SPECIAL STUDY

Financing a Software Infrastructure for Highly Parallelised Codes — IDC FINAL Report for the DG Information Society of the European Commission

Earl C. Joseph, Ph.D.
Gabriella Cattaneo
Steve Conway

IDC OPINION

This study, "Financing a Software Infrastructure for Highly Parallelised Codes," was conducted under the SMART 2010/0052, contract number 30-CE-0394162/00-51, for the European Commission Information Society and Media Directorate-General, Emerging Technologies and Infrastructures.

IDC believes that Europe needs to adopt a parallel software development strategy to realize the full value of Europe's existing and planned investments in high-end supercomputers. Sufficient parallel software development expertise exists within Europe to seize global leadership in HPC in a number of key areas. Leadership, defined as sustained performance on real-world applications, will depend far more on software advances than hardware progress. Previous IDC research has established the linkage between HPC leadership, and scientific and industrial/commercial competitiveness.

The study participants are clearly ready to collaborate within a European Commission framework to advance key parallel applications used within their scientific and industrial/commercial organizations. The opportunity now exists for the European Commission to coordinate European hardware and software initiatives for the benefit of Europe and the Member States. In this Final Report (WP2) of the study, IDC recommends a specific strategy for accomplishing this. Here are the key recommendations:

- Use HPC Parallel Software Development to Help Close Europe's Innovation Gap
  - Europe has a limited window-of-opportunity to become a global innovation leader in targeted domains of scientific and engineering research
- Establish an e-Infrastructure for Advancing Parallel Software:
  - Create a new EC body to coordinate the holistic parallel software and holistic HPC strategy for Europe
  - Establish European centers-of-excellence for parallel software development
  - Create a European Web-based parallel software clearinghouse
  - Establish "Tiger Teams" to improve HPC access across Europe
  - Put into place the recommended parallel software funding

Filing Information: July 2011, IDC #SR03S
: Special Study

IDC EMEA, 389 Chiswick High Road, London, W4 4AE, U.K. Tel.+44.20.8987.7100 www.idc.com

Author(s)	IDC: Earl C. Joseph, PhD; Steve Conway; Gabriella Cattane
Deliverable	WP2 Final Report
Date of delivery	July 18, 2011
Version	1.2
Addressee officers	Dr. Timothy Heap Performance Computing Officer European Commission DG Information Society Unit F03, Office BU25 04/087 25 Avenue Beaulieu, B-1160 Bruxelles
Contract ref.	Contract Number 2009/S99-142914

Note: The opinions expressed in this report are those of the authors and do not necessarily reflect the views of the European Commission.

TABLE OF CONTENTS

P

TABLE OF CONTENTS — Continued

P

LIST OF TABLES

P

LIST OF FIGURES

P

1.0 EXECUTIVE SUMMARY

1.1 Summary of the Research Findings

Europe has a number of globally successful scientific and engineering software firms, a larger number of nationally and regionally successful software firms, and is strong in many important areas of parallel software development. These findings clearly emerged from the survey of EU HPC research centers and parallel software development organizations IDC conducted for this study.

The underlying problem is that modern HPC hardware with large numbers of CPU cores, each with decreasing levels of memory and memory bandwidth, is causing a mismatch with existing application software, driving a need to fundamentally redesign and rewrite HPC application software for greater parallelism, in order to perform well on future HPC systems.

But Europe (like the rest of the world) has been overly focused on funding parallel hardware to the detriment of parallel software, and on "big science" to the near-exclusion of industry. Where funding has been made available for parallel software development, the funding typically has been for only a year or two, compared with at least 5-10 years of funding needed to develop robust, production-quality software that can remain useful for 10-20 or even 30 years and across multiple generations of HPC hardware systems.

The vast majority (83%) of the most important parallel software applications in use at the surveyed European HPC sites were created in Europe. Intellectual property rights for a substantial majority of the sites' most important application codes (66%) were exclusively owned by European organizations. This finding runs counter to the frequent assumption within the worldwide HPC community that European HPC sites rely primarily on the U.S. for their most important software applications.

Only a few of the site's single most important application codes scale beyond 1,000-2,000 processor cores today, but 77% of the sites thought they could exploit tens of thousands of cores and 36% said they believed they could scale to hundreds of thousands of cores. The largest HPC systems today contain more than 200,000 cores, and million-core systems are expected to begin arriving before the end of this decade.

