important but secondary concern.

For business reasons, the needs of HPC users are often an important but secondary concern.

In former times, HPC hardware vendors operated on larger margins and invested substantial human and financial resources in collaborating with ISVs to improve the scalability and performance of applications 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 these investments. 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 to improve the scalability of ISV software that is crucial for the competitiveness of U.S. businesses.

Market forces alone will not address the gap between HPC users' needs and ISV application software capabilities.

Q: What are your organization's annual sales revenues for all products and services, not just ISV applications?

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 (35%). Few ISV applications (3%) are associated with organizations in the $25-50 million range. See Table 9.

TABLE 9

Number of ISV Applications and Companies by Total Company Revenue

	Companies		Applications	
Total Company Revenue	Number	Percent of Responses	Number	Percent of Responses
Under $1M	6	14.3%	9	9.8%
$1M to $5M	10	23.8%	27	29.3%
$5M to $10M	7	16.7%	11	12.0%
$10M to $25M	5	11.9%	10	10.9%
$25M to $50M	3	7.1%	3	3.3%
Over $50M	11	26.2%	32	34.8%
Total	42	100.0%	92	100.0%

Source: IDC, 2005

Q: How many employees are there in your organization?

A similar, related bifurcated pattern appears in company headcounts. (See Table 10.) Large numbers of ISV applications are offered by organizations with 25 of fewer employees (42%) and by organizations with more than 500 employees (32%), but only two of the applications (2%) come from organizations in the 250-500 employee range.

TABLE 10

Number of ISV Applications and Companies by Total Company Headcount

	Companies		Applications	
Total Company Headcount	Number	Percent of Responses	Number	Percent of Responses
Under 10	10	19.2%	14	13.3%
10 to 25	13	25.0%	30	28.6%
25 to 50	5	9.6%	9	8.6%
50 to 100	4	7.7%	5	4.8%
100 to 250	6	11.5%	11	10.5%
250 to 500	2	3.8%	2	1.9%
Over 500	12	23.1%	34	32.4%
Total:	52	100.0%	105	100.0%

Source: IDC, 2005

Tables 9 and 10 show that a substantial portion of the applications, approximately 40%, come from ISVs can be categorized as small businesses, according to standards set by the Small Business Administration. Regardless of growth or profitability, these companies could lack the resources that would be necessary to completely rewrite their applications.

Q: What are your organization's annual revenues for ISV applications alone?

Table 11 again shows a bifurcated pattern, this time for annual sales revenues of ISV applications alone, rather than for all products and services (Table 9). Nearly half (48%) of the applications come from organizations that qualify as small businesses with annual ISV applications revenues of less than $5 million. (The U.S. Small Business Administration frequently uses an upper limit of $6 million in annual revenue to define small businesses.) About one-fifth (18%) of the applications belong to organizations with $50 million or more in yearly applications sales. Only 13% of the ISV codes reside within organizations with applications sales in the $5-50 million range.

If these companies invest 10% of their yearly revenues for a particular application package into R&D, it would still be far from sufficient in most cases to do a full or even partial rewrite of the code. This gets worse, since the majority of the R&D has to go to testing and certification on a large number of different computers. 82% of the applications in this study would have less than $5 million a year for R&D per application package.

If these companies invest 10% of their yearly revenues for a particular application package into R&D, it would still be far from sufficient in most cases to do a full or even partial rewrite of the code.

TABLE 11

Annual Revenue of the ISV Applications

Revenue	Number of Applications	Percent
Under $500K	10	9.5%
$500K to $1M	24	22.9%
$1M to $5M	14	13.3%
$5M to $10M	5	4.8%
$10M to $25M	5	4.8%
$25M to $50M	4	3.8%
Over $50M	19	18.1%
Total:	81	100.0%

Source: IDC, 2005

Q: What has happened with your organization's revenues for ISV applications alone over the past five years?

Nearly three-quarters (71%) of the ISV applications reside in organizations whose revenue from ISV software alone grew 10% or more annually during the past five years. (See Table 12.) For 41% of the applications, the ISV organization standing behind them had yearly growth exceeding 25% in this period. Only 8% of the applications come from organizations with flat or declining growth; i.e., more than 90% are associated with growing organizations. It is worth noting again that revenue growth (at any rate) does not ensure profitability, a topic that was outside the scope of this study.

TABLE 12

Five-Year Growth of Total ISV Revenue, by Application

Growth Range	Number of Applications	Percent
Declined	3	3.5%
Flat	4	4.7%
Under 5%	10	11.6%
5% to 10%	8	9.3%
10% to 25%	26	30.2%
Over 25%	35	40.7%
Total:	86	100.0%

Source: IDC, 2005

Q: Are your organization's revenues for ISV applications alone growing or declining today?

Again here, the vast majority (94%) of respondents providing information about applications said that their organizations are experiencing moderate or high growth in ISV applications revenues today. (See Table 13.) Based on our findings, it is safe to conclude that ISVs serving the HPC sector typically are growing. But as we stated in the previous question, even when business in their mainstream markets is doing well, the study shows that 40% of ISVs cannot afford the expense that would be needed to rewrite their applications software to meet the demanding requirements of HPC users.

The study shows that 40% of ISVs cannot afford the expense that would be needed to rewrite their applications software to meet the demanding requirements of HPC users.

TABLE 13

Current Growth Rate of ISV Applications

Growth Range	Number of Applications	Percent
In decline	1	1.3%
No growth	4	5.2%
0% to 5% growth	35	45.5%
Growth over 5%	37	48.1%
Total:	77	100.0%

Source: IDC, 2005

One potential cause of this growth is the influx of clusters into the HPC market. IDC defines clusters used in technical markets as a set of independent computers combined into a unified system through systems software and networking technologies. In recent years, clusters using commodity off-the-shelf microprocessors have made HPC hardware pricing more attractive and have captured a substantial, growing share of the overall market for HPC computer servers.

Because of this favorable microprocessor pricing, end-users can afford to acquire more processors in a cluster than they could if they were to buy a traditional HPC system. But many ISV applications are no more capable of exploiting larger clusters than larger HPC systems of other kinds. In other words, end users may be purchasing more processors but not necessarily the ability to solve larger, more complex problems. In IDC's "High Performance Technical Computing Cluster Multi-Client Study," completed in 2004, "the ability to run larger problems" and "application availability" were among the top challenges cited by cluster users.

Yet the ability to buy more processors also tends to increase the cost of application software—i.e., ISV revenues—when ISVs charge on a per-processor basis or for per-computer licensing (the additional processors may be spread out over multiple cluster systems). The implication here is that ISVs may sometimes benefit financially from improved hardware price/performance, even without investing time and money to improve software scalability.

Q: What is the overall market share for this application?

One-quarter (24%) of the 78 ISV applications for which this information was known and disclosed have a commanding presence in their markets, defined here as more than a 50% market share. (See Table 14.) The remaining three-quarters (76%) of the applications are smaller players in target markets that may be dominated by a single competitor (no doubt one of the fortunate 24% in some cases) or fragmented among a greater number of participants. Attesting to the frequency and extremes of fragmentation, more than one of every four applications (28%) holds less than a 5% share of its market.

TABLE 14

Market Share of ISV Application

Market Share	Number of Applications	Percent
Under 5%	22	28.2%
5% to 10%	7	9.0%
10% to 25%	7	9.0%
25% to 50%	23	29.5%
Over 50%	19	24.4%

TABLE 14

Market Share of ISV Application

Market Share	Number of Applications	Percent
Total:	78	100.0%

Source: IDC, 2005

Q: How many clients are there for your application?

Three-quarters (75%) of the applications have 100 or more clients, and nearly half (47%) have at least 500 paying clients. (See Table 15.) The fact that there are fewer than 10 clients for 12% of the applications (one of every eight) is not necessarily a sign of weakness—a single client might in some instances be a large, multinational business. (Exploring the nature of the clients would have expanded this study beyond reasonable proportions for respondents.)

TABLE 15

Number of Clients for ISV Application

Range of Clients	Number of Applications	Percent
Under 10	9	12.3%
10 to 25	4	5.5%
25 to 50	4	5.5%
50 to 100	1	1.4%
100 to 250	9	12.3%
250 to 500	12	16.4%
Over 500	34	46.6%
Total:	73	100.0%

Source: IDC, 2005

Q: How many licenses are there for your application?

The same precaution applies to the number of licenses for the ISV applications: a single license may bring in considerable revenue or relatively little, and exploring this was beyond the scope of the current study. Prior knowledge and common sense tell us it is significant, however, that more than half (54%) of the 78 applications for which this information was disclosed command 1,000 or more licenses, with one in every seven (14%) having more than 10,000 licenses. (See Table 16.) Looking at the numbers of clients and licenses together underscores the broad impact of this ISV community—the number of licenses for the 78 applications represented in Table 15 is minimally 143,000 and might approach half a million.

TABLE 16

Number of Licenses for ISV Application

Range of Licenses	Number of Applications	Percent
Under 10	7	9.0%
10 to 100	11	14.1%
100 to 250	4	5.1%
250 to 500	6	7.7%
500 to 1,000	8	10.3%
1,000 to 10,000	31	39.7%
Over 10,000	11	14.1%
Total:	78	100.0%

Source: IDC, 2005

Top ISVs in Auto/Aero, Bio/Pharma and Oil/Gas

When this data is combined, we can analyze it to determine which are the most critical ISVs in each sector. Because no single measurement is adequate in describing an ISV, IDC used a multi-dimensional chart to plot them. Figure 5 shows the chart and its methodology.

Figure 5: Plotter for Importance Ranking of ISV Applications

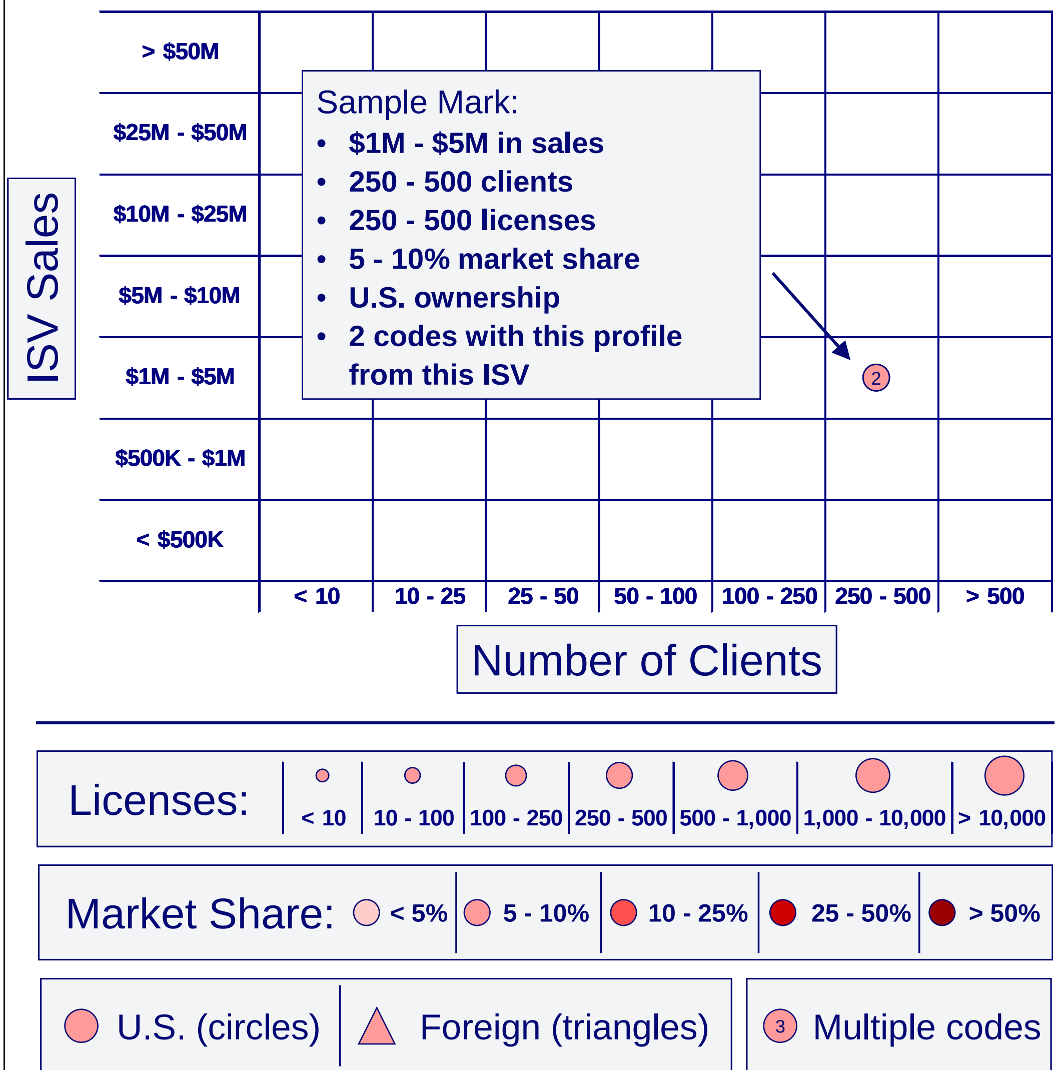

There are six dimensions measured on this plot:

- Vertical axis: Annual sales revenues for the code
- Horizontal axis: Number of clients for the code
- Size of mark: Number of licenses for the code
- Color (darkness) of mark: Market share for the code
- Shape of mark: Parent company is U.S.-owned (circles) or foreign-owned (triangles)

- Number with mark: When a number appears with the mark, a single ISV gave matching answers for multiple codes. This repetition is indicated with a number rather than multiple marks.

Marks were plotted any time data was collected for at least five of the six dimensions. When one dimension is missing, its lack is indicated. For example, if no market share data was collected, the circle is unfilled (no color). If no sales revenue data was collected, the mark appears below the chart in the correct column, but the proper row is unknown.

Figure 6 gives the plot for ISV applications in automotive and aerospace.

Figure 6: Plotter for Importance of Applications in Automotive/Aerospace

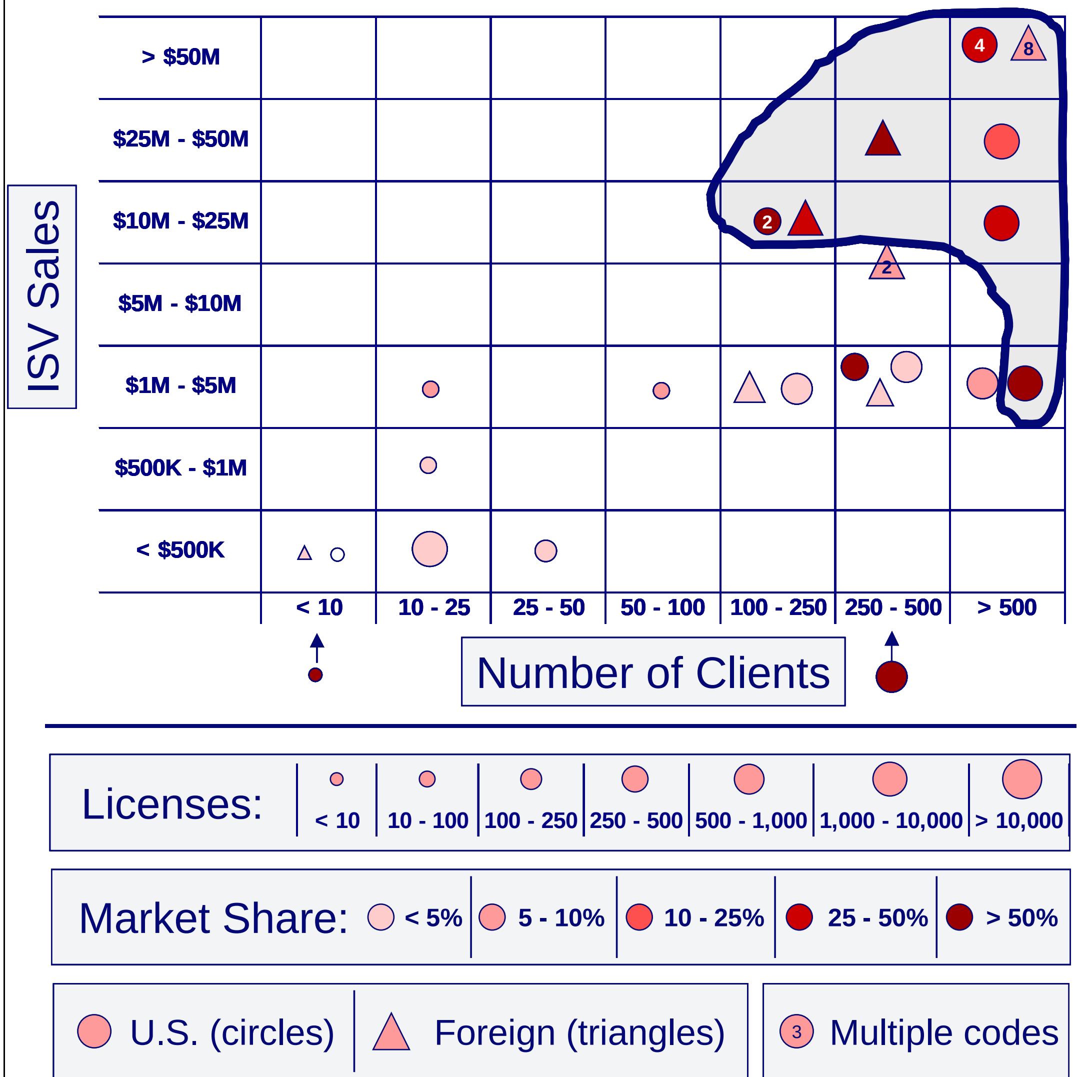

IDC considers the eight marks in the shaded region to represent the most significant applications among those that responded. (The mark below the chart in the 250 – 500 client column may be significant as well, depending on sales.) Of those eight, it is noteworthy that three (37.5%) represent codes from foreign-owned companies.

Figure 7 shows a similar plot for codes in biotechnology or pharmaceuticals.

Figure 7: Plotter for Importance of Applications in Biotechnology/Pharmaceutical

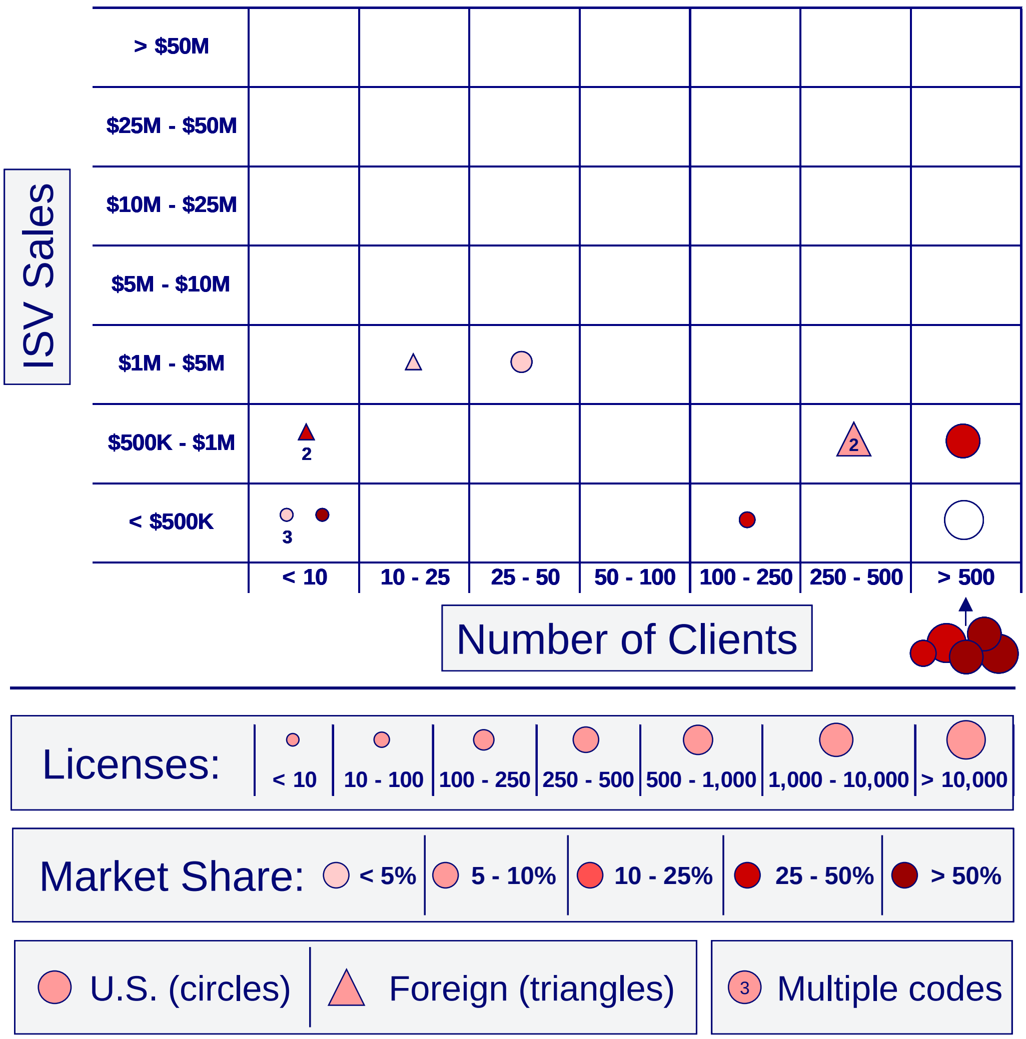

In Figure 7, no ISV in bio/pharma reported over $5M in annual sales for its application(s). However, there is a cluster of five otherwise-significant marks – over 500 clients, high market share, many licenses – for which the respondents did not share revenue data, and another mark in the same column indicates over 10,000 licenses, but for under $500,000 in revenue, with no market share data.

Given the rate at which applications are entering the bio/pharma market (data in the next section of this report), IDC believes that ISVs in this sector do not wish to advertise large revenues, for fear of attracting other entrants.

Figure 8 shows the importance plot for ISV application codes in Oil and Gas. In this case, the relatively smaller number of responses prevents detailed analysis or conclusions.

