

WHITE PAPER

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

Sponsored by: Defense Advanced Research Projects Agency, Council on Competitiveness

Earl Joseph, Ph. D. Addison Snell
Christopher G. Willard, Ph.D. Suzy Tichenor
Dolores Shaffer Steve Conway
July 2005



An Initiative of the Council on Competitiveness

EXECUTIVE SUMMARY

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. 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.

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.

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.

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

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.

TABLE OF CONTENTS

	P
Executive Summary	2
Key Findings.....	6
Definitions and terminology	12
Application Software (or Application Software Package).....	12
Capability-Class and Capacity-Class HPC Computer Servers	13
HPC	13
Clusters	13
Heterogeneous Problem.....	14
Highly Scalable Systems	14
ISV (Independent Software Vendor).....	14
Last Technology Update.....	14
Legacy Code	14
Major Rewrite	15
Middleware	15
Open Source Software	15
Petascale Computer	15
Scalability	15
Technical Computing and Commercial Computing.....	16
Study Background	16
Methodology	17
Study Limitations	18
Study Results	18
ISV Technical Software Demographics	18
Business Models of ISV Organizations.....	21
Financial Condition of ISV Organizations	23
Top ISVs in Auto/Aero, Bio/Pharma and Oil/Gas.....	30
Age and Status of ISV Applications	35
Current Scalability of ISV Applications.....	38
Ability To Improve Scalability of ISV Applications.....	40
Willingness To Collaborate and Preferred Partners.....	42
Conclusions	46
Appendix A: Analysis Of ISV Application Usage Between Industry, Government and Academia	48
Appendix B: List of the Application Packages AND ISVs in the Study	54

LIST OF TABLES

	P
1 Number of ISV Applications and Companies by Total Company Revenue	10
2 Partners ISVs Selected as Potentially Most Useful, by Application	11
3 Primary Industry of ISV Applications	18
4 Types of Organizations Supporting Each Application	19
5 Organization's Geographical Region, by Application.....	20
6 Parent Company's Geographical Region, by Application	21
7 Ownership Model of Organizations Supporting Each Application	21
8 Pricing Models of ISV Applications.....	22
9 Number of ISV Applications and Companies by Total Company Revenue	24
10 Number of ISV Applications and Companies by Total Company Headcount	25
11 Annual Revenue of the ISV Applications.....	26
12 Five-Year Growth of Total ISV Revenue, by Application.....	27
13 Current Growth Rate of ISV Applications	27
14 Market Share of ISV Application	29
15 Number of Clients for ISV Application	29
16 Number of Licenses for ISV Application	30
17 Original source of ISV Application.....	35
18 Age of ISV Application (Since First Release).....	36
19 Time Since Last Major Technical Update.....	37
20 Typical Number of Processors the ISV Applications Use for Single Jobs	39
21 Ability of Application to Scale to Hundreds of Processors.....	40
22 Ability of Application to Scale to Thousands of Processors.....	41
23 Ability of Application to Scale to Tens of Thousands of Processors.....	41
24 ISV Willingness to Improve Application (by Application Count).....	42
25 Key Factors Needed for ISVs to Improve Applications.....	43
26 ISV Willingness to Develop Partnerships to Improve Applications	44
27 ISV Willingness to Work with U.S. Government.....	44
28 Most Helpful Types of Partners for ISV Applications.....	45
B1 Directory of ISV Applications and Organizations Included in the Study	54

LIST OF FIGURES

	P
1 Original Source of ISV Application	8
2 Percentage of Applications by Type of Ownership	9
3 Willingness to Collaborate with U.S. Government, by Application.....	12
4 Application Pricing Models by Industry.....	23
5 Plotter for Importance Ranking of ISV Applications.....	31
6 Plotter for Importance of Applications in Automotive/Aerospace.....	32
7 Plotter for Importance of Applications in Biotechnology/Pharmaceutical.....	33
8 Plotter for Importance of Applications in Oil/Gas.....	34
9 Original Source of ISV Application by Industry.....	36
10 Most Helpful Types of Partners for ISV Applications by Industry	46
A1 Chart to Show Percentage of Business, Government, and University Usage of HPC Applications by Industry	48
A2 Percentage of Business, Government, and University Usage of HPC Applications in Auto/Aero	50
A3 Percentage of Business, Government, and University Usage of HPC Applications in Bio/Pharma	51
A4 Percentage of Business, Government, and University Usage of HPC Applications in Oil/Gas	52
A5 Percentage of Business, Government, and University Usage of HPC Applications in All Industries.....	53

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.

Few Independent software vendor (ISV) applications today can take advantage of more than 128 processors.

"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 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.

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

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.

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

Total Company Revenue	Companies		Applications	
	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.

"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.)

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

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

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 conference 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. Through 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 reached 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%
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 within 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

Total Company Revenue	Companies		Applications	
	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 or 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

Total Company Headcount	Companies		Applications	
	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%
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



Source: IDC, 2005

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




Source: IDC, 2005

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



Licenses: ○ < 10 | ○ 10 - 100 | ○ 100 - 250 | ○ 250 - 500 | ○ 500 - 1,000 | ○ 1,000 - 10,000 | ○ > 10,000

Market Share: ○ < 5% | ○ 5 - 10% | ○ 10 - 25% | ○ 25 - 50% | ○ > 50%

○ U.S. (circles) | △ Foreign (triangles) | ③ Multiple codes

Source: IDC, 2005



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



Source: IDC, 2005

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%
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 21.) 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 22.) 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 23.) 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 24.) 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 25.) 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%

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 26.)

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 27.)

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 28.)

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%


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.

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

In recent years, an alarming gap has developed between the needs of HPC users and the capabilities of ISV applications.

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.

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.


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.

APPENDIX A: ANALYSIS OF ISV APPLICATION USAGE BETWEEN 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



Source: IDC, 2005

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 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 A2

Percentage of Business, Government, and University Usage of HPC Applications in Auto/Aero




Source: IDC, 2005

Figure A3 shows the mix of sales for ISV applications in the biotechnology and pharmaceutical industries.

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 A3

Percentage of Business, Government, and University Usage of HPC Applications in Bio/Pharma




Source: IDC, 2005

Figure A4 shows the mix of sales for ISV applications in the oil and gas industry.

As with the auto/aero industries, oil/gas HPC applications are sold primarily to businesses, with more government sales than university sales.

FIGURE A4

Percentage of Business, Government, and University Usage of HPC Applications in Oil/Gas




Source: IDC, 2005

Figure A5 combines Figures A2, A3, and A4 and also includes applications for other 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.

FIGURE A5

Percentage of Business, Government, and University Usage of HPC Applications in All Industries



Source: IDC, 2005

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
BioInformatIQ	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

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2005 IDC. Reproduction without written permission is completely forbidden.