Half of the sites' first top codes were 10 or more years old, and 14% were more than 20 years old. Although most (71%) of the first top codes have undergone a technical update within the past year, and nearly all (97%) have been updated within the past five years, this does not imply that the codes have been substantially adapted or fundamentally rewritten to efficiently exploit today's and tomorrow's large-scale high performance computing (HPC) hardware systems.

Nearly all of the surveyed sites said they would be willing to invest in improving the parallel software application codes that are most important for their users. The key missing ingredients for improving the sites' important parallel codes are adequate internal expertise and financial investment. Almost all of the sites were willing to partner with outside organizations, including the European Commission, in these development/advancement efforts.

Respondents associated with industrial/commercial HPC use stressed that the development models for scientific and industrial/commercial parallel software need to be different. They said industrial/commercial representatives need to be involved in developing/advancing codes intended for their use, and that they would be willing to contribute expertise, money, and other resources to these initiatives, assuming that arrangements for intellectual property rights could be worked out satisfactorily.

Study participants responded that parallel software development for scientific and industrial/commercial use is important enough to establish a coordinating body within the European Commission. Working closely with PRACE, EESI, EGI, and others, they said, this body should coordinate R&D for the European HPC ecosystem.

Note that both the parallel applications codes and all of the supporting software tools, libraries, algorithms, etc. need to be improved in order to advance Europe's software environment., IDC estimates that there are at least 5,000 to 10,000 application codes and software tools in use across Europe that potentially need to be addressed. In this study a sample of 79 codes were carefully selected to provide a broad representative sample of the many codes in use across Europe today.

1.2 Benefits of an Expanded European Parallel Software Strategy

The main benefits of adopting a European strategy for the development/advancement of server-based parallel software for scientific and engineering research are as follows:

- **Parallel software advances will be more important for future HPC leadership than progress in designing parallel hardware systems.** For multiple decades the HPC community, in Europe and elsewhere, has focused primarily on highly parallel hardware systems, to the detriment of parallel software. As a result, today only about 1% of HPC application codes can exploit 10,000 or more processor cores. The largest HPC hardware systems today contain more than 200,000 cores, and million-core supercomputers will begin to arrive before the end of this decade.

- **A parallel software development strategy gives Europe the opportunity to create the world's best computational resources to enable scientific and industrial/commercial leadership.** Multiple worldwide studies, many conducted by IDC, have firmly established the close linkage between HPC use, and scientific and industrial competitiveness. As the Washington, DC-based Council on Competitiveness has said, in today's global economy "to out-compute is to out-compete." Without capable, highly parallel software, large supercomputers are less useful.

- **Europe has already made substantial investments in HPC hardware systems.** Through the PRACE program, Europe has invested large sums in high-end supercomputers to advance scientific, and to a much lesser extent engineering, capabilities. A complementary parallel software development strategy is needed to maximize the returns on the existing and planned hardware investments.

1.3 Key Recommendations

- Use HPC Parallel Software Development to Help Close Europe's Innovation Gap
  - Europe has a limited window-of-opportunity to become a global innovation leader in targeted domains of scientific and engineering research
- Establish an e-Infrastructure for Advancing Parallel Software:
  - Create a new EC body to coordinate the holistic parallel software and holistic HPC strategy for Europe
  - Establish European centers-of-excellence for parallel software development
  - Create a European Web-based parallel software clearinghouse
  - Establish "Tiger Teams" to improve HPC access across Europe
  - Put into place the recommended parallel software funding

2.0 INTRODUCTION AND STUDY BACKGROUND

2.1 Study Background and Research Methodology

This is the interim report of the study, "Financing a Software Infrastructure for Highly Parallelised Codes," conducted under the SMART 2010/0052, contract number 30-CE-0394162/00-51, for the European Commission Information Society and Media Directorate-General, Emerging Technologies and Infrastructures. The aim of the Contract is to perform a study that maps the parallel codes landscape in Europe and sets the framework for the creation of a European software infrastructure available for research and science. In addition, the study aims to analyze the strengths and weaknesses of the above-mentioned infrastructure for parallel scientific computing and point out the opportunities from the establishment of such infrastructure. The final report will identify the appropriate management structure, funding and governance model for the infrastructure.

Note that both the parallel applications codes and all of the supporting software tools, libraries, algorithms, etc. need to be improved in order to advance Europe's software environment., IDC estimates that there are at least 5,000 to 10,000 application codes and software tools in use across Europe that potentially need to be addressed. In this study a sample of 79 codes were carefully selected to provide a broad representative sample of the many codes in use across Europe today.