Figure 8: Plotter for Importance of Applications in Oil/Gas

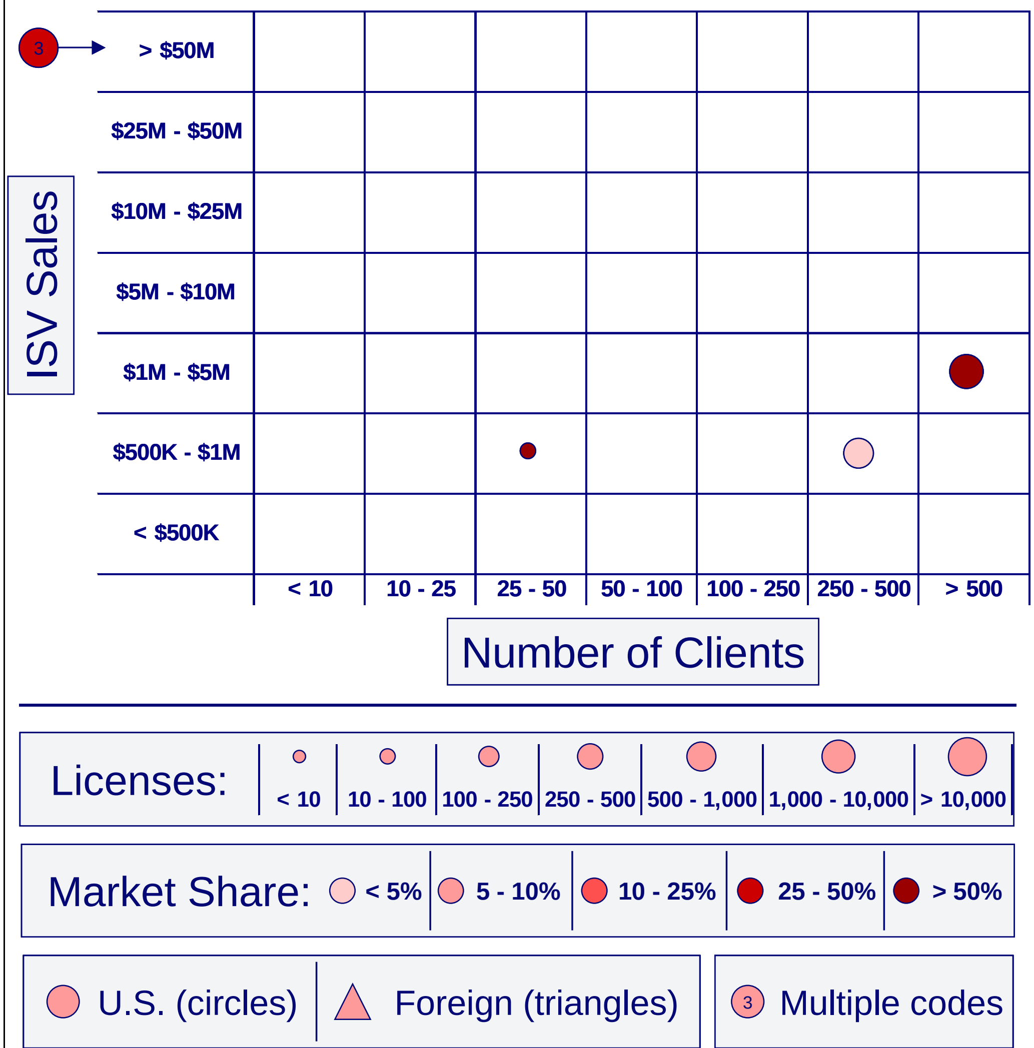

Age and Status of ISV Applications

Q: In what kind of organization did your application originate?

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% of the applications are based on open source software. (See Table 17.)

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.

TABLE 17

Original source of ISV Application

Source	Number of Applications	Percent
In-house developed	77	72.6%
University	18	17.0%
National lab	8	7.5%
Open community	3	2.8%
Total:	106	100.0%

Source: IDC, 2005

When this data is examined by industry, we see a continued affiliation of universities and bio/pharma. (See Figure 9.) 13 of the 18 codes that originated in universities are in the bio/pharma industry, accounting for one-third of the total number of applications in that sector.

FIGURE 9

Original Source of ISV Application by Industry

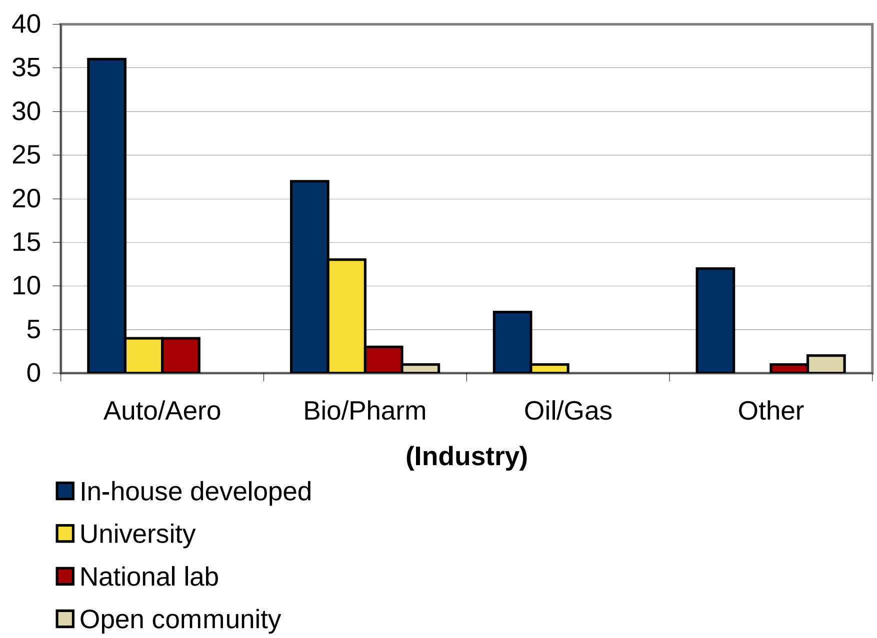

Source: IDC, 2005

Q: What is the age of your application?

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. (See Table 18.) Separate IDC research has indicated that some current ISV applications date back 20 years or more.

TABLE 18

Age of ISV Application (Since First Release)

Age of Code	Number of Applications	Percent
Less than 1 year	6	6.1%
1 to 2 years	5	5.1%
2 to 3 years	2	2.0%
3 to 5 years	13	13.3%
Over 5 years	72	73.5%
Total:	98	100.0%

Source: IDC, 2005

Q: When was the last major technical update made to your application?

Although the majority of the codes are older, the vast majority of ISV applications (85%) have had major technology updates within the past one or two years. (See Table 19.) Note that a major technology update is a substantial enhancement that typically does not involve fundamentally rewriting the code (the latter often results in a new name for the application).

TABLE 19

Time Since Last Major Technical Update

Time Since Update	Number of Applications	Percent
Less than 1 year	69	73.4%
1 to 2 years	11	11.7%
2 to 3 years	6	6.4%
3 to 5 years	5	5.3%
Over 5 years	3	3.2%

TABLE 19

Time Since Last Major Technical Update

Time Since Update	Number of Applications	Percent
Total:	94	100.0%

Source: IDC, 2005

Combined with other data gathered in this study, the age of these codes and the method in which they are updated lead to some dramatic conclusions. With growing revenues, most companies have the resources they need to provide updates to their products, in terms of new feature upgrades or other enhancements. However, the majority of the applications are legacy codes (older than five years), and they have not been re-written in this time period.

The majority of the applications are legacy codes (older than five years), and they have not been re-written in this time period.

Within the last five years there have been significant architectural changes in the market, such as the fast adoption of industry standard components, clustering, and new processor families (viz., 32-bit processors with 64-bit extensions). Over the next five years, other equally significant changes are imminent, such as multi-core processors and the goal of petascale systems.

Constraints to Advancing ISV Application Software

IDC concludes that ISVs will be slow to react to these changes. The scalability and usage of ISV applications is already limited on today's scalable systems. If this trend continues, it would likely take at least five years for most ISVs to rewrite their codes to take advantage of the petascale HPC computer servers the U.S. Government and hardware vendors plan to make available by the end of this decade.

If this trend continues, it would likely take at least five years for most ISVs to rewrite their codes to take advantage of the petascale HPC computer servers

The key limiting factors are the ISVs' resources and their motivation. As small companies, many of them cannot afford major rewrites, even in the face of growing revenues. Furthermore, even if they could afford it, most of them would not have the motivation to do so. High-performance computing is a small subset of the overall computing market, and ISVs might prefer to invest their limited R&D dollars in product improvements for the broader marketplace, especially if their HPC revenues will continue to grow regardless.

Current Scalability of ISV Applications

IDC market research shows that for most ISV organizations, HPC is only a small part— typically less than five percent—of the overall market for their applications. ISV applications frequently are designed primarily to run effectively on single processor desktop systems (workstations, PCs, Macs) and servers, and only secondarily to exploit multiprocessor HPC computer servers. As the tables in this section illustrate, the applications typically use 32 or fewer processors when running single HPC problems, and only a handful of the applications are able to "scale up" to exploit more than 128 processors for large single problems, despite the fact that the largest contemporary HPC computer servers may have up to 10,000 processors.

For U.S. industries that need to out-compete their non-U.S. competitors by out-computing them, the current constraints on scalability has become a limiting factor (see Table 19). 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.

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

The current situation is not the fault of the ISV organizations. As for-profit businesses (for the most part), they are pursuing the economic models they need to follow to remain profitable and cannot afford to make investments that are unlikely to contribute to that profitability.

Although processors are not the only components of HPC computer servers that help accelerate problem-solving speed, counting how many processors an application can exploit is the most convenient measure of the application's scalability. It is also worth keeping in mind that the underlying problems ISV applications address vary greatly in complexity, and for this reason it is far easier to scale up some applications than others.

It is worth repeating the point about clusters entering the market. From a hardware perspective only, it is easier and considerably cheaper to scale the number of processors in a distributed-memory cluster than it is to design larger shared-memory SMP systems. Through the industry's adoption of clusters, some of the scalability burden has shifted from the hardware and operating system provider to the application provider, who must now adapt the code to scale well on loosely coupled, commodity components, rather than specialized HPC architectures. Since this change in the market has taken place over the past five years, and most ISV codes are older than that, it should not be surprising to see codes fall short on scalability metrics. The cost and difficulty of rewriting HPC application software, combined with the secondary importance of the HPC market suggest that ISVs will not make this investment without external support. Unless this investment is made, ISV codes will lag even further as petascale systems are introduced into the market.

Q: How many processors does your application typically use for single problems?

About one-quarter of the applications (24.4%) typically run on only a single processor of an HPC computer server, and fewer than 7% use more than 128 processors. (See Table 20.)

TABLE 20

Typical Number of Processors the ISV Applications Use for Single Jobs

CPU Range	Number of Applications	Percent
1	19	24.4%
2-8	25	32.1%
9-32	20	25.6%
33-128	9	11.5%
129-1024	4	5.1%
Unlimited	1	1.3%
Total:	78	100.0%

Source: IDC, 2005

The significant number of single-processor usage again raises the question of how dual-core and multi-core CPUs will be used, licensed and charged. If a usage model emerges in which each core is assigned its own jobs, it can lead to sizable increases in software license fees. The per processor pricing models makes great business sense for ISV companies, but could create limitations in scaling combined with growing software costs as processor core counts grow.

Ability To Improve Scalability of ISV Applications

ISV applications that 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"). Conversely, some of the most complex and consequential problems are far more difficult to scale up. In some cases, applications of crucial, competitive importance to industry can exploit only a handful of processors and would require a fundamental rewriting to advance beyond this state.

Applications of crucial, competitive importance to industry can exploit only a handful of processors and would require a fundamental rewriting to advance beyond this state.

Q: Do you know how to scale your application to hundreds of processors?

Fewer than half (46%) of the ISV applications scale to hundreds of processors today. (See Table 23.) 37% feel that they could scale but it would be hard or they have no plans to scale to this level. Responses to later questions make it clear that a lack of interest is not a major factor here.

TABLE 21

Ability of Application to Scale to Hundreds of Processors

Status	Number of Applications	Percent
Already does	43	46.2%
Yes, and plans in place	13	14.0%
Yes, but hard	20	21.5%
Yes, but no plans	14	15.1%
No, not possible	3	3.2%
Total:	93	100.0%

Source: IDC, 2005

Q: Do you know how to scale your application to thousands of processors?

Not surprisingly, when the scalability goal is raised an order of magnitude to thousands of processors, the percentage of those claiming this ability today declines markedly. (See Table 24.) Less than one third of applications (32%) scales to thousands of processors today. 44% feel that they could scale but it would be hard or they have no plans to scale to this level, and 12% indicated it would not be possible.

TABLE 22

Ability of Application to Scale to Thousands of Processors

Status	Number of Applications	Percent
Already does	28	31.8%
Yes, and plans in place	10	11.4%
Yes, but hard	19	21.6%
Yes, but no plans	20	22.7%
No, not possible	11	12.5%
Total:	88	100.0%

Source: IDC, 2005

Q: Do you know how to scale your application to tens of thousands of processors?

When we asked whether the applications can scale today to tens of thousands of processors, the set of application packages claiming this ability dropped significantly to 19%, less than one in five. (See Table 25.) Although there was not a significant increase in the number of "not possible" responses, the number of applications with no immediate plans to scale to this level increased to 60%, and 14% indicated that it would be impossible. This again suggests that ISV application software will not be able to take advantage of petascale systems when they are delivered.

TABLE 23

Ability of Application to Scale to Tens of Thousands of Processors

Status	Number of Applications	Percent
Already does	17	19.3%
Yes, and plans in place	6	6.8%
Yes, but hard	19	21.6%
Yes, but no plans	34	38.6%
No, not possible	12	13.6%
Total:	88	100.0%

Source: IDC, 2005

Willingness To Collaborate and Preferred Partners

For the substantial percentage of respondents who said they know how to make their applications more scalable, it is important to determine whether they are willing to make the needed effort (presumably those with plans in place are), what additional ingredients would be needed to accomplish this goal, whether they are willing to collaborate with outside parties and, finally, what types of outside partners they would prefer to collaborate with.

Q: Are you willing to improve your application?

Nearly all (98%) of the respondents said they are willing to improve the scalability of their applications. (See Table 26.) Almost as many (86%) said the work has already begun, though this says nothing about how fast it is proceeding or how far it has gotten. About one in eight (12%) said the expense prevents them from improving their applications.

TABLE 24

ISV Willingness to Improve Application (by Application Count)

Willingness	Number of Applications	Percent
Yes, already underway	89	86.4%
Yes, but it's too expensive	12	11.7%
Maybe/Uncertain	2	1.9%
Never or very hard	0	0.0%
Total:	103	100.0%

Source: IDC, 2005

Q: What additional things would you need to improve your application?

ISVs need more money for R&D investments, a stronger business case or more customers to offset investment cost, and more qualified staff and/or access to outside experts to improve their applications. (See Table 27.) It is useful to note that 15% of the responses pointed to a lack of external expertise, and about 10% to the need to re-think their software code, a process that would presumably result in a fundamental re-writing of the software.

ISVs need more money for R&D investments, a stronger business case or more customers to offset investment cost, more qualified staff and/or access to outside experts.

TABLE 25

Key Factors Needed for ISVs to Improve Applications

Factor	Number of Applications	Percent of Responses
Money / investments	50	24.9%
Business case / many customers	39	19.4%
Internal people or experts	35	17.4%
External tech expertise	30	14.9%
Partnerships to share costs & risks	28	13.9%
A whole new approach to their code	19	9.5%
Total:	201	100.0%

TABLE 25

Key Factors Needed for ISVs to Improve Applications

Factor	Number of Applications	Percent of Responses

Note: Multiple responses permitted.

Source: IDC, 2005

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

"The requirement from government or industry for this advanced functionality is not present"

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

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

Q: Are you willing to develop partnerships to improve your application?

Five out of six (83%) of the respondents declared themselves open to developing partnerships with other organizations, and when the "maybe" responses are added in, the percentage climbs to 98. Only two of 104 respondents provided an outright no to this question. (See Table 28.)

TABLE 26

ISV Willingness to Develop Partnerships to Improve Applications

Willing to Partner?	Number of Applications	Percent
Yes	86	82.7%
Maybe	16	15.4%
No	2	1.9%
Total:	104	100.0%

Source: IDC, 2005

Q: Are you willing to work with the U.S. Government to improve your application?

In past studies, respondents have sometimes staunchly resisted 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 at least open to the possibility of working with the government, and 93 of them (89%) gave a definite yes. (See Table 29.)

TABLE 27

ISV Willingness to Work with U.S. Government

Willing to Work with Government?	Number of Applications	Percent
Yes	93	89.4%
Maybe	11	10.6%
No	0	0.0%
Total:	104	100.0%

Source: IDC, 2005

Q: What types of partners would help you most to improve your application?

The ISV organizations preferred other code developers (25%), government labs (25%) and universities (22%) as partners for helping to improve their applications. (See Table 30.)

TABLE 28

Most Helpful Types of Partners for ISV Applications

Partner	Number of Applications for which Partner Would Be Useful	Percent of 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	100.0%

TABLE 28

Most Helpful Types of Partners for ISV Applications

Partner	Number of Applications for which Partner Would Be Useful	Percent of Responses

Note: Multiple responses permitted.

Source: IDC, 2005

When we look at the partnering preferences in relation to the industries that the ISV applications target, significant differences emerge. (See Figure 10.) For ISV respondents targeting the automotive/ aerospace sector, government labs are the most preferred partners, and the "other code developers" category is tied for third. The order is reversed for applications software used in the bio-pharmaceutical market. Oil and gas ISV companies ranked all types of partnerships equally desirable.

FIGURE 10

Most Helpful Types of Partners for ISV Applications by Industry

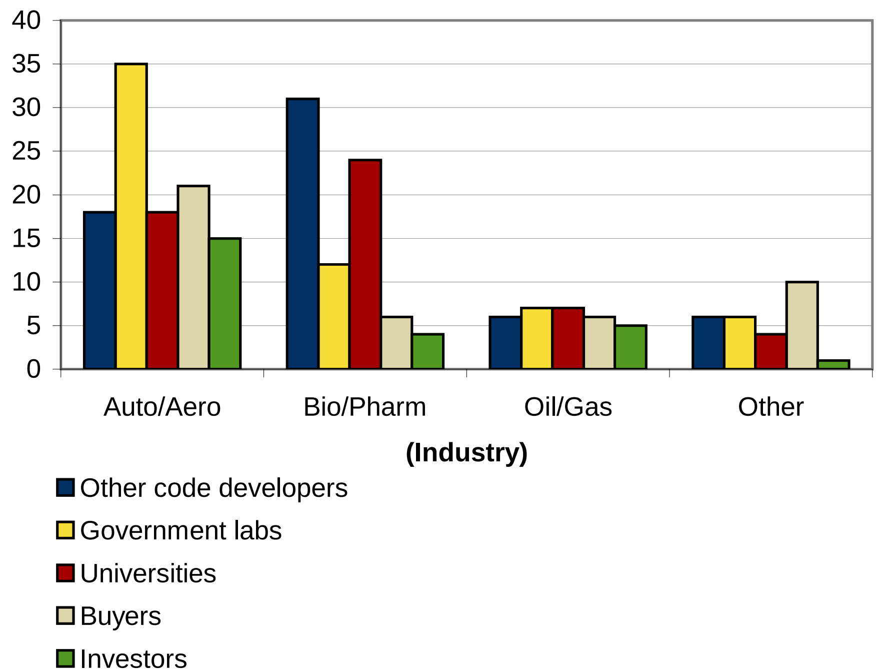

Note: Multiple responses permitted.

Source: IDC, 2005

CONCLUSIONS

In recent years, an alarming gap has developed 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.

In recent years, an alarming gap has developed 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 technical HPC market often represents less than five percent of their overall revenues, and in some cases this figure is less than one percent. As implied earlier, 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.

Although the ISVs are making rational business choices, the implications for U.S. competitiveness are sobering. 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, important problems, such as those encountered in designing new cars and airplanes and pharmaceuticals, or extracting more oil from reservoirs, often cannot be solved today in reasonable timeframes, or possibly at all. 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.

In former times, HPC hardware vendors operated on larger margins and invested substantial human and financial resources in collaborating with ISVs to improve the scalability and performance of applications 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 these investments, and U.S. businesses historically have not funded R&D for ISV application software. Even given proper investment, many ISVs cited a need for either internal or external technical expertise to improve their applications. Money alone cannot solve the problem.

Market forces alone will not address the serious gap between HPC users' needs and ISV application software capabilities. Market forces need to be supplemented with external funding and/or expertise to improve the scalability of ISV software that is needed for improving the competitiveness of U.S. businesses. Without proper funding or a more compelling business case, ISVs are unlikely to rewrite their codes to accommodate current scaling limitations, much less take advantage of petascale systems when they are available.

Appendix A: Analysis Of ISV Application Usage In Industry, Government and Academia

For each application code, the ISV was also asked what their mix of sales is for that code, as a percentage between business, government, and universities. To visualize the responses to that question, IDC used the triangular chart depicted in Figure A1.

Figure A1: Chart to Show Percentage of Business, Government, and University Usage of HPC Applications by Industry

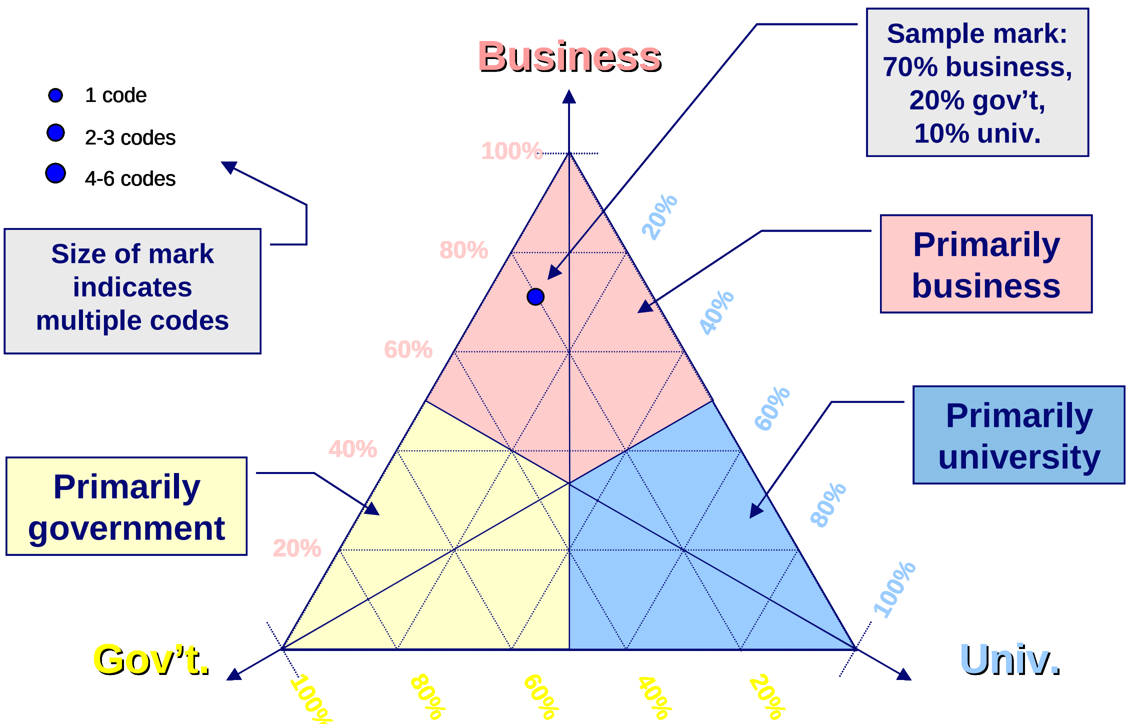

There are three dimensions to the figure. The vertical axis plots the percentage of sales that goes to businesses. The bottom (base) of the triangle represents 0% business, and as you move up the chart, you pass through horizontal hash marks that represent 20%, 40%, 60%, and 80%, until you reach 100% at the top vertex. These percentages are shown along the left side, increasing upwards.

Similarly, percentages of university sales begin at 0% on the left side of the triangle. As you progress down to the right, you pass through increasing percentages (indicated to the right of the triangle, next to the corresponding hashes) until you reach 100% at the bottom-right vertex. Percentages of government sales begin at 0% along the right side, increasing down and to the left, reaching 100% at the bottom-left vertex.

In this way, any mix of sales can be plotted on the chart, provided that the respondent gave numbers that added to 100%. The sample mark in Figure A1 provides an example. Moving vertically, this mark is above the 60% line, halfway to 80%, indicating 70% of sales to businesses. Moving from top-right to bottom-left, the mark is on the 20% government line. And finally, it lies halfway between the 0% line (left side) and 20% line for universities. This mark therefore represents 70% business sales, 20% government sales, and 10% university sales. Whenever a single ISV provides the same data for multiple codes, this is reflected with a larger mark.

Once the data is plotted in this way, the casual observer can see at a glance the overall mix of sales. Any marks in the top (pink) section are primarily sales to business. Marks in the left (yellow) section are primarily to government. Marks in the right (blue) section are primarily to universities. The central intersection represents a point that is one-third to each.

Each of these primary regions is also split in half, and which half of the region the mark is in shows which the second-most significant sales category is. For example, our sample mark in Figure 1 is in the pink region (most of sales go to business). The fact that it lies left of the centerline shows that it indicates more government sales (second most) than university sales (third most).

Figure A2 shows the mix of sales for ISV applications in the automotive and aerospace industries.

Figure A2: Percentage of Business, Government, and University Usage of HPC Applications in Auto/Aero

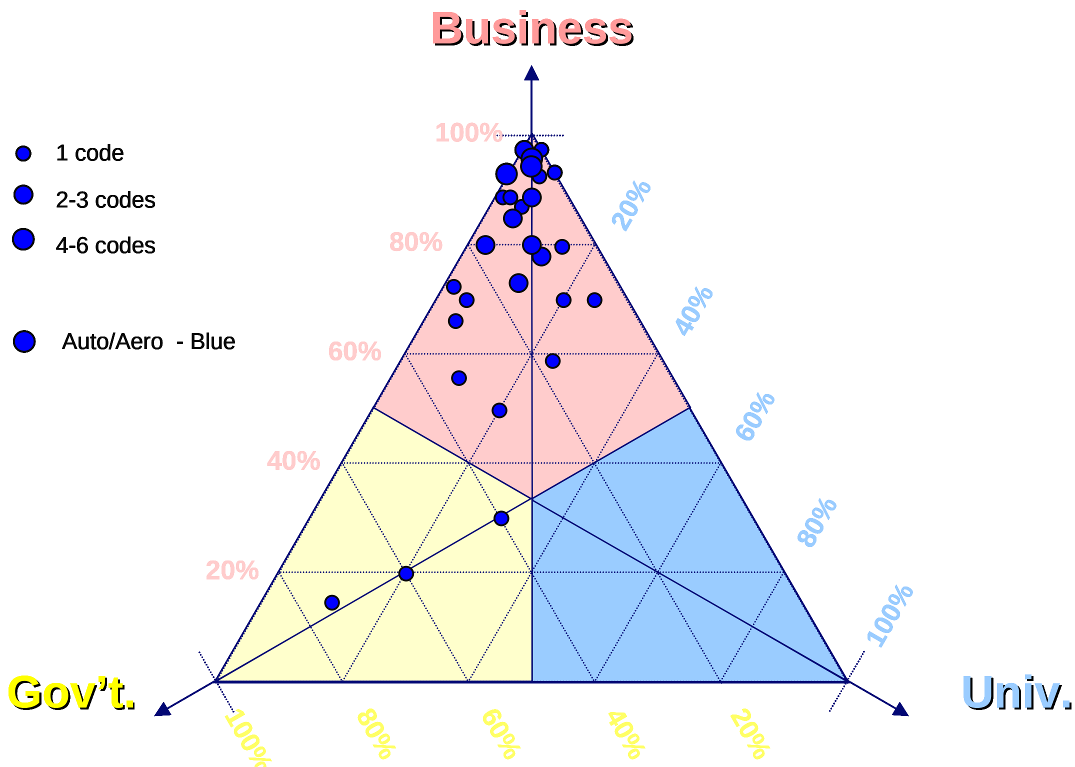

Figure A2 clearly shows that the predominant customers for ISV software in automotive and aerospace are businesses. Only three codes were sold primarily to government, and none to universities. Overall, more sales went to government than universities, even for those applications primarily sold to business.

Figure A3 shows the mix of sales for ISV applications in the biotechnology and pharmaceutical industries.

Figure A3: Percentage of Business, Government, and University Usage of HPC Applications in Bio/Pharma

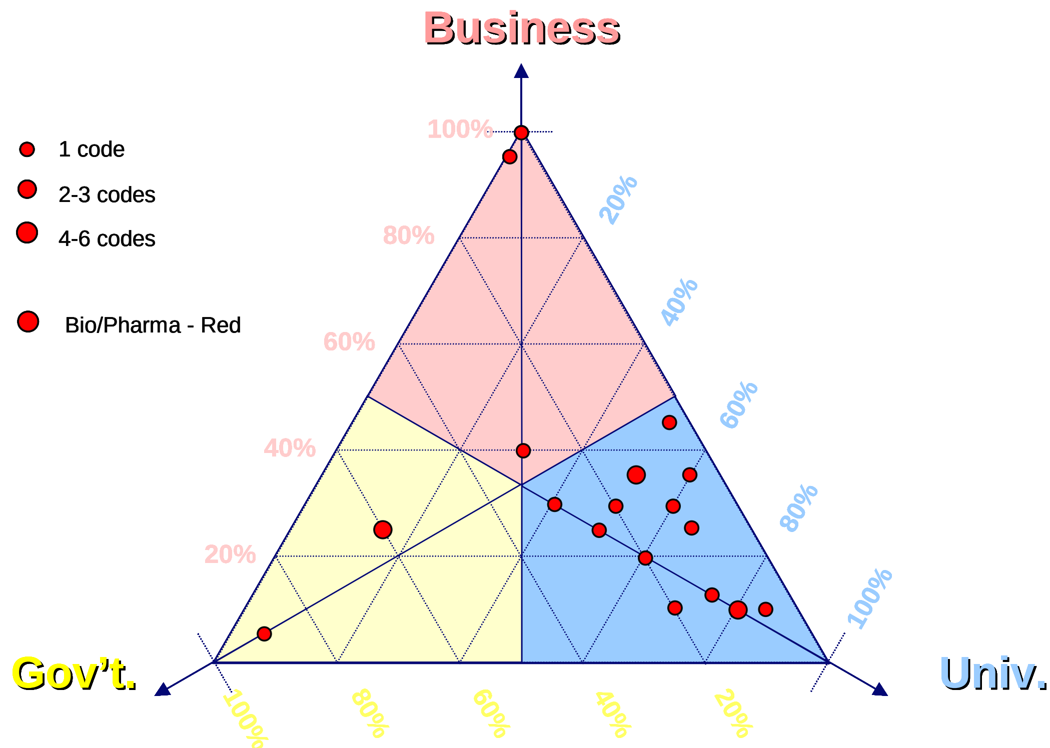

Figure A3 shows a dramatically different picture for bio/pharma than we saw for auto/aero. Here universities are the dominant customers. The results also clearly show more business sales than government sales.

Figure A4 shows the mix of sales for ISV applications in the oil and gas industry.

Figure A4: Percentage of Business, Government, and University Usage of HPC Applications in Oil/Gas

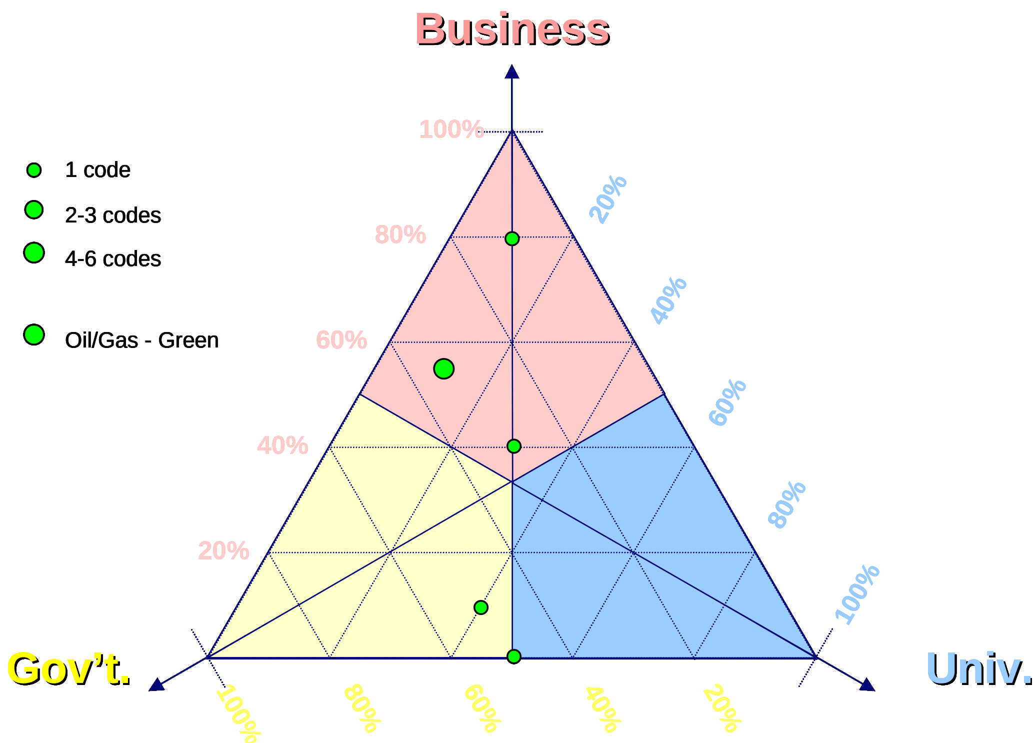

As with the auto/aero industries, oil/gas HPC applications are sold primarily to businesses, with more government sales than university sales.

Figure A5 combines Figures A2, A3, and A4 and also includes applications for other industries.

Figure A5: Percentage of Business, Government, and University Usage of HPC Applications in All Industries

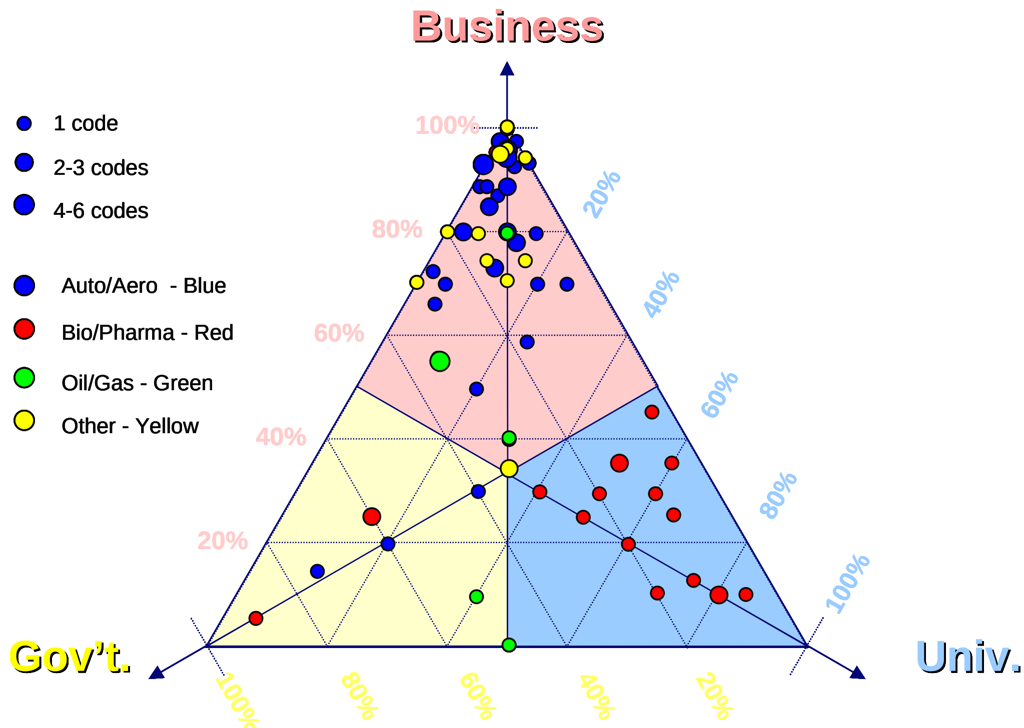

All of the industries other than bio/pharma follow the same pattern as auto/aero: business is the most significant sector for sales, with more government than university. With its predominant university sales, bio/pharma is the exception. Only six codes total, three of which are from the same ISV, have more than 50% sales to government.

Appendix B: List of the Application Packages AND ISVs in the Study

TABLE B1

Directory of ISV Applications and Organizations Included in the Study

ISV Package Name	ISV Supplier / Company Name	Supplier Type	Company Location	Primary Industry	Primary Regions the Code Is Sold In	Original Source of Code
ACUSOLVE	ACUSIM	Company	U.S.	Auto/Aero	U.S., Japan, Europe	In-house
Adams	MSC Software	Company	U.S.	Auto/Aero	Worldwide	In-house
ADF	Scientific Computing & Modelling	Company	Europe	Bio/Pharm	Worldwide	University
AMBER	Scripps Research Institute	University	U.S.	Bio/Pharm	U.S.	University
ANSYS	Ansys	Company	U.S.	Auto/Aero	North America, Europe, Japan	In-house
APBS	Washington Univ., St. Louis	University	U.S.	Bio/Pharm	U.S., Europe	University
ArcGIS Server	ESRI	Company	U.S.	Oil/Gas	Worldwide	In-house
ArcIMS	ESRI	Company	U.S.	Oil/Gas	Worldwide	In-house
ArcSDE	ESRI	Company	U.S.	Oil/Gas	Worldwide	In-house
Aspen Plus	AspenTech	Company	U.S.	Auto/Aero	Worldwide	University
Autoform	Autoform Engineering USA Inc.	Company	U.S.	Auto/Aero	Worldwide	In-house
AVS/Express	Advanced Visual Systems Inc.	Company	U.S.	Oil/Gas	Worldwide	In-house
AVS5	Advanced Visual Systems Inc.	Company	U.S.	Oil/Gas	Worldwide	In-house
BAND	Scientific Computing & Modelling	Company	Europe	Bio/Pharm	Worldwide	University
Bioconductor	Multi-university	Open-source community	U.S.	Bio/Pharm		Open-source community

TABLE B1

Directory of ISV Applications and Organizations Included in the Study

ISV Package Name	ISV Supplier / Company Name	Supplier Type	Company Location	Primary Industry	Primary Regions the Code Is Sold In	Original Source of Code
Biofacet	Gene-IT	Company	Europe	Bio/Pharm	U.S., Europe	National lab
BioInfomatIQ	Proteome Systems Inc.	Company	U.S.	Bio/Pharm	Worldwide	In-house
BLACS	University of Tennessee	University	U.S.	Other	Worldwide	Open-source community
BLAS	University of Tennessee	University	U.S.	Other	Worldwide	Open-source community
BLAST	Blast Inc.	Company	U.S.	Bio/Pharm	North America	In-house
BLAT	Kent Informatics	Company	U.S.	Bio/Pharm	U.S., Japan, Europe	In-house
Calibre	Mentor Graphics	Company	U.S.	Auto/Aero	Worldwide	In-house
CASE	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
CFD++	Metacomp Technologies	Company	U.S.	Auto/Aero	U.S., Europe, Japan	In-house
CFD-ACE	ESI US R&D	Company	U.S.	Other	Worldwide	In-house
CFD-FASTRAN	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	In-house
CFX-5	Ansys	Company	Canada	Auto/Aero	Worldwide	In-house
CHARMM	Scripps Research Institute	University	U.S.	Bio/Pharm	U.S., Europe	University
Checkmate	Mentor Graphics	Company	U.S.	Auto/Aero	Worldwide	In-house
Chemkin	Reaction Design	Company	U.S.	Auto/Aero	U.S., Japan, Europe	National lab
Cobalt	Cobalt Solutions	Company	U.S.	Auto/Aero	U.S., Europe	National lab
COMAZ	3DGeo Development Inc.	Company	U.S.	Oil/Gas	North America, Europe, West Africa	In-house

TABLE B1

Directory of ISV Applications and Organizations Included in the Study

ISV Package Name	ISV Supplier / Company Name	Supplier Type	Company Location	Primary Industry	Primary Regions the Code Is Sold In	Original Source of Code
DataStage Enterprise Edition Orchestrate	Ascential Systems	Company	U.S.	Other	U.S., Europe, Asia-Pacific	In-house
emu	CoBegin Inc.	Company	U.S.	Auto/Aero	U.S.	National lab
EON	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
FASTA	Univ. of Virginia	University	U.S.	Bio/Pharm	U.S., Europe, Japan	University
FEAP	Engineering Mechanics Research Corp.	Company	U.S.	Auto/Aero	Worldwide	In-house
FEKO	EMSS	Other	Other	Auto/Aero	U.S., Europe	In-house
FIELDVIEW	Intelligent Light	Company	U.S.	Auto/Aero	US, Japan	In-house
FILTER	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
FIPER	Engineous Software Inc.	Company	U.S.	Auto/Aero	Worldwide	In-house
FLOW-3D	Flow Science	Company	U.S.	Auto/Aero	North American, Japan, Europe	National lab
FRED	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
GAMESS	Iowa State Univ.	University	U.S.	Bio/Pharm	Worldwide	National lab
Gaussian 03	Gaussian Inc.	Company	U.S.	Bio/Pharm	U.S., Europe, Japan	In-house
Gaussian 94	Gaussian Inc.	Company	U.S.	Bio/Pharm	Worldwide	
Gaussian 98	Gaussian Inc.	Company	U.S.	Bio/Pharm	Worldwide	
GrailEXP	Oak Ridge National Laboratories	National lab	U.S.	Bio/Pharm	Worldwide	National lab

TABLE B1

Directory of ISV Applications and Organizations Included in the Study

ISV Package Name	ISV Supplier / Company Name	Supplier Type	Company Location	Primary Industry	Primary Regions the Code Is Sold In	Original Source of Code
GT-Power	Gamma Technologies	Company	U.S.	Auto/Aero	U.S., Europe, Japan	In-house
HMMER	Washington Univ., St. Louis	Open-source community	U.S.	Bio/Pharm		In-house
Houdini	Side Effects Software	Company	Canada	Other	North America, Japan, Europe	In-house
HYCOM	Miami University	University	U.S.	Oil/Gas	U.S., Europe, South America	University
HyperChem	Hypercube Inc.	Company	U.S.	Bio/Pharm	U.S., Europe	In-house
IC Verify	Mentor Graphics	Company	U.S.	Auto/Aero	Worldwide	In-house
ICEM CFD	Ansys	Company	U.S.	Auto/Aero	U.S., Europe, Japan	In-house
ImageGear Professional	AccuSoft	Company	U.S.	Other	U.S.	In-house
iSIGHT	Engineous Software Inc.	Company	U.S.	Auto/Aero	Worldwide	In-house
MEDINA	T-Systems International GmbH	Company	Europe	Auto/Aero	Europe, Japan, Asia Pacific	In-house
MOPAC	Stewart Computational Chemistry	Company	U.S.	Bio/Pharm	Worldwide	University
MySQL	MySQL Inc.	Company	U.S.	Other	Worldwide	In-house
Nastran / LS-DYNA	MSC Software	Company	U.S.	Auto/Aero	Worldwide	In-house
Net Vault	BakBone Software	Company	U.S.	Other	Worldwide	In-house
NISA / 3D-FLUID	Engineering Mechanics Research Corp.	Company	U.S.	Auto/Aero	Worldwide	In-house
Nisa Family	Engineering Mechanics Research Corp.	Company	U.S.	Auto/Aero	Worldwide	In-house

TABLE B1

Directory of ISV Applications and Organizations Included in the Study

ISV Package Name	ISV Supplier / Company Name	Supplier Type	Company Location	Primary Industry	Primary Regions the Code Is Sold In	Original Source of Code
OEChem	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
OGHAM	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
Omega	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	University
OpenPBS	Altair Engineering	Company	U.S.	Other	Worldwide	National lab
OptiStruct	Altair Engineering	Company	U.S.	Auto/Aero	Worldwide	In-house
PAM-CRASH	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	In-house
PAM-FLOW	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	University
PAM-GEN	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	University
PAM-MEDYSA	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	
PAM-OPT	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	In-house
PAM-SAFE	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	In-house
PAM-STAMP	ESI US R&D	Company	U.S.	Auto/Aero	Worldwide	In-house
PATRAN	MSC Software	Company	U.S.	Auto/Aero	Worldwide	In-house
Phlex	Altair Engineering	Open-source community	U.S.	Auto/Aero	Worldwide	In-house
PRISM	Advanced Systems Controls	Company	U.S.	Other	North America, Asia Pacific	In-house
Project Alexandria - ArcView	ESRI	Company	U.S.	Oil/Gas	Worldwide	In-house
ProteomIQ Access	Proteome Systems Inc.	Company	U.S.	Bio/Pharm	Worldwide	In-house
PSS/Adept	PTI	Company	U.S.	Other	U.S.	In-house
PSS/E	PTI	Company	U.S.	Other	Worldwide	In-house