In 2010, IDC conducted a separate study for the same party, entitled "A Strategic Agenda for Supercomputing Leadership in Europe" (contract number 2009/S99-142914). The final report for this field research-based study noted, among other things, that:

> *"The increasing commoditization of highly parallel HPC hardware systems has greatly improved their price/performance while making them harder to use efficiently. This has shifted the burden of innovation much more onto the software side…IDC believes that in the next five years many HPC applications will need to be fundamentally rewritten so they can efficiently exploit HPC systems that will soon feature one million processor cores or more…European stakeholders from industry, research, and academia believe that Europe has a chance to jump back to the forefront of development for the next generation of HPC-based research, and for the applications and other software technologies required for the transition to petascale and exascale computing."*

The present study -- an entirely separate initiative -- set out to characterize Europe's strengths and weaknesses with respect to highly parallel software, and to propose an infrastructure for advancing Europe's capabilities in this software realm that will be increasingly important for Europe's scientific and economic competitiveness.

The specific goals of this study, as set out in the contract, are as follows:

- Map the parallel codes landscape in Europe
- Set the framework for creating a European software infrastructure for research and science.

- Analyze the strengths and weaknesses of the infrastructure
- Point out the opportunities arising from the infrastructure
- Identify the appropriate management structure and governance model for the infrastructure

2.1.1 Research Approach and Methodology

The study methodology combined qualitative-quantitative research to gather the necessary evidence, and analysis and assessment needed to produce recommendations. The methodology is structured in two main phases:

The Research Phase

The research phase consisted of:

- **Desk research** of relevant reports and other documents, including those found on the websites of the European Exascale Software Initiative (http://www.eesi-project.eu/pages/menu/homepage.php), the International Exascale Software Project (http://www.exascale.org/iesp/Main_Page), and other sources.
- **Field survey.** IDC completed 47 separate surveys, representing 79 critical parallel software application codes, for this report. In addition, IDC completed six in-depth interviews with HPC software leaders in Europe, along with dozens of separate discussions with software writers and users across the EU. IDC believes that the completed surveys fairly represent the common situations to be found among HPC users of parallel and highly parallel software. The variety of situations uncovered in this survey closely corresponded to those IDC found in prior worldwide studies with larger samplings. Before conducting the surveys, IDC reviewed the draft survey questionnaire with the European Commission. The questionnaire focused heavily on the top two software codes in use at the surveyed organizations. A copy of the questionnaire appears in the appendix of this report. Sites were selected carefully to provide a broad representation of the different software across Europe, balanced between industry/academic/government sites; balanced between the many different domains; and balanced across the different countries in Europe. Table 1 shows the sites surveyed using the standard questionnaire.
- **Telephone interviews with key stakeholders.** IDC believed it was also important to probe key stakeholders on the survey topic in a more open-ended, qualitative way than would be practical with a written questionnaire alone. The goal of this probe, as that word suggests, was not to interview each of the many individuals with a stake in software development within Europe, but to provide a reality check for the written survey responses by ensuring that no important part of the "bigger picture" would be missed in IDC's analysis and recommendations. IDC interviewed software experts with leading roles in the European Exascale Software Initiative (EESI) and, for an outside perspective, the International Exascale Software Project (IESP) as well. In addition, IDC interviewed European-headquartered developers/vendors of highly parallel software and prominent European industrial users of parallel software, whose perspectives on appropriate models for software development and governance often differered markedly from the viewpoints of scientific/academic software users. The telephone interviews focused on these questions (note: IDC deliberately used the term "parallel software" rather than "highly parallel software," because many

industrial applications codes are not highly parallel and our goal was to include industrial software):

- Where do you think there are strengths within Europe in parallel software?
- If you had money to invest in parallel software within Europe, where would you invest?
- If you had money to invest in parallel software within Europe, what selection process and business model would you use for these investments?
- What governance structure would you use to manage European investments in parallel software?

2.2 Sites Interviewed

IDC completed 47 separate surveys, representing 79 critical parallel software application codes (shown in Table 1), for this report. In addition IDC completed 6 in-depth interviews with parallel software leaders in Europe, along with dozens of separate discussions with software writers and users across the EU.