TABLE B1

Directory of ISV Applications and Organizations Included in the Study

ISV Package Name	ISV Supplier / Company Name	Supplier Type	Company Location	Primary Industry	Primary Regions the Code Is Sold In	Original Source of Code
PYOCHEM	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
QUAC PAC	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
RADIOSS	MECALOG S.A.R.L.	Company	Europe	Auto/Aero	U.S., Europe, Japan	In-house
RADIOSS-CFD	MECALOG S.A.R.L.	Company	Europe	Auto/Aero		
ROCS	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
SAMCEF Linear	SAMTECH sa	Company	Europe	Auto/Aero	Europe	In-house
SAMCEF Mecano	SAMTECH sa	Company	Europe	Auto/Aero	Europe	In-house
SAMCEF Thermal	SAMTECH sa	Company	Europe	Auto/Aero	Europe	In-house
SHAPE	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
SMACK	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
SMART	Daylight Chemical Info. Systems	Company	U.S.	Other	Worldwide	In-house
SMILES	Daylight Chemical Info. Systems	Company	U.S.	Other	Worldwide	In-house
Spartan	Wavefunction Inc.	Company	U.S.	Bio/Pharm	U.S., Europe, Japan	University
SpeedUp	AspenTech	Company	U.S.	Auto/Aero	Worldwide	University
SPS Cross Match	Southwest Parallel Software	Company	U.S.	Bio/Pharm	U.S., Japan, Asia Pacific	University
SPS Phrap	Southwest Parallel Software	Company	U.S.	Bio/Pharm	U.S., Japan, Asia Pacific	University
SPS SWAT	Southwest Parallel Software	Company	U.S.	Bio/Pharm	U.S., Japan, Asia Pacific	University

TABLE B1

Directory of ISV Applications and Organizations Included in the Study

ISV Package Name	ISV Supplier / Company Name	Supplier Type	Company Location	Primary Industry	Primary Regions the Code Is Sold In	Original Source of Code
SuperForge / SuperForm	MSC Software	Company	U.S.	Auto/Aero	North America, Europe, Japan	In-house
SZYBKI	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
TEA Mecano	SAMTECH sa	Company	Europe	Auto/Aero	Europe	In-house
Time Navigator	Atempo	Company	U.S.	Other	Worldwide	In-house
TimeLogic DeCypher Biocomputing Solution	Active Motif, Inc.	Company	U.S.	Bio/Pharm	U.S., Europe, Japan	In-house
VIDA	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
VisiQuest / Khoros	AccuSoft	Company	U.S.	Auto/Aero	U.S., Europe	In-house
WABE	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	In-house
XtremeAutoRoute	Mentor Graphics	Company	U.S.	Auto/Aero	U.S., Europe, Japan	In-house
ZAP	OpenEye Scientific Software	Company	U.S.	Bio/Pharm	U.S., Japan	University

Source: IDC, 2005

IDC
Analyze the Future

WHITE PAPER

Council on Competitiveness Study of ISVs Serving the High Performance Computing Market: Part B — End-User Perspectives

Earl Joseph, Ph.D.
Addison Snell
Christopher G. Willard, Ph.D.
Suzy Tichenor, Council on Competitiveness
Steve Conway, Conway Communications

February 2006

Commissioned by the Council on Competitiveness
Sponsored by the Defense Advanced Research Projects Agency

EXECUTIVE SUMMARY

This survey is Part B of the two-part study *Council on Competitiveness Study of the Need for Better Application Software* sponsored by the Defense Advanced Research Projects Agency (DARPA). It is a survey of industrial end users or buyers of HPC systems, and it explores their views and concerns about independent software vendor (ISV) application software and other barriers to using high-performance computing (HPC) more aggressively for competitive advantage. Part A[1] of the study was the first independent, extensive assessment of the landscape and market dynamics surrounding independent software vendors that serve high-performance computing users.

An important impetus for undertaking this study was the July 2004 *Council on Competitiveness Study of U.S. Industrial HPC Users,*[2] also sponsored by the Defense Advanced Research Projects Agency. That 2004 study found, among other things, that 97% of the U.S. businesses surveyed could not exist, or could not compete effectively, without the use of HPC. It also revealed, along with the Council's 2004 *HPC Users Conference Report,*[3] that the lack of production-quality application software is a significant barrier preventing more aggressive use of HPC across the private sector. For U.S. industries that need to outcompete their non-U.S. competitors by outcomputing them, this is indeed a major competitive barrier. In practice, it means that large, complex, and 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 time frames. 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.

Of the U.S. businesses surveyed, 97% could not exist, or could not compete effectively, without the use of HPC.

Part B of the study directly surveyed a select group of well-known U.S. businesses that are highly experienced HPC users. IDC asked them about their requirements for HPC-specific application software and related resources. The HPC end users IDC interviewed for Part B represented a wide range of industries, from defense contractors to an entertainment company and a consumer products supplier. The end users employed a correspondingly broad spectrum of ISV application codes. Not surprisingly, the most frequently named ISV codes were associated with the manufacturing industries (e.g., aerospace and automotive industries), especially to support their shared need to perform structural analysis and computational fluid dynamics simulations.

[1] *Council on Competitiveness Study of ISVs Serving the High Performance Computing Market: Part A — Current Market Dynamics* (IDC #05C4522, July 2005) (Commissioned by the Council on Competitiveness and sponsored by the Defense Advanced Research Projects Agency)

[2] Report available at **http://www.compete.org/pdf/HPC_Users_Survey.pdf**

[3] Report available at **http://www.compete.org/hpc/hpc_conf_report.asp**

KEY FINDINGS

1. HPC-Specific ISV Application Software Is Indispensable for U.S. Industrial Competitiveness

Four out of five (81%) industrial end users indicated that changing current ISV application software was out of the question ("absolutely not") or highly unlikely. Dependence on current ISV suppliers is based on some combination of the software's ability to solve the end user's problem fully and accurately, the software's data format, training and certification requirements, and other factors.

Four out of five (81%) industrial end users indicated that changing current ISV application software was out of the question.

When the enormous time, money, and business disruption that a transition to different software might entail are considered, this fierce loyalty to current ISV software makes sense. Industrial HPC users often spend years using their ISV software before trusting that it consistently produces accurate results on their own crucial problems. Moving to different software could jeopardize this large investment and leave the company competitively vulnerable for a substantial transition period. Worse yet, if the replacement software could not be certified as producing accurate results, the company might no longer be able to meet regulatory and customer requirements needed to sell its products in the marketplace.

If for some reason the end user's current ISV software were no longer available, in three out of four cases (75%) the end user would be able to acquire similar software from another supplier. In the remaining cases (25%), the industrial end users would develop equivalent software themselves or with others in order to maintain business competitiveness. Make no mistake, however: Both of these courses — moving to a different ISV supplier or attempting to develop application software on their own — would expose the companies to the same substantial costs, competitive risks, and major business disruption described above.

If for some reason the end user's current ISV software were no longer available, in three out of four cases (75%) the end user would be able to acquire similar software from another supplier.

It is important to note that in the absence of their familiar ISV application software, none of the industrial end users would cease using HPC to solve these crucial problems. This corroborates the July 2004 *Council on Competitiveness Study of U.S. Industrial HPC Users,* sponsored by the Defense Advanced Research Projects Agency. The 2004 study found, among other things, that 97% of the U.S. businesses surveyed could not exist, or could not compete effectively, without the use of high-performance computing.

"We depend on our ISV software."

"We know and trust the results on our current ISV codes. It takes years to prove the results are correct."

2. Virtually All of the Firms Said They Have Larger Problems That They Can't Solve Today

Respondents indicated that large, complex, and competitively important problems, such as those encountered in designing new cars and airplanes and pharmaceuticals or those increasing the yield from oil reservoirs, often cannot be solved today in reasonable time frames. 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 class.

Most of the firms (83%) said they have unsolvable problems that are 5–100 times larger than the problems they can solve today, although for one U.S. company this figure climbed to 100,000. Also, it is important to note that when industrial HPC end users talk about solving larger problems, they typically don't mean simply doing more of the same thing. In most cases, they mean solving problems with greater resolution that can lead to new insights and superior new products.

Most of the firms (83%) said they have unsolvable problems that are 5–100 times larger than the problems they can solve today.

3. For Many Industrial HPC Users, There Are Substantial Barriers Preventing Them from Using HPC More Aggressively for Competitive Advantage

Respondents cited both financial and technical barriers that are preventing them from using HPC more aggressively. A frequently cited obstacle was inflexible pricing of some ISV application software — a mismatch between the ISV's pricing model and the way in which the industrial buyer would ideally like to use the software. This mismatch echoes concerns IDC raised in Part A of this study about the ability of ISV pricing models to keep pace with current developments in HPC hardware systems. Will ISVs pricing by the number of computers charge the same amount for a 10,000-processor server as for a 100-processor server? How will ISVs pricing by the number of processors count the emerging wave of multicore processors? Clearly, the industrial end users are already starting to wrestle with these issues today.

A frequently cited obstacle was inflexible pricing of some ISV application software.

But it isn't just the cost of software (and hardware) that restricts the use of HPC for gaining competitive advantage in commercial markets. Even if money were no object, the capabilities of currently available hardware were judged inadequate by nearly half (46%) of the industrial firms, and software capabilities were lacking for about one-third (31%) of the respondents. IDC knows from other recent research, for example, that the limited scalability of ISV application software is the main barrier blocking the automotive and aerospace industries from using HPC to even better advantage. Software certification and data formats are training problems that also represent additional barriers.

"Current hardware and software architectures aren't ready to take full advantage of massively parallel operations."

"High-performance hardware and software aren't available."

"Hardware growth is not being matched by software growth in scaling, performance, or business model."

"Software costs are too high."

"We're moderately dependent on the formats of the MAYA file and MENTAL RAY files."

4. A Majority of the U.S. Firms Are Developing Application Software on Their Own, but Only to a Limited Extent and Often with Reluctance

A substantial majority (86%) of the industrial firms indicated they are already developing some application software on their own. The accompanying comments tell the real story, however: Many of the businesses tackle application software only to a limited extent, and even then often with reluctance. In most cases, the application software the industrial firms plan to develop in-house is not intended to replace software they acquire from ISV suppliers. As mentioned in key finding 1, attempting to develop application software on their own would expose the companies to substantial costs, competitive risks, and major business disruption, up to and including the possibility of not being able to introduce or sell their products in the marketplace.

"[Developing our own application software] is too costly in all respects."

"Total cost of ownership is quite expensive for do-it-yourself software."

"We have very little ability to do it ourselves."

"We used to develop our own software, but it is too costly these days."

"We used to develop our own software, but it is too costly these days."

5. Three-Quarters of the U.S. Firms Could Benefit from a Petascale Computer System

Three-quarters of the industrial firms (73%) said they could make use of a petascale computer to run today's crucial problems faster or to tackle next-generation problems of great competitive importance. This is an interesting finding, given that industrial users usually acquire substantially smaller versions of HPC systems than do leading government and academic users. But commercial computer purchases are more heavily dictated by budgets, and the fact that industrial firms have more modest HPC budgets than leading government users does not mean the companies have smaller ambitions for applying HPC. For example, the majority of the industrial end users said they would (83%) or might (91%) use a petascale computer to run heterogeneous problems. A heterogeneous problem, also called a multiphysics or multidisciplinary problem, is one that involves multiple scientific disciplines — for example, studying the complex interaction between the structure of an automobile and the fluid dynamics of air flow around it.

The majority of the industrial end users said they would (83%) or might (91%) use a petascale computer to run heterogeneous problems.

One of the major findings of Part A of this study was that there is a lack of readiness for petascale systems among the ISV suppliers. 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 time frames continue, when petascale systems become available, the majority of ISV codes will not be able to take full advantage of them for at least three to five years.

Fewer than half (46%) of the ISV applications scale even to hundreds of processors today.

"We could do a full engine simulation."

"We could do real-time rendering of full-resolution frames of CG [computer-generated] films. We could also do real-time simulation of clothing, hair, fur, feathers, fluids, etc."

"We could push the envelope with new science and more refined models."

"We could do calculations to define the parameters and do physics that we don't understand today."

6. Market Forces Alone Will Not Address This Problem

Previously, an ISV attempting to improve the performance of application software on a vendor's hardware product collaborated directly with that vendor. The vendor operated on large margins and invested substantial human and financial resources in the collaborative effort. In today's commoditized, lower-margin market for HPC hardware, neither ISV organizations nor HPC hardware vendors can afford to make major new R&D investments to fundamentally rewrite application software to take advantage of highly scalable systems. At the same time, U.S. firms are engaged in intense competition for global market leadership in their own industries and generally don't have the time, resources, or desire to be "in the business" of also developing application software. Their current work on application software is limited and often undertaken with reluctance; in most cases it is not intended to replace the end user's ISV application software. In sum, market forces alone will not address the gap between HPC end users' needs and ISV application software capabilities. Market forces need to be supplemented with external funding support and expertise to improve the scalability of ISV software that is needed for improving the competitiveness of U.S. businesses.

"Few [ISV] vendors are looking out as far as the leading edge of our industry needs to look. The investment from our commercial vendors is limited and is generally focused on market share and commodity, pro-sumer markets."

"We explored solid and fluid mechanics with our ISV and didn't reach a satisfactory conclusion. They said it was too much of an investment and I'd have to foot the entire development bill to address my problem. The ISVs won't take the risk."

"[Developing our own application software] is too costly in all respects."

"[Developing our own application software] is too costly in all respects."

7. Nearly Half the U.S. Firms Would Be Willing to Partner with Outside Parties to Develop Application Software, Though Some Have Concerns

Among the partner types, U.S. firms are most interested in collaborating with ISVs (67%). National labs and universities were tied (50% positive responses each) as the second-most-popular potential partners for collaborations related to application development. Government agencies were also viewed favorably (42% positive responses), though not as favorably as ISV suppliers, national labs, or universities. Easily the least-favored type of category for potential collaboration consisted of "competitors in the company's own industry." To help complete the picture, we asked the industrial end users whether they had concerns about working with non-U.S. partners. A small majority (54%) replied yes.

"We'd had very good experience [with ISV partnerships] across our industry and will continue doing this."

"We'd had very good experience [with ISV partnerships] across our industry and will continue doing this."

"We love working with labs."

"Our experience with agencies has been very positive."

"We'd want exclusive use of the results for a period of time."

"We're open to partnerships but would want to maintain our competitive advantage in our core industry."

"We'd want a partnership where the scope, goals, and methods are agreed upon in the beginning and we are an active partner."

IMPLICATIONS OF THE FULL STUDY: OBSERVATIONS FROM PART A AND PART B

How the End Users' Views Align with the ISVs' Views

Table 1 lists the key findings from Part A (the ISVs' view) and Part B (the end users' view) of the study. The juxtaposed findings reveal the large disparity between the need of U.S. businesses to use HPC more aggressively for competitive advantage, and the current plans of ISV suppliers to meet this crucial need. The limitations of HPC-specific ISV application software are not the only barrier to fuller exploitation of HPC but are regularly cited by industrial end users as the most important constraint.

TABLE 1

Comparison of Key Findings

Study Part A: ISV Suppliers	Study Part B: HPC End Users
The business model for HPC-specific application software has all but evaporated in the last decade.	HPC-specific ISV application software is indispensable for U.S. industrial competitiveness.
ISV applications can exploit only a fraction of the problem-solving power of today's high-performance computers.	Virtually all of the firms said they have larger problems that they can't solve today.
For many applications, the ISVs know how to improve scalability but have no plans to do so because the HPC market is too small to justify the R&D investment.	The lack of scalable application software is preventing many industrial users from using HPC more aggressively for competitive advantage.
There is a lack of readiness among ISV suppliers for petascale systems.	Three-quarters of the U.S. firms could benefit from a petascale computer system.
Market forces alone will not address the gap between HPC users' needs and ISV software capabilities.	Market forces alone will not address the gap between HPC users' needs and ISV software capabilities.
Most ISVs would be willing to partner with outside parties to accelerate application software development.	Nearly half the firms would be willing to partner with outside parties to accelerate application software development.

Source: IDC, 2006

Implications of the Study

U.S. industry has long been a global leader in exploiting HPC to drive product innovation. This innovation has given U.S. businesses an important marketplace advantage over competitors from nations with lower labor costs or lesser ability to exploit HPC. In the past decade, however, the ISVs' business model for developing HPC-specific application software has nearly evaporated. As a result, key U.S. industries often do not have the ISV software they want and need. This has rendered them incapable of innovating as aggressively as they would like, and has compromised U.S. industrial competitiveness. Market forces alone can no longer correct this problem; they need to be supplemented with external funding and expertise. Fortunately, most ISVs and a substantial portion of the U.S. businesses are willing to partner with outside parties to speed progress in creating more capable HPC application software.

PART B METHODOLOGY

Definitions: Additional Terminology for Part B

Certification

Certification guarantees that the computer simulation of a scientific or engineering problem will match experimental ("real world") results within a certain window of accuracy. Certification in this context refers to an ISV supplier's "stamp of approval" that one of its applications runs compatibly and accurately, and in compliance with any applicable regulatory requirements, on a specific HPC computer product (e.g., Cray XT3, HP Integrity Superdome, IBM Blue Gene/L, SGI Altix, Voltaire Pinnacle). Given the substantial number of HPC hardware and operating system variants (Unix and Linux derivatives, Windows) in the market today, ISV suppliers — many of which are small businesses — may lack the financial and human resources to certify, in a timely manner, a version of their codes for every variant. Considering that one of the main purposes of certification is to ensure that an application is producing accurate results, lack of certification can pose major problems for U.S. businesses that are engaged in life-and-death battles to bring higher-quality products to market ahead of their global competitors. If application software cannot be certified as producing accurate results, the company might not be able to meet regulatory and customer requirements needed to sell its products in the marketplace.

Data Formats

In the context of this study, data formats refer to unique ways in which data are arranged for processing, storage, and visual display. Familiar examples of data formats from the world of desktop applications include Microsoft PowerPoint (.ppt format), Microsoft Word (.doc format), and Adobe Reader (.pdf format). In the high-performance computing realm, there are hundreds of ISV applications and scores of unique, proprietary data formats. A major complaint of HPC end users, according to recent IDC research, is that ISV software applications often cannot "talk to" each other; that is, data from one application can't be deciphered, for the purposes of combining or comparing results, by applications with different data formats.

Study Background — Part B

Part A of this study provided the first extensive, independent assessment of the landscape and market dynamics surrounding ISVs that serve HPC users. In summary, Part A concluded that the ISVs' business model for creating HPC-specific application software has all but evaporated in the past decade, and market forces alone will not address the need of U.S. businesses for applications they can use to exploit HPC more aggressively for global competitive advantage. Market forces need to be supplemented with external funding support and expertise to improve the scalability of ISV software that is needed for improving the competitiveness of U.S. businesses. If this support were available, most ISV organizations would be willing to partner with outside parties to accelerate progress.

Part B of the study directly surveyed a select group of well-known U.S. businesses that are highly experienced HPC users. IDC asked them about their requirements for HPC-specific application software and related resources to supplement the broader discussion that occurred at the Council on Competitiveness Users Conference and Software Workshop (July 13 and14, 2005). Part B complements and expands on the results of IDC's July 2004 *Council on Competitiveness Study of U.S. Industrial HPC Users*. This study found that among the criteria for purchasing high-performance computers, "performance on our applications" was most prominent. In the battles U.S. businesses wage for global market supremacy, faster application performance often means faster time to market and superior products.

In the July 2004 study, IDC asked the U.S. businesses about their interest in a hypothetical high-performance computer 100 times faster than today's models — analogous to questions in the present study about petascale computers. The 2004 respondents first cautioned that it is not sufficient to imagine a machine that is merely 100 times faster theoretically — the machine would need to be 100 times faster than today's high-performance computers when running users' own ISV applications. As the chief benefit of a dramatically faster computer like this, they then cited the ability to produce higher-quality products. Those who were able to quantify this benefit named figures ranging from $10 million to several billion dollars for their organization alone. Other positive effects of better high-performance computers included more powerful pharmaceutical drugs and faster disease cures, more environmentally friendly manufacturing, reduced litigation expense, and more-entertaining animated films.

As a preface to this study, it is important to note that the U.S. businesses participating in the July 2004 study stressed that for the foreseeable future, the crucial competitive benefits of HPC will remain heavily dependent not just on faster hardware servers, but on the capabilities of ISV application software.

Methodology — Part B

This study is based on interviews of 13 well-known U.S. commercial firms that are important end users of high-performance computing, combined with five in-depth interviews of end users. Part B was not intended to be a large industrial user survey. IDC selected a targeted group of highly experienced HPC users for a benchmark response to Part A of the study. These industrial HPC end users represented a wide range of industries and employed a correspondingly broad spectrum of ISV applications codes. The primary industries of the survey respondents spanned seven sectors, from defense contractors in the aerospace industry to an entertainment company and a consumer products supplier. The interviews included a mix of multiple choice, quantitative, and qualitative questions. To the extent possible within the pragmatic time limits of the interviews, we elicited additional comments on the topics in question. This document includes a representative sampling of the comments.

Study Limitations — Part B

While IDC aims to provide an accurate, comprehensive view of the subject being studied, certain limitations inevitably affect the results. Based on other recent IDC research and HPC activities, IDC believes that the opinions of the 13 industrial end-user organizations interviewed for this study fairly represent the general thinking of the larger community of U.S. industrial HPC users. It would be presumptuous to claim, however, that there are no users in the larger U.S. industrial community whose opinions differ in certain respects from any in the group IDC interviewed for this study. Also, consistent with this study's purpose, we elected to focus on U.S. end-user companies. This study therefore does not purport to represent the opinions of non-U.S. industrial end users of HPC. Finally, with a group size of 13 end users, some less-popular options for responding to questions are thinly represented, occasionally with only one or two responses. IDC has tried to exercise extreme caution in generalizing from such results and cautions readers to do the same.

PART B STUDY RESULTS

The Part A study results are available from the Council on Competitiveness at **www.compete.org/hpc**.

Industry HPC End-User Demographics

The U.S. industrial HPC end users IDC interviewed for this study represented a wide range of industries and employed a correspondingly broad spectrum of ISV application codes.

Primary Industry of End Users

When IDC asked for the primary industry of the survey respondents (see Figure 1), the replies spanned seven sectors, from defense contractors in the aerospace industry to an entertainment company and a consumer products supplier. More than half of the respondents (61%) belonged to the aerospace and automotive industries, which are often grouped together in an HPC context because they share many of the same problems (e.g., aerodynamic design, fuel efficiency) and therefore use many of the same ISV applications.

Primary Application Codes of End-Users

Not surprisingly (see Table 2), the most frequently named ISV codes are associated with the aerospace and automotive industries, especially to support their shared need to perform structural analysis (e.g., LS-DYNA, ABAQUS) and computational fluid dynamics (e.g., FLUENT) simulations. Even within the aerospace/automotive sector, preferences for favorite structural analysis and CFD applications varied considerably. At the other end of the spectrum is a substantial list of ISV applications mentioned only once each, such as MAYA and MENTAL RAY, both of which are used by filmmakers and others for visual rendering.

The most frequently named ISV codes are associated with the aerospace and automotive industries.

FIGURE 1

Primary Industry of End Users

Q. What is your primary industry?

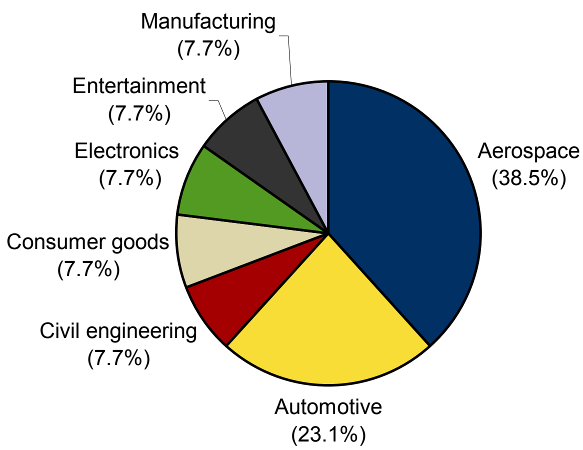

Source: IDC, 2006

TABLE 2

Primary Application Codes of End Users

Q. What are the primary ISV application codes you use for HPC?

Application Code	Most Important	Second	Third	Total Mentions
FLUENT	2	3	1	6
ABAQUS	2	1	1	4
LS-DYNA	2	2	–	4
ANSYS	2	1	–	3
IDEAS Simulation	–	1	1	2
NASTRAN	–	–	2	2
ACCELRYS	–	–	1	1
ALLSTAR	1	–	–	1
CADENCE	1	–	–	1
CFD++	–	–	1	1
CFL3D	–	1	–	1
GAMBIT	–	1	–	1

TABLE 2

Primary Application Codes of End Users

Q. What are the primary ISV application codes you use for HPC?

Application Code	Most Important	Second	Third	Total Mentions
MAYA (ALIAS)	1	–	–	1
MENTAL RAY	–	1	–	1
MENTOR	–	–	1	1
NASTAR	–	–	1	1
NORDEX	–	–	1	1
OVERFLOW	1	–	–	1
STAR-CD	–	1	–	1
SYNOPSIS	–	1	–	1
TETRIS	–	–	1	1
Y237	–	1	–	1
Total	12	14	11	37

Source: IDC, 2006

Percent of Applications Developed In-House Versus Externally

The ratio of applications created by the businesses themselves to applications developed externally by third parties varied greatly by industry (see Table 3). Overall, three-quarters (75%) of the applications being used by these organizations were acquired from external ISV vendors, with the remainder (25%) developed in-house. Looking beyond the single consumer goods company that used outside codes exclusively (100%), the automotive (97%) and aerospace (72%) firms relied most heavily on external ISV application software. The major ISV codes serving these two sectors, and to some extent the manufacturing sector as well, typically are of older origin and have stood the test of time. At the opposite extreme, the entertainment industry firm depended minimally (20%) on outside applications that are newer and less well tuned to the company's requirements.

TABLE 3

Percent of Applications Developed In-House Versus Externally

Q. What percent of your applications are developed internally and what percent are developed externally?

		Mean (%)	
Industry	Number	In-House	External
Consumer goods	1	–	100
Automotive	3	3	97
Aerospace	5	28	72
Manufacturing	1	30	70
Electronics	1	40	60
Entertainment	1	80	20
Total	12	25	75

Source: IDC, 2006

Dependence on ISV Application Software Vendors

Lock-In to ISV Supplier

Roughly three-quarters (73%) of the industry organizations said they are not "locked in" to their current ISV application vendors through formal agreements (see Figure 2). Even for this majority, however, the contractual freedom to change is often tempered by practical dependence on the ISV applications for day-to-day business operations. This practical dependence presumably is greatest in sectors, like the aerospace and automotive industries, that make heaviest use of third-party ISV codes (refer back to Table 3).

"We have some two-, three-, and five-year agreements, but none are exclusive."

"We have some two-, three-, and five-year agreements, but none are exclusive."

Training Issues with ISV Supplier

Most (58%) of the businesses said they had training problems with their ISV application vendors (see Figure 3), though a significant number (42%) reported no problems. Comments cited the cost of training and the fact that training requirements can vary greatly by application. Most also (58%) replied yes when asked whether current training problems would affect a decision to change ISV suppliers (see Figure 4).

"Training is available from our ISV, but it would be expensive to retrain all the engineers in our group."

"Our training issues range from significant to limited, depending on the application."

"Yes, but they can be overcome."

"Yes. Our code has a relatively small user base but is highly powerful. It's complex to use and requires good training."

"Yes. Training's a big factor."

"No. Training wouldn't prevent us from changing."

FIGURE 2

Lock-In to ISV Supplier

Q. Are you locked in to your current ISV application suppliers?

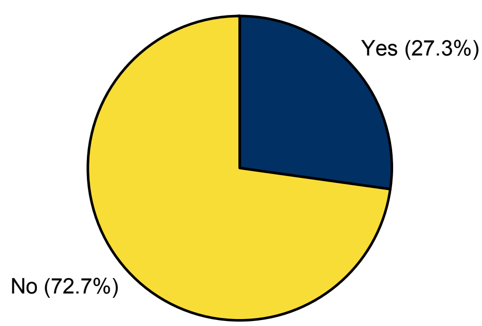

n = 11

Source: IDC, 2006

FIGURE 3

Training Problems with ISV Supplier

Q. Do you have training problems with your current ISV application suppliers?

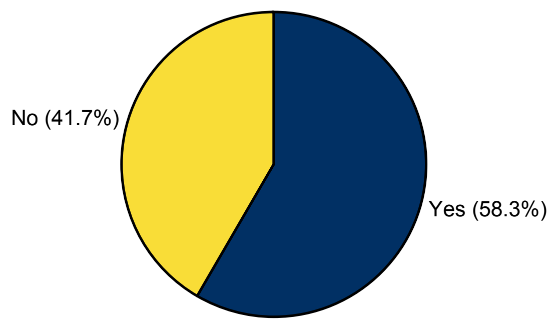

n = 12

Source: IDC, 2006

FIGURE 4

Training Issues Affecting Changing ISV Supplier

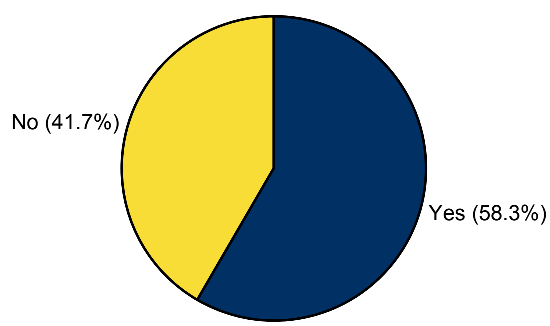

n = 12

Source: IDC, 2006

Certification Issues with ISV Supplier

More than half (55%) of the respondents affirmed that they had certification problems with their ISV suppliers (see Figure 5). Certification in this context refers to an ISV supplier's "stamp of approval" that one of its applications runs compatibly and accurately, and in compliance with any applicable regulatory requirements, on a specific HPC computer product (e.g., Cray XT3, HP Integrity Superdome, IBM Blue Gene/L, SGI Altix, Voltaire Pinnacle). Given the substantial number of HPC hardware and operating system variants (Unix and Linux derivatives, Windows) in the market today, ISV suppliers — many of which are small businesses — may lack the financial and human resources to certify in a timely manner a version of their codes for every variant. Considering that one of the main purposes of certification is to ensure that an application is producing accurate results, lack of certification can pose major problems for U.S. businesses that are engaged in life-and-death battles to bring higher-quality products to market ahead of their global competitors. If application software cannot be certified as producing accurate results, the company might not be able to meet regulatory and customer requirements needed to sell its products in the marketplace.

More than half (55%) of the respondents affirmed that they had certification problems with their ISV suppliers.

FIGURE 5

Certification Issues with ISV Supplier

Q. Do you have certification problems with your current ISV application suppliers?

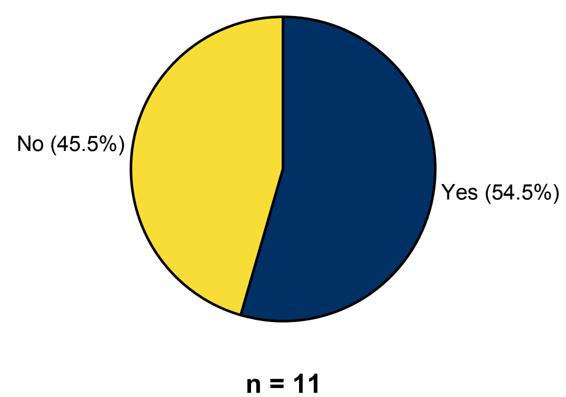

Source: IDC, 2006

Data Format Dependence on ISV Supplier

A major complaint of HPC end users, according to recent IDC research, is that ISV software applications often cannot "talk to" each other; that is, data from one application can't be deciphered by other applications for the purposes of combining or comparing results. This limitation can be especially frustrating for businesses that want to gain competitive advantage by tackling heterogeneous problems, also called multiphysics or multidisciplinary problems. Heterogeneous problems involve multiple scientific disciplines — for example, studying the complex interaction between the structure of an automobile (structural analysis) and the air flow surrounding the vehicle (computational fluid dynamics).

An important reason for the incompatibility among many ISV applications is varying data formats — unique ways in which data are arranged for processing, storage, and visual display. Familiar examples of data formats from the world of desktop applications include Microsoft PowerPoint (.ppt format), Microsoft Word (.doc format), and Adobe Reader (.pdf format). In the high-performance computing realm, there are hundreds of ISV applications and scores of unique, proprietary data formats.

It is therefore not surprising that more than three-quarters (78%) of the industry respondents in this study pointed to data format dependence on ISV suppliers as a serious constraint (see Figure 6).

More than three-quarters (78%) of the industry respondents in this study pointed to data format dependence on ISV suppliers as a serious constraint.

"We're moderately dependent on the formats of the MAYA file and MENTAL RAY files."

"The translation of FEA models between the West Coast and East Coast can be difficult."

FIGURE 6

Data Format Dependence on ISV Supplier

Q. Do you have data format dependence on your current ISV application suppliers?

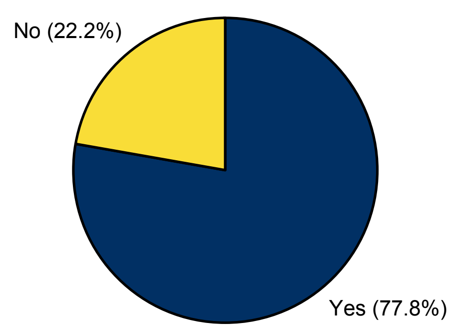

n = 9

Source: IDC, 2006

Likelihood of Changing ISV Software

Four out of five (81%) of the responses to this question indicated that changing current ISV software was out of the question ("absolutely not") or highly unlikely, even if the end users could change if they so desired (see Figure 7). Only one respondent was actively exploring options for changing ISV software.

Four out of five (81%) of the responses to this question indicated that changing current ISV software was out of the question ("absolutely not") or highly unlikely.

When the enormous time, money, and business disruptions that a transition to different software might entail are considered, this fierce loyalty to current ISV software makes sense. Industrial HPC users often spend years using their ISV software before trusting that it consistently produces accurate results on their own crucial problems. Moving to different software would jeopardize this large investment and leave the company competitively vulnerable for a substantial transition period. Worse yet, if the replacement software could not be certified as producing accurate results, the company might no longer be able to sell its products in the marketplace. From a business standpoint, a major disruption like this is highly impractical and to be avoided at almost any cost.

FIGURE 7

Likelihood of Changing ISV Software

Q. What is the likelihood you will change your current ISV application software?

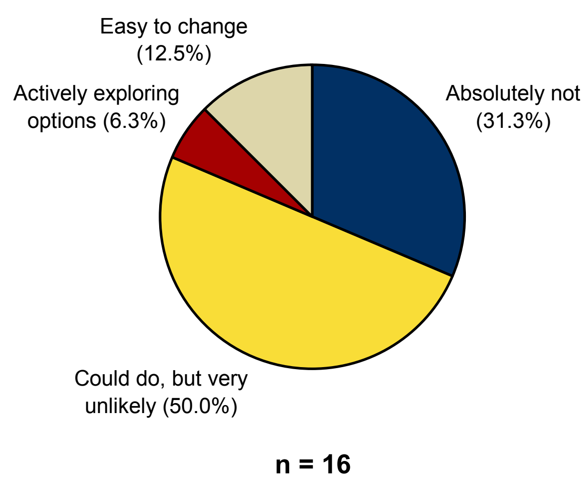

Note: Multiple responses were allowed.

Source: IDC, 2006

Action If ISV Software Weren't Available

If for some reason the end user's current ISV software were no longer available, in three out of four cases (75%), the end user would be able to acquire similar software from another supplier (see Figure 8). In the remaining cases (25%), the industrial end users would develop equivalent software themselves, either because they preferred the do-it-yourself route or lacked other options. Make no mistake, however: Both of these courses — moving to a different ISV supplier or attempting to develop application software on their own — would expose the companies to the same substantial costs, competitive risks, and major business disruption described in the preceding question (refer back to Figure 7). It could take hundreds, even thousands, of person years to recreate an ISV software application, and even then questions would remain for some time about its efficacy in solving company-specific problems.

What is interesting about responses to the current question, however, is that in the absence of their familiar ISV application software, none of the industrial end users would cease using HPC to solve these crucial problems. This corroborates IDC's July 2004 *Council on Competitiveness Study of U.S. Industrial HPC Users,* which was sponsored by the Defense Advanced Research Projects Agency. The 2004 study found, among other things, that 97% of the U.S. businesses surveyed could not exist, or could not compete effectively, without the use of high-performance computing.

FIGURE 8

Action If ISV Software Weren't Available

Q. What would you do if your current ISV application software were no longer available?

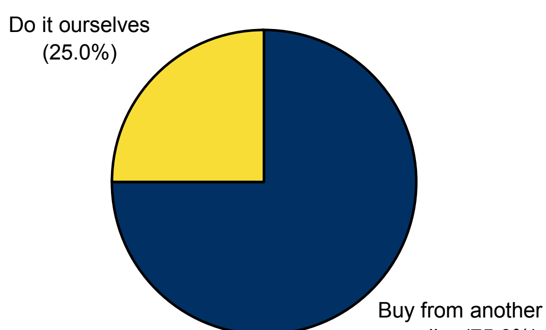

n = 16

Note: Multiple responses were allowed.

Source: IDC, 2006

Criteria for Selecting Current ISV Software

As Figure 9 shows, the most important criterion for selecting the industrial end user's current ISV application software was the "key attributes of the code" (average rating: 59%). "Key attributes" refers to the fit between the application software and the technical computing problem the user is attempting to tackle (i.e., the software's ability to solve the problem fully and accurately). Because the industrial problems in question are tightly linked to product development and therefore corporate success, it is not surprising that application codes with the right key attributes emerged as the top criterion for end users.

FIGURE 9

Criteria for Selecting Current ISV Software

Q. On a scale of 0–100%, where 100% defines the highest possible importance, rate each of the criteria you used for selecting your current ISV application software.

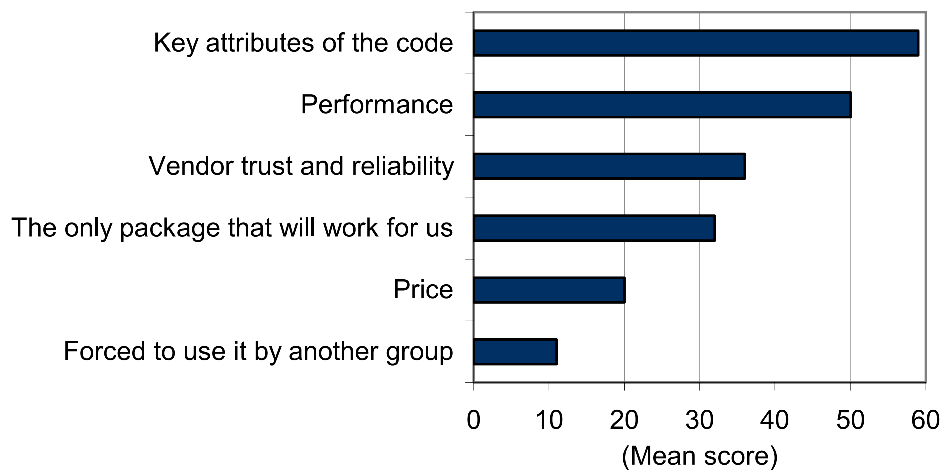

n = 12

Source: IDC, 2006

Next in importance (50% average rating) was the application software's speed ("performance") in solving the problem. Faster problem-solving speed can mean faster time to market for new products. Running a distant third and fourth in the ratings were vendor reliability (36%) and software uniquely able to solve the problem (32%).

Interestingly, price as a criterion for selecting ISV application software came in next to last (20%), ahead of only the situation in which a group was forced to use an ISV software package by another group in the same company (11%). Yet other IDC research shows that the costs of ISV application software are especially burdensome for industry. In Part A of this study[4], the ISV organizations reported that applications software costs approached 50% of overall HPC costs for businesses, versus only 5–15% for academic users and 5–10% for government users, who often develop their own software codes for the research they pursue. IDC also knows from other research that software costs become even more onerous as users scale up the size of their hardware systems. In spite of this, "key attributes of the code" take precedence over price and other factors because industrial firms are highly dependent on ISV codes for their survival and success. ISV software typically plays an important but less crucial role in academic and government organizations.

Interestingly, price as a criterion for selecting ISV application software came in next to last (20%).

Key Weaknesses of Current ISV Software

"Limited scalability" and "inflexible licensing model" (23% of the responses each) were the most frequently cited weaknesses of industrial end users' current ISV application software (see Table 4).

As used in this study, scalability means the ability of application software to effectively exploit more processors on an HPC computer server in order to solve current problems faster and with greater resolution and to effectively address more-complex next-generation problems. The largest HPC systems may have hundreds or even thousands of processors today, growing to tens of thousands in the near future, yet Part A of this study (see key finding 2) showed that many applications popular in industry today can exploit only 1–4 processors in practice (some scale to 16–32 processors or more). This limited scalability retards a company's ability to answer the "what if" questions that provide new insights needed to drive the process of innovation.

Industrial end users' complaints about "inflexible licensing models" generally point to a mismatch between the ISV's pricing model and the way in which the industrial buyer would ideally like to use the software. This mismatch echoes the concerns raised in Part A of this study, in which the ISV organizations reported their current pricing models: charging by the number of users, by the number of processors the application can be run on, and by issuing site licenses for unrestricted use. IDC said then, and repeats here, that it will be interesting — and important — to see how ISV organizations grapple with current developments in HPC hardware systems. Will those pricing by the number of computers charge the same amount for a 10,000-processor server as for a 100-processor server? How will those pricing by the number of processors count the emerging wave of multicore processors? Clearly, the industrial end users are already starting to wrestle with these issues today.

Industrial end users' complaints about "inflexible licensing models" generally point to a mismatch between the ISV's pricing model and the way in which the industrial buyer would ideally like to use the software.

Beyond these two frequently cited weaknesses, a long list of other drawbacks received single mentions: performance, price/costs, lack of openness, limited future functionality, show enhancements, lack of support for problems, and incompatibility with other codes.

---

[4] See *Council on Competitiveness Study of ISVs Serving the High Performance Computing Market: Part A — Current Market Dynamics*, available at **www.compete.org/hpc**.

TABLE 4

Key Weaknesses of Current ISV Software

Q. What are the key weaknesses of your current ISV application software?

	Number	% of Respondents
Limited scalability	3	23.1
Inflexible licensing model	3	23.1
Performance	1	7.7
Price/costs	1	7.7
Lack of openness	1	7.7
Limited future functionality	1	7.7
Slow enhancements	1	7.7
Lack of support for problems	1	7.7
Incompatibility with other codes	1	7.7
Total	13	100.0

Source: IDC, 2006

Stress Points with Current ISV Software

As a cross-check, we used a different format to ask again about end users' problems with current ISV application software. Once again (see Table 5), "scaling ability" (limited scalability) clearly emerged as the most important problem (42% average rating), following by pricing model (32%) and price level (31%). Other issues were less important.

TABLE 5

Stress Points with Current ISV Software

Q. On a scale of 0–100%, where 100% defines the highest possible importance, rate each of the stress points with your current ISV application software.

Stress point	Mean Score
Scaling ability	42
Pricing model	32
Price level	31
Performance	25
Missing features/functions	22
Vendor service/support	10

n = 12

Source: IDC, 2006

Change to a Different ISV Supplier

Although in response to previous questions (see Table 6) the industrial end users assigned higher importance to software scalability than price, answers to question number 6 showed that priorities can shift in a different context. When asked what would motivate them to change ISV suppliers, end users said a relatively modest 27% price improvement would be enough to trigger the move, whereas a 46% improvement in scalability or a 122% performance speedup would be needed to motivate a change in ISV suppliers.

TABLE 6

Change to a Different ISV Supplier

Q. What would it take for you to move to a different ISV supplier?

Criteria	Required Change (%)
Performance improvement	122
Scaling improvement	46
Price savings improvement	27

n = 12

Source: IDC, 2006

Desired Functionality or Features

Respondents were asked, What desired functionality or features are missing from your current ISV application software? The wish list of functions and features industrial end users would like to see included in their ISV application software ranged from the general (improved algorithms, better ease of use, more compatibility between codes) to the specific (interactive lighting interface, support for direct transient rocket analysis).

"We'd like new physics and usability features."

"The codes don't talk to each other. There are no coupled solutions."

"The codes don't talk to each other. There are no coupled solutions."

"For film rendering, we need an interactive lighting interface and better scan line performance. For another important code, scaling to complexity is a big issue, and so is working in parallel. The code needs to open access to libraries instead of through the existing limited interface."

"We need support to direct transient solid rocket analysis."