TABLE 1

Organizations Surveyed With The Fixed Questionnaire

(47 Surveys Representing 79 HPC Application Codes)

Organization	Organization
Audi AG	Leibniz Supercomputing Centre (4)
AWE PLC	Met Office
CEA – Fusion and Plasma	Numerical Algorithms Group
CEA -- Materials Sciences	Prometeus GmbH
CEA DEN - Nuclear Energy	Research Centre Juelich
CILEA	SARA Computing and Networking Services
CINECA	Technische Universität München
CSC - IT Center for Science (2)	The University of Edinburgh
CSCS – Swiss National Supercomputing Centre	UCD
Daresbury Laboratory	University College Dublin
ECMWF	University of Amsterdam
EDF R&D (2)	University of Coimbra
Erlangen Regional Computing Center / Univ. of Erlangen	University of Edinburgh

TABLE 1

Organizations Surveyed With The Fixed Questionnaire

(47 Surveys Representing 79 HPC Application Codes)

Organization	Organization
Forschungszentrum Jülich GmbH, IEK-8: Troposphäre	University of Genova
GENCI	University of Groningen
HLRS/University of Stuttgart (2)	University of Manchester
Imperial College London – Dept of Materials and Physics	University Of Sofia (2)
INRIA	University of Sussex
IPSL/CNRS	University of Tennessee
KIT	VSB-Technical University of Ostrava

Source: IDC, 2011

The Analytical and Assessment Phase

This phase evaluated the novel field research conducted for this study in light of IDC's prior HPC research and general understanding of the European and worldwide HPC markets, in order to propose recommendations in the final report for the advancement of parallel software development in Europe.

2.1.2 Report Layout

This interim report is divided into these primary sections:

- Executive Summary
- Introduction and Study Approach
- Field Research Results
- IDC Recommendations for The EU
  - Situation Background
  - Key Actions
  - Funding
  - Some out-of-the-box ideas
- Conclusions
- Appendices

2.3 Market and Technology Trends and Issues

2.3.1 HPC Trends

- The Worldwide Economic Recovery Is Restoring Healthy HPC Growth. In 2009, the low point of the global economic downturn, worldwide HPC server revenue dropped 13% year over year, to €6.0 billion ($8.6 billion) -- although the market for supercomputers priced at €2.1 million ($3 million) and up increased by 65%. In contrast, full-year 2010 HPC server revenue grew 10% to reach €6.6 billion ($9.5 billion). This was still below the 2007 high point of €7.5 billion ($10.1 billion). IDC forecasts that HPC server revenue will increase at a healthy rate (CAGR) of about 7% annually during the next five years to reach €8.5 billion ($12.3 billion) in 2014.
- Major challenges for datacenters are as follows:
  - Power, cooling, real estate, system management
  - Storage and data management continue to grow in importance
- Software hurdles will rise to the top for most users
- SSDs will gain momentum and could redefine storage
- GPUs are seeing real tractions in certain verticals
- The worldwide "Petascale/Exascale Race" is proceeding at full speed

2.4 The General HPC Parallel Software Landscape

Observations about Highly Parallel Software

IDC offers the following observations about parallel software:

- **The increasing commoditization of highly parallel HPC hardware systems has greatly improved their price/performance while making them harder to use efficiently**. This has shifted the burden of innovation much more onto the software side. IDC believes that in the next five years, many HPC applications will need to be fundamentally rewritten so they can more efficiently exploit HPC systems the largest of which will soon feature one million processor cores or more. As IDC's 2010 study for the European Commission confirmed, European stakeholders from industry, research, and academia believe that Europe has a chance to jump back to the forefront of development for the next generation of HPC-based research, and for the applications and other software technologies required for the transition to petascale and exascale computing.
- **Few scientific and engineering software codes are highly parallel today.** Hence, there is a major opportunity for further parallelizing scientific and engineering codes so that they can better exploit large HPC systems and the substantial investments being made in these systems. It is important to put the term "highly parallelized" in perspective. While a handful of academic and government codes have already been run on more than 100,000 processor

cores, most scientific and engineering software codes used with HPC systems are far less parallel than this today. A recent IDC worldwide study found that 52% percent of the HPC applications reported by the 188 surveyed sites were running on only one node (averaging 3.4 cores), or a fraction of one node. Only about 12 percent of the codes were running on more than 1,000 cores. And just one percent of the codes were able to exploit 10,000 or more cores. (The results for European HPC sites did not differ markedly from results for other global regions.) This is not much fodder for today's largest HPC systems that feature more than 200,000 cores, not to mention exascale systems (10^{18} operations/second) with more than one million cores that are expected before the year 2020. There is a major opportunity for further parallelizing scientific and engineering codes so that they can better exploit large HPC systems and the substantial investments being made in these HPC systems.