Willingness to Develop Application Software

Four out of five (79%) of the industrial end users said they were willing to develop application software themselves (see Figure 10), but even within this group there were strong reservations. Industrial HPC users often spend years using their ISV software before trusting that it consistently produces accurate results on their own crucial problems. Attempting to develop application software on their own would expose the companies to substantial costs, competitive risks, and major business disruption. It could take hundreds, even thousands of person years to recreate an ISV software application, and even then, questions would remain for some time about its efficacy in solving company-specific problems. If the replacement software could not be certified as producing accurate results, the company might no longer be able to sell its products in the marketplace. From a business standpoint, a major disruption like this is highly impractical and to be avoided at almost any cost.

"Yes, but we hate to do it."

"It's too costly in all respects."

"Total cost of ownership is quite expensive for do-it-yourself software."

"Total cost of ownership is quite expensive for do-it-yourself software."

"We have very little ability to do it ourselves."

FIGURE 10

Willingness to Develop Application Software

Q. Are you willing to develop application software yourself?

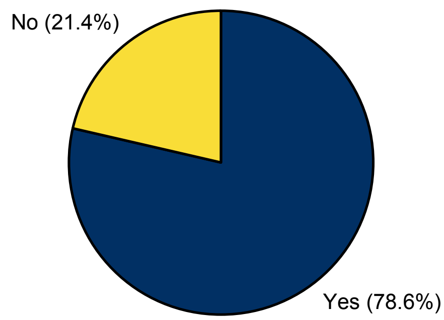

n = 14

Note: Multiple responses were allowed.

Source: IDC, 2006

Developing Application Software

A slightly larger majority (86%) indicated they are already developing application software on their own (see Figure 11). Once again, however, the comments tell the real story: Many of the businesses tackle application software only to a very limited extent, and even then often with great reluctance.

"We generally have relied on our own software because few vendors are looking out as far as the leading edge of our Industry needs to look. The investment from our commercial vendors is limited and is generally focused on market share and commodity, pro-sumer markets."

"Yes, but only the user interface."

"Yes, but very limited."

"Yes, to a limited extent."

FIGURE 11

Developing Application Software

Q. Are you currently developing application software yourself?

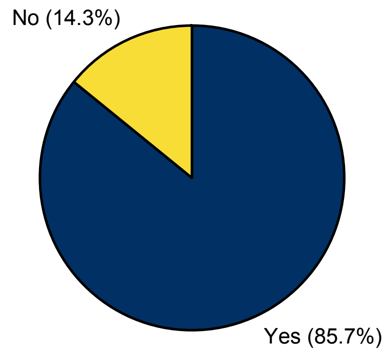

n = 14

Note: Multiple responses were allowed.

Source: IDC, 2006

Developing Application Software with National Lab or University

As Figure 12 indicates, nearly half (46%) of the businesses are already developing application software with a national laboratory or university. This provides a strong baseline for future collaboration among these parties.

FIGURE 12

Developing Application Software with National Lab or University

Q. Are you currently developing application software with a national lab or university?

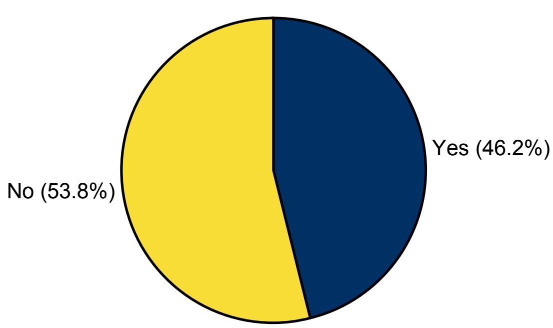

n = 13

Source: IDC, 2006

Developing Application Software in the Future

More than three-quarters (77%) of the industrial firms said they plan to develop application software in the future, as Figure 13 illustrates. Again, the comments show that, at least in some cases, the software in question may be very limited.

"Yes, but only user interfaces."

"To a limited extent."

FIGURE 13

Developing Application Software in the Future

Q. Do you plan to develop application software in the future?

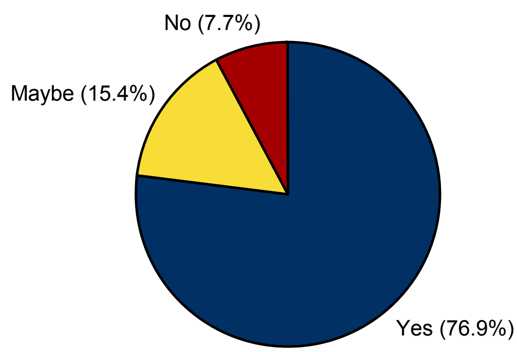

n = 13

Source: IDC, 2006

Replace External Software

In the majority of these cases (60%), the application software the industrial firms plan to develop is not intended to replace software they acquire from ISV suppliers, but in an impressive 4 out of 10 instances, the software is aimed at doing precisely that (see Figure 14).

FIGURE 14

Replacing External Software with Application Software

Q. Will your application software replace external application software from ISVs?

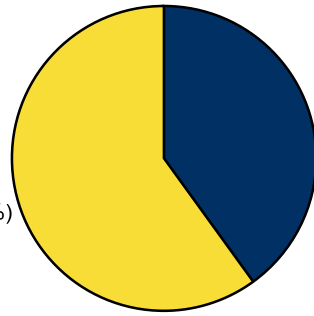

n = 10

Source: IDC, 2006

Interest in Working with National Lab, University, or Government Agency

When "yes" (69%) and "maybe" (23%) responses are combined, 92% of the industrial end users declared an interest in working with a national lab, university, or government agency to create application software (see Table 7). Previously in this study, we found that virtually all ISV organizations were willing to collaborate on application software develop with the government or other outside parties.

"Yes. We already work with all three types."

"Possibly. We would have to look at how much they offered versus any possible slowdown in moving to a larger team."

TABLE 7

Interest in Working with National Lab, University, or Government Agency

Q. Are you interested in working with a national laboratory, university or government agency to develop application software?

	Number	% of Respondents
Yes	9	69.2
Maybe	3	23.1
No	1	7.7
Total	13	100.0

Source: IDC, 2006

Preferred Partners for Developing Application Software

Table 8 displays the types and names of the organizations the industrial firms said they would prefer to collaborate with in developing application software (ISVs were not included in the choices for this question). As a category, government easily won out over universities by a factor of 15 to 2. Within government, the Department of Energy (DOE) emerged as the department most often preferred, obtaining 10 mentions compared with three for the Department of Defense (DOD) and two for NASA. Sandia National Laboratories and Oak Ridge National Laboratory, with the closely allied University of Tennessee, stood out as popular choices in the DOE family.

While IDC cautions against making detailed conclusions based on this polling of a limited number of industrial end users, we believe the general preference for government as an application development partner and the strong interest in the DOE labs are instructive.

Some of the respondents continued to express reservations about these external partnerships, pointing to existing relationships with ISV suppliers ("We mostly buy from ISVs") or the costs of external collaboration ("Coordination may make it more difficult and expensive in opportunity costs than doing it entirely in-house").

The "no preference" and "not sure" responses indicate that work is still needed to make HPC users in industry aware of the opportunities for partnering with public sector organizations.

"We have worked with 30 different ones in various categories."

"We like CRADA umbrellas. All that's currently required is a statement of work."

"We're currently looking at various opportunities."

"We don't know what's available."

"We have no preference."

TABLE 8

Preferred Partners for Developing Application Software

Q. Who would you prefer to partner with to develop application software?

Partner Type/Organization	Number
DOE	
Sandia	4
ORNL/UT	2
Argonne	1
LANL	1
PNNL	1
DOE unspecified	1
Subtotal	10
DOD	
Air Force	1
Navy	1
DOD unspecified	1
Subtotal	3
NASA	2
Universities	
Stanford	1
University of Minnesota	1
Subtotal	2
No preference	1
Not sure	1
Total	19

Note: Multiple responses were allowed.

Source: IDC, 2006

Larger Problems to Solve

Virtually all (92%) of the industrial firms said they have larger (or more intractable) problems they would like to be able to solve (see Figure 15).

FIGURE 15

Larger Problems to Solve

Q. Do you have larger problems you'd like to solve?

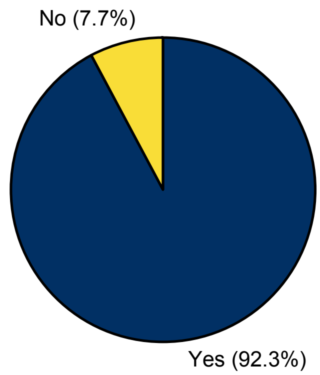

n = 13

Source: IDC, 2006

Size of Current Problems

To provide a baseline for comparisons with currently unsolvable problems, we asked the industrial end users about the sizes of the largest problems they are solving today (see Table 9). Although two of the respondents had large problems requiring 10 or fewer hours of CPU (processor) time to complete, the great majority (78%) of the larger problems required 100 or more CPU hours. For four of the nine (44%) companies responding to this question, problem size exceeded 10,000 CPU hours. In the case of a filmmaker, a single large problem — presumably the processing of a feature-length film — consumed more than 10 million CPU hours. It is worth noting here that many industrial firms routinely tackle problems of varying size.

"There's no one answer. We have some large, medium, and rather small problems."

"Ours are greater than two weeks to solution."

"Very large!"

"Our large problems take 1,000+ CPU hours routinely."

"Rendering a film takes 10–100 gigaflops per frame, needs more than 300 terabytes of memory and more than 10 million CPU hours."

"Currently, we have problems in the several-CPU-hour range."

TABLE 9

Size of Current Problems

Q. How large are the largest problems you are solving today, in terms of CPU hours?

CPU Hours	Number	% of Respondents
<5	1	11.1
5–10	1	11.1
100	1	11.1
1,000	1	11.1
3,000	1	11.1
>10,000	4	44.4
Total	9	100.0

Source: IDC, 2006

Need to Solve Larger Problems

The answers to this question varied dramatically, ranging from 5x to 100,000x. For most of the firms (83%), however, currently unsolvable problems were in the 5–100x range (see Table 10).

Software Cost Limitations

More than half (54%) of the industrial end users said their ability to use HPC as aggressively as they would like for competitive advantage is limited by software costs (see Table 11). When the "at times" responses are added in, this figure climbs to 77%.

"Software costs are too high."

"Yes. Because of this, we sometime can't make a business case."

"Not really. We would likely use in-house software."

Hardware Cost Limitations

An even larger number of companies (69%, or 84% with "at times" responses included) are constrained by the cost of hardware (see Table 12).

"We also base hardware expenditures on the business case."

"Yes. To reduce our production cycles, we could use far more than we can afford."

TABLE 10

Need to Solve Larger Problems

Q. How much larger are the problems you'd like to solve?

Multiple	Number	% of Respondents
5–10x	7	58.3
100x	3	25.0
10,000x	1	8.3
100,000x	1	8.3
Total	12	100.0

Source: IDC, 2006

TABLE 11

Software Cost Limitations

Q. Are you limited by software costs?

	Number	% of Respondents
Yes	7	53.8
No	3	23.1
At times	3	23.1
Total	13	100.0

Source: IDC, 2006

TABLE 12

Hardware Cost Limitations

Q. Are you limited by hardware costs?

	Number	% of Respondents
Yes	9	69.2
No	2	15.4
At times	2	15.4
Total	13	100.0

Source: IDC, 2006

Other Critical Limitations

The outlines of the problem emerged more clearly when we asked about other critical limitations to the use of HPC. As Table 13 illustrates, it isn't just the cost of hardware and software that restricts the use of HPC for gaining competitive advantage in commercial markets. Even if money were no object, the capabilities of currently available hardware were judged inadequate by nearly half (46%) of the industrial firms, and software capabilities were lacking for about one-third (31%) of the respondents.

TABLE 13

Other Critical Limitations

Q. Are there other critical limitations preventing you from using HPC more aggressively?

	Number	% of Respondents
Hardware capabilities	6	46.2
Software capabilities	4	30.8
Problems don't parallelize well	2	15.4
Incompatible with others in company	1	7.7
Total	13	100.0

Source: IDC, 2006

For some companies, hardware capabilities are fine and software is the laggard; for other firms, the opposite is true; and for still others, both hardware and software fall short of their requirements. Much depends on the mix of problems each business needs to run. IDC knows from other recent research, for example, that the limited scalability of ISV application software is the main barrier for the automotive and aerospace industries.

The issue of hardware and application software scalability does not affect all industries or all applications equally. The challenge of increasing an application's scalability is closely related to the nature of the application, and ultimately to the mathematical complexity of the underlying technical problem, or set of problems, the application was designed to address.

At one end of the spectrum are certain "embarrassingly parallel" seismic analysis problems in the petroleum industry, for example, that can be neatly divided into distinct subproblems (e.g., single acoustic soundings). Each of the many subproblems can be run independently on a different processor of a high-performance computer, resulting in good scalability on many types of computer architectures, including more affordable clusters.

At the other extreme are challenging structural analysis problems and the emerging class of multidisciplinary (multiphysics, multiscale) problems, such as coupled fluid-structure interactions in the automotive and aerospace industry, that are not easily divisible and can be mapped onto only a small number of processors that must operate interdependently while the application is running. Current applications addressing problems like these have limited scalability and perform better on SMP (symmetric multiprocessing) computers and traditional supercomputers than on clusters. It follows from these examples that no single HPC computer architecture is ideal for all applications.

"Huge software costs."

"Software costs are too high."

"Very high software costs."

"Hardware growth is not being matched by software growth in scaling, performance, or business model."

"Hardware growth is not being matched by software growth in scaling, performance, or business model."

"The hardware is not there yet."

"The single image architectures [SMPs, traditional supercomputers] don't scale high enough, and our problems don't decompose for cluster architectures."

"Current hardware and software architectures aren't ready to take full advantage of massively parallel operations."

"High-performance hardware and software aren't available."

"The functionality is missing."

"We require compatibility with other groups in our company."

Need for a Petascale Computer

Three-quarters of the industrial firms (73%) said they could make use of a petascale computer (see Figure 16) to run today's crucial problems faster or to tackle next-generation problems of great competitive importance. This is an interesting finding, given that industrial users acquire substantially smaller versions of HPC computers, on average, than do leading government and academic users. But commercial computer purchases are more heavily dictated by budgets, and the fact that industrial firms have more modest HPC budgets than leading government users does not mean the companies have smaller ambitions for applying HPC.

Three-quarters of the industrial firms (73%) said they could make use of a petascale computer.

In fact, the comments accompanying this question cover a wide range of ambitions for using petascale computing capability, from running today's problems with greater resolution to achieving scientific and engineering breakthroughs.

"Yes, to do more extensive analysis."

"We could push the envelope with new science and more refined models."

"We could do calculations to define the parameters and do physics that we don't understand today."

"We could do real-time rendering of full-resolution frames of CG [computer-generated] films. We could also do real-time simulation of clothing, hair, fur, feathers, fluids, etc."

"We could do a full engine simulation."

FIGURE 16

Need for a Petascale Computer

Q. Would you have use for a petascale computer?

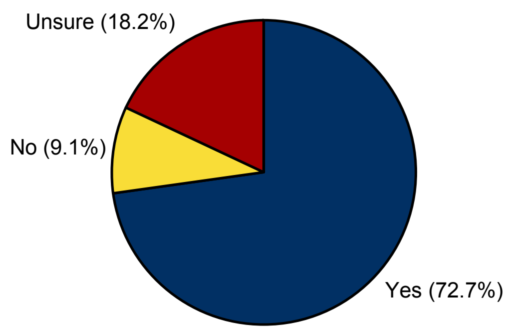

n = 11

Source: IDC, 2006

Running Heterogeneous Problems

The vast majority of the industrial end users (see Figure 17) said they would (83%) or might (91%) use a petascale computer to run heterogeneous problems. IDC defines a heterogeneous problem, also called a multiphysics or coupled or multidisciplinary problem, as one that involves multiple scientific disciplines — for example, studying the complex interaction between the structure of an automobile and the fluid dynamics of air flow around it. HPC users are increasingly interested in solving heterogeneous problems — and successfully doing so could provide substantial competitive advantage — but currently available software and hardware systems are very limited in their ability to address the complexity of this type of problem.

The vast majority of the industrial end users said they would (83%) or might (91%) use a petascale computer to run heterogeneous problems.

FIGURE 17

Running Heterogeneous Problems

Q. Would you use a petascale system to run heterogeneous problems?

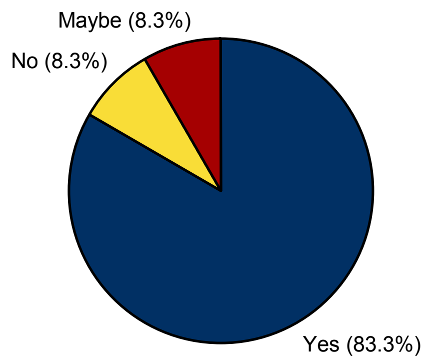

n = 12

Source: IDC, 2006

Explored Heterogeneous Problems with Your ISVs

All but one (91%) of the end users have already discussed their desire to run heterogeneous problems with their ISV suppliers (see Figure 18). This demonstrates that the desire to solve these problems is real and immediate, not simply futuristic. Unfortunately, ISVs by and large lack the ability to address heterogeneous problems today. As the comments below indicate, a few industrial firms are currently running simple heterogeneous problems, but for most commercial end users solving heterogeneous problems remains a more distant, sometimes frustrating goal.

"Yes. We've discussed coupled FEA/CFD models."

"Yes. We explored solid and fluid mechanics with our ISV and didn't reach a satisfactory conclusion. They said it was too much of an investment and I'd have to foot the entire development bill to address my problem. The ISVs won't take the risk."

"Yes, we are currently solving heterogeneous problems."

"Yes. We are running ISV codes coupled with our own internal software and routines."

"Yes, for chemistry and flow."

"No. We have other, more pressing problems, and this level of complexity has a relatively small market."

FIGURE 18

Explored Heterogeneous Problems with ISVs

Q. Have you explored heterogeneous problems with your ISV suppliers?

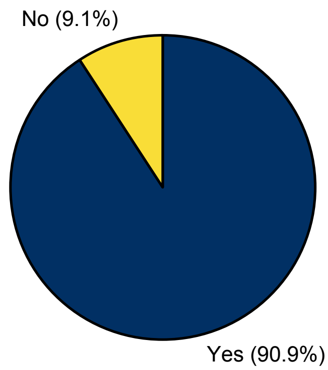

Source: IDC, 2006

Competitive Risks of Standard ISV Codes

As Figure 19 shows, one-third (33%) of the industry respondents replied that using standard ISV codes presents a competitive risk for their companies, while the majority (67%) said it does not. For the minority who replied in the affirmative, the competitive risk is seen as an inability to differentiate the company's products. Here the assumption is that if everyone uses the same ISV software, everyone's products will be indistinguishably similar. In sharp contrast to this assumption, one of the respondents who saw no competitive risk underscored the human factor and noted, "It's how we use the software that gives us a competitive edge."

One-third (33%) of the industry respondents replied that using standard ISV codes presents a competitive risk for their companies.

"Yes. Running industry-standard software limits the visual complexity that is possible [in making films] and limits our creative appetite."

"Yes. It gives our competitors the same capabilities as us."

"No. It's how we use the software that gives us a competitive edge."

"Not using it risks slowing down our CAE development."

FIGURE 19

Competitive Risks of Standard ISV Codes

Q. Does running standard ISV codes pose a competitive risk for your company?

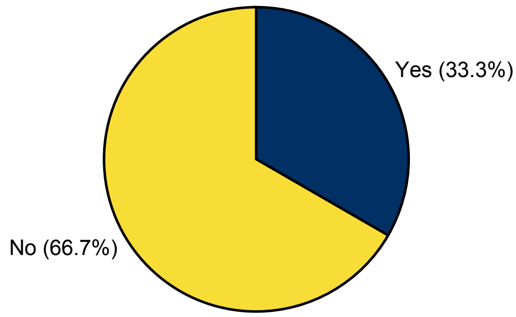

n = 12

Source: IDC, 2006

Views on Partnering

End-User Views of Potential Partnerships for Improving Applications

At this point in the study, we asked the industrial end users to respond to a more comprehensive set of questions regarding potential partners for collaborations to improve application software. Tables 14 and 15 compare the reactions to five partner types: national labs, universities, government agencies, business competitors, and ISV suppliers.

The "total" column, far right, shows that in general, positive responses (47%) outnumbered negative responses (20%) by more than two to one. Nearly half of the industrial firms had favorable views of external partnerships to improve application software. This represents a strong pool for forming collaborative relationships.

A minority of the respondents (27%) had at least "some concerns" about outside partnerships, however. Comments revealed that many of the concerns about partnering understandably had to do with maintaining control and competitive advantage. Some industrial end users wanted an assurance that they would gain competitive advantage through the partnership, such as from exclusive use of the results for a specified period, or from rights (exclusive or shared) to any new intellectual property. Others were more concerned that they might lose competitive advantage by sharing their knowledge and technologies with outside organizations. Still others simply wanted to be assured that the partnerships would really improve their simulation capabilities.

"We'd want exclusive use of the results for a period of time."

"We'd expect shared risk, with joint ownership of the intellectual property or credit for the IP we contribute."

"We're open to partnerships but would want to maintain our competitive advantage in our core industry."

"We'd want a partnership where the scope, goals, and methods are agreed upon in the beginning and we are an active partner."

"Ones that will improve our ability to conduct design analyses."

"We'd want a partnership where the scope, goals, and methods are agreed upon in the beginning and we are an active partner."

TABLE 14

End-User Views of Potential Partnerships for Improving Applications (Number of Respondents)

Q. Please give us your views of various potential types of partners for collaborating with you to improve application software.

	National Labs	Universities	Government Agencies	Competitors	ISV Suppliers	Total
Positive	6	6	5	3	8	28
Some concerns	3	4	4	5	–	16
Neutral	–	–	–	1	1	2
Negative	2	2	3	2	3	12
Not sure	1	–	–	1	–	2
Total	12	12	12	12	12	60

Source: IDC, 2006

TABLE 15

End-User Views of Potential Partnerships for Improving Applications (% of Respondents)

Q. Please give us your views of various potential types of partners for collaborating with you to improve application software.

	National Labs	Universities	Government Agencies	Competitors	ISV Suppliers	Total
Positive	50.0	50.0	41.7	25.0	66.7	46.7
Some concerns	25.0	33.3	33.3	41.7	–	26.7
Neutral	–	–	–	8.3	8.3	3.3
Negative	16.7	16.7	25.0	16.7	25.0	20.0
Not sure	8.3	–	–	8.3	–	3.3
Total	100.0	100.0	100.0	100.0	100.0	100.0

Source: IDC, 2006

Partnering with ISV Suppliers

Among the partner types, ISV suppliers elicited the highest percentage (67%) of positive responses (refer back to Tables 14 and 15) — although they also scored comparatively high in the negative response category (25%). None of the respondents said they had "some concerns" or were "not sure" about working with ISV suppliers. In other words, respondents appeared to know what to expect from ISV collaborations.

"We'd had very good experience across our industry and will continue doing this."

"For the most part, it's problematic. They want me to assume all the risk for development."

"They're focused on the commodity space, not the future of computing. Getting a larger number of less-sophisticated customers is their goal."

Partnering with National Laboratories and Universities

National labs and universities were tied (50% positive responses each) as the second-most-popular potential partners for collaborations related to application development. These two categories also had the same number (17%) of negative responses. The two chief complaints about these partner types were that they tend to overestimate the value of their own ideas and failed to understand commercial market requirements (e.g., ease of use, functionality, intellectual property protection).

"We love working with labs."

"I'm concerned that the labs will see it as a scientific experiment, while we are faced with time-critical real world deliverables."

"We have limited experience working with labs on competitive technologies."

"In our experience, relationships with universities take time to get started. We usually enter into one-to-one agreements."

"Working with universities is okay for research projects."

"Collaborating with universities is problematic. They're obsessed with obtaining value and have a disproportionate idea of the value of their ideas. They don't get the concept of commercial value in the marketplace."

"I worry about the IP risk in working with universities."

Partnering with Government Agencies

Government agencies were also viewed favorably (42% positive responses), though not as favorably as ISV suppliers, national labs, or universities. Where agencies were concerned, the primary fear was that lengthy bureaucratic approval processes ("red tape") and slow working methods might seriously impede progress.

"*Our experience with agencies has been very positive."*

"Positive."

"We have some experience, with a mixed bag of results. It's difficult to generalize."

"We have little direct experience. We've heard from other parties that agencies appear to have conflicts in their mission about collaboration with industry. There may be too much paperwork and red tape."

"They're generally slow and laborious, at least based on our experience. There's too much regulation and too little action. The process and paperwork would likely make it difficult to succeed by our standards."

Partnering with Others in Your Industry

Easily the least-favored type of category for potential collaboration consisted of "competitors in your industry," although even here there were some interested parties (25% positive responses). The comments showed that at least some industrial end users have had successful collaborations with direct competitors. Most of the end users, not surprisingly, viewed the prospect of sharing knowledge and technology with competitors askance.

"We have done so successfully for several years."

"I'd love to work with a noncompetitive automotive or aerospace company but am very wary of our own industry."

"These are the areas of competitive advantage. It is unlikely any of us would share results with enough specificity to be useful."

"Our experience doing this is mixed: some good, some not so good."

"We have done a couple of very large shared projects with our competitors in the area of large longer term R&D that no single company could afford to do by itself."

Collaborating with Non-U.S. Partners

To help complete the picture, we asked the industrial end users whether they had concerns about working with non-U.S. partners. A small majority (54%) replied yes. Comments highlighted the need to safeguard sensitive information in dealings with non-U.S. collaborators (see Table 16).

"We already do. We just need to be aware of sensitive subjects and information."

TABLE 16

Concerns Regarding Working with Non-U.S. Partners

	Number	% of Respondents
Yes	7	53.8
No	6	46.2
Total	13	100.0

Source: IDC, 2006

FOR ADDITIONAL INFORMATION

For additional information, go to:

www.compete.org

www.idc.com

www.hpcuserforum.com

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Council on Competitiveness Study of HPC As An Innovation Driver For Industry

Sponsored by: DARPA and the Council on Competitiveness

Earl Joseph, Ph.D. Christopher G. Willard, Ph.D.
Addison Snell Steve Conway, Conway Communications
Suzy Tichenor, Council on Competitiveness Jei Wu
February 2007

EXECUTIVE SUMMARY

Investigating and advancing the use of HPC to increase U.S. industrial productivity and global competitiveness is the main purpose of the Council on Competitiveness' HPC Initiative, a coordinated program of original research, conferences and workshops that began three years ago. .

The current study, sponsored by the Council on Competitiveness and the Defense Advanced Research Projects Agency (DARPA), was designed to measure the penetration of high performance computing (HPC) among leading companies and their suppliers in selected industries, as a way of helping to determine whether they are adequately equipped to accelerate innovation and competitiveness. For this study, IDC conducted surveys of U.S. firms (and one U.S.-Canadian firm) in the automotive, aerospace, energy and biotech industries. The study also gathered concrete information on the role HPC has already played as an innovation driver in the surveyed companies, and on important problems HPC could enable the companies to solve in the future. With the help of the information and metrics gathered in this study, the Council hopes to establish HPC as a gauge to evaluate the country's capacity for innovation.

American industry is in the midst of a new, twenty-first century industrial revolution, driven by the application of computer technology to industrial and business problems. HPC plays a key role in designing and improving many industrial products—from automobiles to airplanes, pharmaceutical drugs, microprocessors, computers, implantable medical devices, golf clubs and household appliances—as well as industrial-business processes (e.g., finding and extracting oil and gas, manufacturing consumer products, modeling complex financial scenarios and investment instruments, planning store inventories for large retail chains, creating animated films, forecasting the weather). HPC users typically pursue these activities with *virtual prototyping and large-scale data modeling*, i.e., using computers to create digital models of products or processes, and then evaluating and improving the design of the products or processes by manipulating these computer models. Given their broad and expanding range of high-value economic activities, HPC users are increasingly crucial for U.S. innovation, productivity and competitiveness.

Six in 10 of the surveyed companies use HPC to help drive innovation in both their R&D exploration and production work, meaning work directly related to the company's revenue-producing products or services that must to be accomplished under time deadlines. All but one of the firms has U.S. headquarters. Three-quarters of the firms said their HPC usage occurs primarily in the U.S. Another 22% reported that they exploit HPC both in the U.S. and beyond.

HPC has gained enormous importance for most of these firms. All but one view HPC as a strategic asset for innovation and competitiveness. For this group of companies, HPC reduces the costs of engineering and testing; makes possible new ideas and insights, superior products, faster time-to-market and a host of industry-specific advances noted by the survey respondents. The risks without HPC range, according to the surveyed firms, range from the inability to meet product regulations, to falling behind competitors, to "being priced out of the market, completely." Nearly all (86%) said expanded use of HPC could help bring about a dramatic increase in their future innovation.

A substantial majority of the respondents said their managements, customers and owners/shareholders also recognize HPC as an innovation driver. More than half the firms reported that their customers require them to use HPC, although few of the surveyed companies require their own suppliers to use HPC. All but one of the industrial firms said their competitors used HPC.

Of the major factors motivating the surveyed companies' most recent HPC system purchase, the top three all were related to providing new insights or making new problems tractable. The biggest barriers to expanded HPC use are the unavailability of more-capable application software and of people skilled at applying HPC to challenging business and industrial problems.

This study reaffirms the central findings of prior IDC research conducted for the Council on Competitiveness, which revealed that virtually all U.S. businesses that have already adopted HPC consider this technology indispensable for their competitiveness and corporate survival. The current study expands on earlier research by illuminating HPC's specific role in driving innovation.

Heightened competition from other nations has made it more urgent to accelerate innovation and elevate productivity within the nation's private sector. IDC believes that the failure of companies of all sizes to exploit HPC more thoroughly for increased innovation will put major U.S. industries at greater risk—and sacrifice a rare opportunity for the U.S. to make a quantum leap forward in innovation, productivity and competitiveness.

DEFINITIONS

Technical Computing

IDC uses the term "technical computing" to encompass the entire market for computers (and related software and services) employed by scientists, engineers and others to address computationally intensive modeling and simulation problems. Technical computing activities can be found in industry, government and academia. Industrial activities include: automotive and aerospace product development, oil and gas exploration, drug discovery, weather prediction and climate modeling, complex financial modeling, consumer product design and optimization, advanced 3D animation and others. Technical computers range from single-user desktop computers (PCs, MACs and workstations) to supercomputers (a continuous spectrum from entry-level to high-end machines). Technical computing is in contrast to commercial computing as used for business operations such as accounting, payroll, sales, customer relations, transaction processing, human resources and purchasing. Other terms for supercomputers are *technical servers* and *high performance computing (HPC) systems.*

High-Performance Computing

High-performance computing (HPC) is the important subset of the technical computing market that addresses the largest, most challenging modeling and simulation problems. The term encompasses both the activities carried out in this market and the computers used to perform these activities: HPC systems (common synonyms: *supercomputers*, *technical servers*). HPC systems include the full spectrum that extends from entry-level to high-end supercomputers, but exclude single-user desktop computers (PCs, MACs and workstations) that are used for technical computing.

Technical Computing (HPC) Compared to Commercial Computing

Technical computing is in contrast to commercial computing as used for business operations such as accounting, payroll, sales, customer relations, transaction processing, human resources and purchasing. The term technical computing, also called scientific and technical computing, also called HPC or HPTC. IDC uses these terms to refer to all technical computer servers used to solve problems that are computationally intensive or data intensive, and also to refer to the market for these servers and the activities within this market. It excludes single user desktop workstations and PCs.

Innovation

IDC uses the term innovation to refer to creating new ideas, products, inventions, manufacturing processes, risk modeling, supply chain optimization, new services, business process innovations, etc. An innovation is something new that creates value for the innovating organization or its customers. In this study, IDC is interested in innovations that were achieved using HPC, especially those that would have been difficult or impossible to accomplish without HPC. Organizations range from entire companies to single departments within a company.

Entry-Level Supercomputers

IDC defines entry-level supercomputers (also called HPC systems) as servers designed for technical computing that are priced from about $5,000 to as high as $250,000. Entry-level supercomputers may be designed as single computers, or as so-called *clusters* that link together multiple smaller computers.

Virtual Prototyping and Large-Scale Data Modeling

IDC defines virtual prototyping and large-scale data modeling as the use of computers (a) to create digital models of products or processes, and (b) to evaluate and improve the design of the products or processes by manipulating these computer models. A growing number of companies and industries have adopted virtual prototyping and large-scale data modeling as part of their R&D, production computing and complex business processes because virtual prototyping and large-scale data modeling typically are much faster, less expensive and more conducive to new insights than the traditional process of designing and testing a series of physical prototypes.

TABLE OF CONTENTS

P

LIST OF TABLES

P

SITUATION OVERVIEW

Motivations For This Study

The aim of this study was to measure the penetration of high performance computing (HPC) among leading companies and their suppliers in selected industries, as a way to determine whether they are adequately equipped to accelerate innovation and competitiveness. The study also gathered concrete information on how HPC has helped the surveyed companies and their industries to innovate (success stories), and how their HPC usage has changed over the years. With these metrics, the Council hopes to establish HPC as a gauge to evaluate the country's capacity for innovation. For this study, IDC conducted surveys of U.S. firms in the automotive, aerospace, energy and biotech industries.

Sample Qualifiers:

Do you currently have any technical servers, computers or clusters installed at your site and which are used to directly support engineering, analysis and/or scientific efforts such as simulation, modeling, scientific research, computer-aided design, financial modeling, economic analysis, etc.? (excluding desktops and single user workstations)

Survey Profile

For this study, IDC approached companies without knowing, in many cases, whether or not the companies currently used HPC. As Table 1 illustrates, slightly more than two-thirds (69%) of the contacted firms reported that they were HPC users, while the remaining one-third (31%) said they were not. Because IDC did not administer the survey to non-HPC users and some of the HPC users had not completed the survey at the time of this writing, the current survey results represent 23 firms, just under half (45%) all the companies (including non-HPC users) contacted for this study.

TABLE 1

Percent of Contacts Using HPC

	Count	Percent
Using HPC, completed survey	23	45.1%
Using HPC, will complete survey	12	23.5%
Not using HPC	16	31.4%
Total	51	100.0%

Source: IDC 2007

The 23 surveyed industrial firms had substantial HPC experience (Table Q01), averaging 18 years of HPC usage and in no case reporting less than five years' experience with high performance computing.

TABLE Q01

Years of HPC Experience

	Average	Minimum	Maximum
Years Using HPC	18.2	5	40

Note: N=23

Source: IDC 2007

SURVEY RESULTS

Use Of HPC

Only two of the 23 respondents (9%) accessed HPC cycles from sources outside of their own firms on an outsourced basis, while more than nine in 10 did not. This lopsided split was not surprising. IDC's 2005 global study, "New Perspectives on HPC Usage, Trends, and Applications for Industrial Users," found that industry use of grid computers linking multiple organizations, and of compute cycles purchased from external sources and delivered via networks, together amounted to just 3.9% in that year and was projected to grow marginally to 4.4% of total HPC industry cycles in 2006. IDC knows from other research that for industrial firms, network security is often a barrier to using HPC on an outsourced basis.

As for the more extreme measure of outsourcing not just access to HPC but HPC activity itself, some firms believe that doing this sacrifices the learning and quality control that is needed for innovation and competitive survival.

"Some companies do their processing out of house, but we feel when this is done out of house you can lose the understanding of the results."

"[If we outsourced HPC, we'd be totally dependant on contractors and there would be lack of quality control over their work."

TABLE Q02

Use of External HPC Cycles

	Count	Percent
Yes	2	8.7%
No	21	91.3%

TABLE Q02

Use of External HPC Cycles

	Count	Percent

Note: N=23

Source: IDC 2007

IDC next asked the industry respondents about the sizes of their largest installed technical computers (Table Q03). On average, the computers had 1,215 processors and peak performance of 9,448 gigaflops, or about 7.8 gigaflops per processor. If that average figure seems outsized in relation to the gigaflop ratings of current standard x86 microprocessors, note that the average presumably includes a substantial number of multi-core processors, as well as vector processors with peak performance of as much as 18 gigaflops each. The average central memory size was 2.8 teraflops. The average number of nodes on the industrial users' largest installed computers was 517, which translates to an average of about 2.4 processors per node.

Although the 1,215-processor average size of the largest technical computers installed at the surveyed industrial sites is impressive, that figure pales compared to the average 4,148-processor, 954-node configuration that IDC found for the entire technical computing market (government, university, industry) in the February 2006 multi-client study, "The Cluster Revolution in Technical Computing." The 1,215-processor systems are only 29% as large as the 4,148-processor configurations. Because of their greater dependence on ISV applications with limited scalability, and not least because of their more modest HPC budgets, industrial HPC users typically acquire smaller versions of HPC computer servers than do, for instance, leading government user organizations, such as national laboratories, and large university-based HPC centers.

It is important to note that processor counts are only rough approximations of computing power, because not all processors are equally powerful. The peak performance of single-core processors found in some contemporary HPC servers ranges from a few gigaflops for some microprocessors to nearly 20 gigaflops for the leading vector processors. When the fast-growing contingent of dual-core and other multiple-core processors is added to the picture, the gigaflop range becomes even greater.

TABLE Q03

Size of Largest Computer Installed

	Average	N=
By Number of Processors	1,215	23

TABLE Q03

Size of Largest Computer Installed

	Average	N=
By Peak GFlops	9,448	13
By GB of Memory	2,818	18
By Number of Nodes	517	22

Source: IDC 2007

HPC Usage And Innovation

Well over half (61%) of the surveyed firms applied HPC for both research and development (R&D)—i.e., upstream exploration and innovation activities that precede specific product development—and for engineering design as part of the product development process (see Table Q05). Nearly half of the companies (48%) said that most of their HPC usage is for R&D. A few declared that they exploit HPC even further "downstream," in the manufacturing and product production processes, or for challenging data management. Interestingly, about one in six of the respondents (17%) require their suppliers to line up with them technologically by using compatible HPC hardware and software.

TABLE Q05

How broadly is HPC or technical computing used in your organization?

	Responses	Percentage
Mostly in R&D	11	47.8%
Both R&D and engineering (design)	14	60.9%
Also heavily in manufacturing or production	3	13.0%
Also in large scale data management	3	13.0%
We also require our suppliers to use compatible HPC computers/software	4	17.4%
Other areas	4	17.4%

TABLE Q05

How broadly is HPC or technical computing used in your organization?

	Responses	Percentage

N=23

Note: Multiple Responses Allow

Source: IDC 2007

The responses were similar (Table Q06) when IDC asked where the companies' HPC usage affected innovation. About six out of 10 (59%) said they use HPC to help drive innovation in both their R&D exploration and production work, meaning work directly related to the company's revenue-producing products or services that must to be accomplished under time deadlines. A lesser, still substantial percentage (41%) reported focusing HPC on innovation within R&D only, while a few used HPC to innovate data mining and analysis.

TABLE Q06

Where do you focus your use of HPC for innovation?

	Responses	Percentage
Only in R&D	9	40.9%
R&D and in production	13	59.1%
R&D, Production and supply chain management	0	0.0%
Large scale data mining and/or analysis	3	13.6%
Other areas	1	4.5%

N=22

Note: Multiple Responses Allow

Source: IDC 2007

All but two of the 23 surveyed companies were headquartered in the United States, and all except one of the technical computing data centers represented in the study was located in the U.S. (the exception: Canada). Not surprisingly, then, Table Q07 shows that three-quarters (74%) of the firms said their HPC usage occurred primarily in the U.S. and another 22% reported that they exploited HPC both in the U.S. and beyond. The total for these two categories was 96%. Only the Canadian respondent said his firm's' HPC activities took place entirely outside the U.S.

TABLE Q07

Geographical Usage of HPC

	Count	Percent
Primarily U.S.	17	73.9%
International	1	4.3%
Mixed	5	21.7%

N=23

Source: IDC 2007

All but one (96%) of the industrial firms viewed HPC as a strategic asset for innovation and competitiveness (Table Q08).

TABLE Q08

HPC Viewed as Strategic Asset Linked to Innovation and Competitiveness

	Count	Percent
Yes	22	95.7%
No	1	4.3%

Note: N=23

Source: IDC 2007

As Table Q09 indicates, about three-quarters of the surveyed industrial firms replied that HPC directly benefits their bottom lines, competitiveness, productivity and innovation. The following comments are representative of the many specific responses about HPC's strategic contribution to innovation and competitiveness:

"HPC reduces computation time, which adds quality and value to the products we sell."

"HPC is integral to the design of vehicles. This is no longer possible without HPC!"

"HPC is an indispensable tool for biotech research. We could not do the work without HPC."

"Customers…continually seek companies that are current with the latest HPC developments and architectures."

"In our company, R&D is seen as having value, but HPC is seen only as a tool."

These additional comments are representative of the responses to the four categories listed in Table Q09:

Bottom line:

"Our customers see an impact on the bottom line through use of numerical simulation to complement field testing, and to reduce the cost and need for building physical prototypes."

"HPC enables increased efficiency and reduced man-hours for analysis and design."

"HPC had made a significant improvement in the fuel efficiency of engines."

Innovation:

"HPC is essential for analyzing the massive volume of data and discovering [drug] pathways, binding sites, and the structure and function of proteins."

"HPC allows problem sizes to push the envelope and allows our users to focus on R&D and innovation instead of worrying about computer resources."

"Accelerated research using HPC opens our eyes to new ways of working on some difficult problems."

Productivity and Competitiveness:

"HPC is an indispensable part of the "minimum ante" for anyone to pursue current biotech research."

"HPC has reduced the product introduction cycle from years to months."

"Once we integrated everything for our scientists, they gained back about one hour per day in which to do other work. Since they are very highly paid, that translates into a lot of savings and increased productivity."

"Without HPC resources, it would be impossible to analyze the massive volume of raw research data, or to explore various simulations, predictions, and models."

TABLE Q09

Direct Benefits of HPC to Organization

Benefit	Count	% of Respondents
Improve bottom line	17	73.9%
Increase competitiveness	18	78.3%
Increase productivity	17	73.9%
Accelerate innovation	17	73.9%

TABLE Q09

Direct Benefits of HPC to Organization

Benefit	Count	% of Respondents

Note: N=23. Multiple responses permitted.

Source: IDC 2007

Specific Innovation And HPC Questions

In the survey, IDC asked respondents to elaborate on HPC's role vis-à-vis innovation. The following series of questions probe this topic more deeply.

1. From an INNOVATION viewpoint, what can you do today that you couldn't do before you had HPC computers? (The following specific comments are representative of responses to this question.)

   "We…can reduce physical [airplane] testing (wind tunnel and flight tests) with improved HPC simulation."

   "The short answer is that we can now do analysis before we build the first [product] prototype."

   "We have the ability to predict [vehicle] safety performance in a variety of scenarios [and can do] component durability prediction."

   "We can dock thousands of molecules into proteins. There is no way we could do it without HPC hardware."

   "Historically, [seismic processing] was done in two-dimensional (x & y) grids. In the 1990s, better hardware allowed for three-dimensional modeling. This provides better understanding of velocity, and better resolution."

   "We can chemical compound modeling. Before, we had to make and test models by hand."

   "We can simulate large-scale clinical trials and search for new bio markers."

   "With each new [HPC] system we bring into the facility, we grow our problems and create new problems that we previously could not tackle."

2. What INNOVATION risks do you have if you DON'T have access to HPC computers systems/tools? Representative comments:

   "We wouldn't be able to create new products in a timely manner."

   "We would risk failing government regulations for [vehicle] safety and emissions."

   "We would be priced out of the market, completely."

"There would be no way to keep up with the rest of the industry."

"Research would become too theoretical, creating too many "ivory tower" algorithms. Without HPC, we wouldn't be able to see if it works in practice."

"Our users would not be able to consider alternative, possibly better, designs."

"It would be fundamentally impossible to pursue research in biotechnology without HPC."

3. What ORGANIZATION OR COMPETITIVE concerns do you have if you DON'T have access to HPC computers systems/tools? Representative comments:

   "In the race to be the first to discover a new drug, you need every possible asset to be innovative. You lose on the business end if you aren't first to market.

   "We will not be able to compete without new ideas. Our competitive position [using HPC] allows us to earn a position with national oil companies when we compete for access to new resources."

   "We will not be able to pass internal design audits without completing certain analysis that is dependent on HPC."

   It would take much longer to perform experiments, and some experiments are not feasible HPC."

   "We'd be smaller company."

   "We would need results from people power, which takes longer. We would lose our competitive edge."

4. What PRODUCT risks do you have if you DON'T have access to HPC computers systems/tools? Representative comments:

   "Products which don't include the integration of HPC capabilities will not be seen as providing the best price performance."

   "We would have Inferior, uncompetitive software products."

   "Without quality seismic data, we may drill unnecessary or poor wells, a very costly problem."

   "Someone else will get there first. First to market usually gets 60% of the business, even if second and third are better products."

   "We use a lot of exotic materials and we must understand how they behave. That is difficult without HPC."

5. Thinking about the important computational problems that you can't solve today, If you could solve these problems, how would it make you more innovative and/or competitive? Representative comments:

"We would have better airplane designs, more derivatives [designs], and improved fuel efficiency."

"We could make our organization much more productive by doing right things the first time: predictability."

"It could get us past the limitations to reach the 'holy grail' of the broad spectrum antiviral."

If we could address adverse events and side effects of drugs. Better. That would be a colossal leap forward."

"This would over time improve our ability to understand the complicated physics taking place in aircraft engines and allow us to improve our designs."

It would further research and understanding into human health, diseases, and therapeutics.

"We could have really significant breakthrough in areas of combustion, fuel efficiency and emissions."

"Whole genome analysis is virtually impossible. If we could analyze this data, it would be a huge advantage."

"It is not going to make us more competitive, because we don't have any problems today that we don't run.

"We believe we are solving everything today that we have a need to solve."

6. What is keeping you from acquiring the HPC resources to solve these problems?

"Technology barriers for scaling the applications".

"Software and understanding of how to model the problems."

"The cost of software licenses."

"Cost and physical/environmental limits."

"Human capital: engineers who can exploit the HPC resources."

"Lack of understanding of the business benefits by upper management."

"Nothing. There are problems out there that are simply not solvable today."

"Nothing. It's not the HPC resources, it's the science."

Benefits Of HPC

7. What would be the additional benefits to innovation if you could increase your use of HPC overnight? Representative comments:

   "Better designs of our products."

   "Reduced physical testing and faster time to market."

   "To be able to look at a potential broad-spectrum antiviral and see if there are compounds that are efficacious across multiple viruses."

   For our customers, greater use of parametric studies and validation through the use of multiple codes and comparison of their results. The increased ability to do "what if" scenarios that are currently cost- or time-prohibitive."

   "Pushing the envelope on problem sizes. Allowing models to be run that have never been possible before."

   "New areas of simulation for manufacturing and production processes, to help us better understand propellants & chemistry."

Recognition Of HPC Importance By Others

Table Q12 depicts the extent to which respondents said HPC is recognized as an innovation driver by key corporate constituencies. Virtually all (96%) of the respondents replied that their managements appreciated this aspect of HPC. IDC knows from recent studies conducted for the Council on Competitiveness and other clients that HPC users within industrial firms often cite the difficulty of justifying HPC purchases to senior managements. It is important to distinguish managements' recognition of HPC's contribution from their budgetary decisions affecting HPC. As is well known, not everything of value within companies can be funded at ideal levels. Sometimes, companies cannot justify things they would like to have, in the face of competing budget priorities. Of course, there are also some cases in which a company's management simply cannot understand the benefits HPC adoption might bring.

A substantial majority (more than 70%) of the respondents said customers and owners/shareholders also see HPC as an innovation driver.

TABLE Q12

Recognition of HPC as a Driver for Innovation

	Yes	No	Yes %
By management	22	1	95.7%

TABLE Q12

Recognition of HPC as a Driver for Innovation

	Yes	No	Yes %
By customers	15	6	71.4%
By owners/shareholders/investors	16	6	72.7%

N=23

Source: IDC 2007

Surprisingly, more than half of the industrial firms (60%) said that their customers required them to use HPC (Table Q13). "Require" in this instance might mean that customers have provided explicit direction to use HPC, or simply that the customers' needs (requirements) cannot be met without the use of HPC. In any case, for 60% of the companies in the current study, an external customer requirement in one form or another already existed. Even so, the remaining 40% of the firms needed no customer requirement to begin exploiting HPC. As noted further on (see Table Q17), virtually all (96%) of the firms said that competitors were already using HPC. For the firms that adopted HPC with a customer requirement to do so, competitors' use of HPC may have been an adequate motivator.

TABLE Q13

Do Customers Require HPC Usage?

	Count	Percent
Yes	12	60.0%
No	8	40.0%

N=20

Source: IDC 2007

Table Q14 depicts customers' attitudes toward HPC in more detail. In nearly half the cases (46%), respondents said customers did not care whether their companies used HPC or not. In cases where customers did care, nearly one-quarter of the companies (23%) reported that their use of HPC was a primary competitive differentiator in business relationships, and another 14% said HPC provided at least some competitive differentiation for them. Together, these two response categories were noted by more than one in three respondents (36%). Another 18% of the respondents said their use of HPC was not a differentiatior with customers, but got them into the game by putting them on a par with their competitors.

TABLE Q14

Customers' View of HPC Usage

	Count	% of Respondents
Major reason for doing business with supplier	5	22.7%
Reason for doing business with supplier	3	13.6%
Supplier is considered on par with others in industry	4	18.2%
Customer is indifferent to supplier's HPC usage	10	45.5%

N=22

Source: IDC 2007

The great majority of the surveyed companies (86%) did not require their own suppliers to use HPC (see Table Q15), although at least a few did.

TABLE Q15

Require Suppliers to Use HPC?

	Count	Percent
Yes	3	13.6%
No	19	86.4%

N=23

Source: IDC 2007

Earlier (Table Q12), nearly three-quarters (73%) of the respondents said that their companies' owners/shareholders/investors recognized HPC's role in driving innovation. As Table 16 shows, about one-third of the companies (35%) reported that their investors actually require them to use HPC. A percentage this high was surprising, because it is not typical, at least not in larger publicly held companies, for shareholders to require the use of specific technologies. It was not within the scope of this study to identify in what form this requirement occurred. It is not difficult to imagine, however, that company owners, who exert more influence than average shareholders, might require HPC usage; or that prospective investors would discover, in performing due diligence, that HPC usage was a *sine qua non* for firms in certain industries, and then insist on this as a prerequisite for their investment.

TABLE Q16

Do your investors/shareholders require you to use HPC?

	Count	Percent
Yes	8	34.8%
No	15	65.2%

N=23

Source: IDC 2007

HPC Usage Compared To Competitors

All but one of the surveyed industrial firms (96%) said their competitors used HPC, and the remaining respondent was unsure (Table Q17).

TABLE Q17

Q17a Do your competitors use HPC?

	Count	Percent
Yes	22	95.7%
No	0	0.0%
Unsure	1	4.3%

N=23

Source: IDC 2007

Just two of the surveyed companies (9%) believed that they used HPC more effectively than their competitors, while two-thirds of the firms (65%) thought their HPC usage was only as effective (35%) or less effective (30%) than their rivals (see Table Q18).

TABLE Q18

Q17 b&C: Do your use HPC more or less effectively than your competitors do?

	U.S. Competitors		International	
	Count	%	Count	%
More effectively	2	8.7%	2	8.7%

TABLE Q18

Q17 b&C: Do your use HPC more or less effectively than your competitors do?

	U.S. Competitors		International	
Same, on average	11	47.8%	8	34.8%
Less effectively	7	30.4%	7	30.4%
Unknown	3	13.0%	6	26.1%

Note: N=23 for U.S.; N=23 for international.

Source: IDC 2007

HPC Usage Drivers

Table Q22 summarizes the major factors motivating the surveyed companies' most recent HPC system purchase. Of these, the top three all were related to providing new insights or making new problems tractable. The most frequently mentioned choice was providing greater insight into current problems (cited by 83% of respondents). Allowing currently intractable problems to be addressed was another popular choice (70% of respondents), as was creating new ideas of inventions (52%). Despite the high proportion of the companies noting earlier on that customers recognized HPC's contribution to innovation, fewer than one in four of the companies (22%) said in response to the current question that external requirements played a major role in their last HPC system purchase.

TABLE Q22

Major drivers for last system purchase

	Count	% of respondents	
Provide additional insight into current problems	19	82.6%	b
Help to solve new problems that cannot be addressed practically by other means	16	69.6%	a
Create new ideas or inventions	12	52.2%	d
Risk modeling	8	34.8%	f
Large scale data mining or analysis	8	34.8%	g
Meet external requirements	5	21.7%	c
Develop new services or business process innovations	4	17.4%	l
Create new manufacturing processes	3	13.0%	e

TABLE Q22

Major drivers for last system purchase

	Count	% of respondents	
Supply chain optimization	0	0.0%	h
Other	3	13.0%	j

Note: N=23. Multiple responses permitted.

Source: IDC 2007

Among the top reasons for adopting HPC (Table Q23), little separated the top seven choices: developing better products/services, doing better analysis, improved competitiveness, testing ideas faster, bringing products to market faster, and speeding innovation. Each of these choices earned at least a 4.0 importance rating on the scale where 5 meant most important of all. When choices are grouped this closely together, it is usually safe to conclude that many or all of them are interrelated in the user's mind.

TABLE Q23

Rate the following reasons for adopting HPC (5=most important; 1=least important.)

	Average	Number responding
Ability to build better products/services	4.6	21
Ability to do more/better engineering, science, or analysis	4.5	23
Increased competitiveness	4.5	22
Ability to improve quality	4.2	22
Ability to test ideas faster, compared to live tests	4.2	22
Faster time to market	4.1	23
Accelerate innovation	4.0	23
Increased profitability or lower costs	3.8	22
Large scale data mining or analysis	2.8	22
Supply chain optimization	1.4	20

N=23

Source: IDC 2007

Ways HPC Resources Are Accessed

Earlier in this study (Table Q02), we learned that only two of the 23 respondents (9%) accessed HPC cycles from sources outside of their own firms on an outsourced basis, while more than nine in 10 did not. Table Q24 delved more deeply into the ways in which respondents accessed HPC resources.

Consistent with the earlier findings, only four respondents stated that they access HPC resources from outside their own firms (Table Q24), with two of them using a Grid, and the other two using other approaches. Of the majority who obtain resources from within their companies, most (52% of all respondents) do so through network connections to other groups within the company; other have grid connections to other groups within their own organizations.

TABLE Q24

Ways in which HPC resources are access from outside the group

	Count	% of respondents
Grid connection to other organization	2	8.7%
Grid connection to other group(s) within same organization	3	13.0%
Network connection to other group(s)	12	52.2%
Other method for external access	2	8.7%
No external access; only access to HPC resources within the group	9	39.1%

Note: N=23. Multiple responses permitted.

Source: IDC 2007

Barriers To Expanded Use Of HPC

The greatest barriers to additional HPC usage (Table Q25) were cost-related: budget constraints (cited by 74% of respondents) and ISV software expense (52%). Environmental issues (facility space, power and cooling) was also mentioned frequently (44% of respondents), as were limited performance and technical expertise (39% each) and the availability and maturity of application software (35%). Interestingly, system management and ease-of-use figured less prominently (26%) as a barrier to greater HPC usage.

TABLE Q25

Barriers to Greater HPC Adoption

	Count	% of respondents	

TABLE Q25

Barriers to Greater HPC Adoption

	Count	% of respondents	
Financial (budgets, system costs, other costs)	17	73.9%	a
Third-party software costs	12	52.2%	b
Facilities issues: space, power, cooling	10	43.5%	d
Technical limitations (performance)	9	39.1%	g
Skilled staff / experts / personnel	9	39.1%	k
Application availability or maturity	8	34.8%	h
Hard to justify expense to upper management	6	26.1%	c
Ease of use / system management	6	26.1%	e
Supported data storage mechanisms	4	17.4%	I
Maintenance or uptime issues	3	13.0%	j
Complexity to expand and/or use	2	8.7%	f
Other	1	4.3%	l

Note: N=23. Multiple responses permitted.

Source: IDC 2007

Source Of HPC Applications

Table 29 depicts the sources of application software used by the surveyed industrial firms. The importance for industry of accurate, relevant mathematical models of physical processes—and of software applications embodying these models—can hardly be overstated. Solving problems—often by running the same application repeatedly to close in on an optimal solution—can be far more time-critical for industry than for government and university organizations pursuing scientific research. Industrial and other business firms are driven by external competition in a never-ending race to be first to market with the best products. In these battles for global market supremacy, more-capable computing resources can translate into faster time-to-market, superior product quality, and novel insights that create lasting competitive advantage.

Table 26 shows that the predominant sources of application software for the surveyed group of industrial HPC users were codes developed by the firms themselves (averaging 53% of the software used) and application software purchased from independent software vendors, or ISVs (45%). About one in five application codes used by the firms were freeware (21%), and a smaller percentage came from through collaborations with universities or other not-for-profit organizations.

In the 2005 "Council on Competitiveness Study of ISVs Serving the High Performance Computing Market, Part B: End-User Perspectives," IDC found that a substantial majority (86%) of the industrial firms indicated they were already developing some application software on their own. The accompanying comments told the real story, however: many of the businesses tackled application software only to a limited extent, and even then often with reluctance. In most cases, application software the industrial firms planned to develop in-house was not intended to replace software they acquired from ISV suppliers. Attempting to develop application software on their own, the companies said, exposed them to substantial costs, competitive risks and major business disruption, up to and including the possibility of not being able to introduce or sell their products in the marketplace. IDC knows from other research that some industries (e.g., financial services, entertainment) typically prefer to develop and maintain proprietary application software for competitive reasons, but most industrial users do this mainly because commercial software with needed capabilities is simply unavailable in the market place.

TABLE Q26

Source of Applications

Source	Percent	Minimum	Maximum
Developed in-house	52.6%	5%	100%
Purchased from external source (e.g. ISV)	44.7%	2%	100%
Acquired externally at no cost: free software	20.7%	10%	93%
Collaboration with academia or research consortium	9.4%	5%	33%

N=23

Source: IDC 2007

Responses to: What would it take to dramatically increase your level of innovation?

The following are representative responses to this question:

"We believe we have created a leadership position in our industry, and will work to extend this advantage."

"Higher budgets."

"Commitment of $2-3 million for emphasis in HPC."

"Cheaper software licenses."

"Better predictability. Better understanding of the basic principals that go into molecular binding."

"Higher-capability [HPC] systems."

"People and petaflop system performance."

"Human capital. More engineers and greater expertise."

"People who know how to use the HPC tools."

"More work requirements from our clients."

"HPC alone will not increase the level of innovation. There is a demand for manpower to develop these ideas and a large amount of testing and validation required."

Can HPC play any role in making a dramatic increase in your innovation?

Nearly all (86%) of the surveyed industrial firms believed that HPC could help bring about a dramatic increase in their innovation (see Table 28). This is a strong testament to the credibility that HPC usage has gained within these companies.

Representative comments related to Table 28.

"Yes, assuming funding [becomes available] to modernize legacy codes to enable use of more concurrent processors or accelerators, which would enable increased fidelity and larger problem sizes with reasonable time to solution."

"Yes, if coupled with new approaches to research that really take advantage of the power."

"Yes, if we could drive turnaround time down a lot, into hours (from days & weeks), we could do a lot more things that are not currently practical."

"HPC plays an important role, but I'm not sure dramatic is the word that would describe its role."

"We need to continue to provide for both research groups and for production processing. Upper management will never see HPC as added value, just a tool, but we need to maintain that tool."

"We need the science perspective, not hardware. The problem is figuring out how to get the answers you want, to be able to predict what you want to predict."

TABLE Q28

Can HPC play any role in making a dramatic increase in your innovation?

	Count	% of respondents
Yes	18	85.7%
Uncertain/Maybe	2	9.5%
No	1	4.8%

Note: N=21

Source: IDC 2007

APPENDIX: QUESTIONS AS ASKED IN THE STUDY

Please respond for your division or your company overall based on how your group uses HPC within your organization, e.g. if your group is a separate division with its own HPC resources, respond based on only your division.

BACKGROUND QUESTIONS:

1. How long has your company been using technical computing or HPC (in years)?

   _________ (years)

2. Do you purchase HPC cycles outside of your group or company?:

   ___Yes, ___No

3. What is the size of your largest technical computer?

   - ❑ In number of processors: _____________
   - ❑ In peak GFLOPS: __________________
   - ❑ In total memory (GBs): ______________
   - ❑ In number of nodes: _________________

4. What are the main applications or areas of use for your company's HPC computer:

5. How broadly is HPC or technical computing used in your organization?

   - ❑ ____ Mostly in R&D
   - ❑ ____ Both R&D and engineering (design)

- ❑ ____ Also heavily in manufacturing or production
- ❑ ____ Also in large scale data management
- ❑ ____ Also in _____________,
- ❑ ____ We also require our suppliers to use compatible HPC computers/software

6. Where do you focus your use HPC for innovation?

- ❑ ___ Only in R&D
- ❑ ___ R&D and in production
- ❑ ___ R&D, Production and supply chain management
- ❑ ___ Large scale data mining and/or analysis
- ❑ ___ Please list all areas that you use HPC for innovation:

7. If you have international subsidiaries or parents —> do you primarily do HPC in the US, outside the US or at multiple locations in multiple countries — please explain:

The Linkage of HPC and Innovation

8. Do you view computational capability or HPC as a strategic asset and do you link it to your overall competitiveness and innovation within your industry? ___Yes, ___No

9. What has been the direct benefit of HPC on your organization?

- ❑ ___ Impact on bottom line—can you quantify:
- ❑ ___ Increased competitiveness—describe how:
- ❑ ___ Increased productivity—in what way:
- ❑ ___ Accelerate Innovation—in what way:
- ❑ ___ Other: ______________________________

10. Thinking about how your organization creates value / remains competitive — How does HPC help accomplish these goals? (For either your company or your business unit or your department)

11. Innovation and HPC:

- ❑ From an INNOVATION viewpoint, what can you do today that you couldn't do before you had HPC computers?

- ❑ What INNOVATION risks do you have if you DON'T have access to HPC computers systems/tools?

- ❑ What ORGANIZATION OR COMPETITIVE concerns do you have if you DON'T have access to HPC computers systems/tools?

- ❑ What PRODUCT risks do you have if you DON'T have access to HPC computers systems/tools?

12. Do others recognize HPC as a driver for your innovation?

    - ❑ Does your management believe that HPC drives innovation? ___Yes, ___No.

    - ❑ Do your customers recognize it? ___Yes, ___No.

    - ❑ Do your investors, shareholders or owners recognize it? ___Yes, ___No.

13. Do your customers require you to use HPC? ___Yes, ___No —> Please explain:

14. Do your customers view your use of HPC as an advantage or just required to be competitive?:

    - ❑ ___ They see our use of HPC as a **major reason** they do business with us

    - ❑ ___ They see our use of HPC as a **reason** they do business with us

    - ❑ ___ They see our use of HPC as a similar to others in our industry

    - ❑ ___ They are NOT aware of our use of HPC or they don't care as long as we meet their requirements

15. Do you require your suppliers to use HPC?

    ___Yes, ___No —> Please explain:

16. Do your investors, shareholders or owners require you to use HPC?

    ___Yes, ___No —> Please explain:

17. a) Do your competitors use HPC?

    ___Yes, ___No —> Please explain:

☒ 17) b) Do your competitors **in the US** use HPC more effectively or less effectively than your group? ___Better, ___Less, ___Same —> Please explain:

☒ 17) c) Do your **INTERNATIONAL** competitors use HPC more effectively or less effectively than your group? ___Better, ___Less, ___Same —> Please explain:

18. Is your organization using HPC tools as aggressively as it could?: _____ Yes/No Please explain:

19. What would be the additional benefits to innovation if you could increase your use of HPC overnight?

20. Thinking about the important computational problems that you have today, that you can't solve today — If you could solve these problems, how would it make you more innovative and/or competitive?

21. What is keeping you from acquiring the HPC resources to solve these problems?

22. What **were the major reasons** for purchasing your last HPC system? (multiple responses are allowed):

- ❑ ____ Help to solve new problems that cannot be practically addressed through other means
- ❑ ____ Provide additional insight into current problems (i.e. better understanding of problem characteristics and solution spaces prior to physical test or experimentations), or to address current problems more efficiently (i.e. faster time to solution, lower cost, etc.)
- ❑ ____ Meet external requirements (i.e. regulatory data requirements, standards of practice, etc.)
- ❑ ____ Create new ideas or inventions
- ❑ ____ Create new manufacturing processes
- ❑ ____ Risk modeling
- ❑ ____ Large scale data mining and/or analysis
- ❑ ____ Supply chain optimization
- ❑ ____ Develop new services and/or business process innovations
- ❑ ____ Other, Please explain:_______________________________

23. Please rate the following **potential reasons** for adopting HPC computers or expanding your use of HPC systems, in terms of their importance to your organization or division.

Use the following scale:
5 = Very important
4 = Important
3 = Sometimes important
2 = Rarely important
1 = Unimportant

- ❑ ____ Increased competitiveness
- ❑ ____ Ability to build better products and/or services
- ❑ ____ Ability to improve quality

- ____ Ability to test ideas faster compared to live tests
- ____ Ability to do new more/better analysis, engineering or science
- ____ Faster time-to-market
- ____ Increase profitability, or lower costs
- ____ Large scale data mining and/or analysis
- ____ Supply chain optimization
- ____ Accelerate innovation
- ____ Other: ________________________________

24. Do you also use HPC resources over a Grid or network?:

- ___ We use a Grid to access HPC resources from other organizations
- ___ We use a Grid to access HPC resources from other parts of OUR organization
- ___ We have a network connection to HPC resources from other parts of OUR organization
- ___ We use this approach to access other HPC resources:
- ___ We only use the HPC resources within our group

25. What do you see as the **barriers** to broader HPC adoption for your organization —> check all that apply:

- _____ Financial — budgets, system costs, other costs
- _____ 3rd Party Software costs
- _____ Budgets – upper management doesn’t appreciate the value/hard to justify the expense with upper management
- _____ Space limitations, facility issues power, cooling
- _____ Ease-of-use — System management capability -- Management software
- _____ Complexity to expand and/or use
- _____ Technical limitations — system performance, interconnect performance, complexity/cable, cards, switches
- _____ Application availability/ Lack of maturity of the solution

- ❑ _____ Supported data storage mechanisms (Databases, Parallel file systems etc)
- ❑ _____ Maintenance/Availability issues
- ❑ _____ Having a skilled staff and/or other experts available
- ❑ _____ Other (please specify ___________)

26. What percentage of the applications are purchased from external companies/ISVs or developed in-house (based on CPU hours)?

    - ❑ ____ % Applications developed in-house
    - ❑ ____ % Applications PURCHASED from external companies, e.g. ISVs
    - ❑ ____ % Applications obtain at no cost from external sources
    - ❑ ____% Applications acquired through collaborations with academic or research consortia

    Note = Must add up to 100%

SUMMARY QUESTIONS

27. What would it take to dramatically increase your level of innovation?
28. Can HPC play any role in making a dramatic increase in your innovation?