- **It is important to include industrial/commercial software codes, which could also greatly benefit from further parallelization and other advancements.** Many widely used industrial engineering codes are only modestly parallel today, with the ability to exploit only eight or fewer processor cores. These codes often have high economic value, especially for the manufacturing sector that in 2010 represented 27% of European GDP. Like their scientific counterparts, these codes could benefit greatly from a financed European software infrastructure, assuming it is established with their sometimes different requirements also in mind. To boost their impact on European industrial/commercial innovation and economic competitiveness, the scalability of key industrial application codes need not be raised to the petaflop or exaflop levels; more modest improvements in scalability could have a very salient effect.

- **The parallelization effort is different for each software code.** As is well known, some scientific and engineering codes are inherently more difficult to parallelize than others. Certain applications domains such as quantum chromodynamics (QCD), the particle physics codes addressing the strong interaction, are "embarrassingly parallel." They are likely to remain parallelization (scalability) champions even on future exascale computers with more than one million processor cores. (Getting these codes to efficiently exploit ultrahigh-end HPC systems is nevertheless a difficult, impressive achievement.) Most HPC codes were originally written to run on one single-threaded processor. Many of these codes have never been fundamentally re-written. Instead, these "dusty deck" codes have been laboriously modified over time to enable modest scalability. In other cases, the lack of more scalable algorithms curbs sustained, parallel performance. This is a solvable problem, although there may not be enough people on Planet Earth with the right combination of brainpower and experience to solve it for every deserving application.

- **More and more often, parallel scalability is constrained by the limitations of the underlying, known science.** Where the science doesn't support large single runs of the problem, smaller iterative runs often allow users to home in on productive solutions. Long-established examples include stochastic modeling in the financial services sector and parametric modeling in the design engineering realm. Newer examples are the increasingly complex ensemble models in the weather/climate domain, and a whole host of genomics and proteomics applications. Biology, in particular systems biology, is a scientific domain of such overwhelming complexity that the known science often fails to support single, deterministic runs of large problems.

- **Each software code needs to be assessed based on its own merits and status.** Software codes are not necessarily less valuable simply because they are less scalable than the highest-performing HPC codes. As noted earlier, some codes are inherently more difficult to parallelize than others, and modestly parallel design engineering codes may be indispensable for manufacturing and other economically important industrial sectors. No matter how highly parallelized they are today, many codes could benefit from investment in further parallelization, resolution, robustness, and other attributes. The small and moderate scaling codes needs the most help today, in addition many of the largest scaling codes also need assistance and improvement.

- **Application codes are not the only important software solutions needing further parallelization.** Other software categories that are crucial for computational science and engineering work include tools, programming languages, parallel file systems, mathematical libraries, commands, workload schedulers and balancers, grid and cloud software, and other elements of the HPC management software stack. For highly parallel hardware systems to operate efficiently, all software running on the systems needs to be sufficiently parallel, robust, and resilient. Ideally, application codes and the rest of the software stack should be developed together.

- **There is a major opportunity for further parallelizing software codes.** Few scientific and engineering software codes are highly parallel today. A recent IDC worldwide study of HPC data centers revealed that 52% of their software codes are running on one node or less of an HPC system. Only 12% of the codes could exploit 1,000 or more cores, and only 1% could exploit 10,000 or more cores. Today's largest HPC systems contains more than 200,000 cores, and supercomputers with more than one million cores will begin to arrive before the end of this decade. (See appendix*: IDC HPC End-User Study of the Industry/Application Trends in Technical Computing*, December 2010.)

3.0 THE EUROPEAN PARALLEL SOFTWARE LANDSCAPE: FIELD RESEARCH RESULTS

The field research IDC conducted for this study was fundamentally consistent with prior IDC research findings, but also shed considerable new light on the topic of European parallel software for science and research (including engineering).

For this study, IDC interviewed representatives of leading European and international HPC software initiatives, along with a representative sampling of HPC data-center managers and end-users of HPC scientific and engineering software. In addition, IDC reviewed software-related documents and other information available on the websites of the European Exascale Software Initiative (EESI), the International Exascale Software Project (IESP), and other organizations involved in assessing and advancing parallel software capabilities in Europe, as well as materials available from European parallel software vendors.

IDC completed 47 separate surveys, representing 79 critical parallel software application codes, for this report. In addition, IDC completed six in-depth interviews with HPC software leaders in Europe, along with dozens of separate discussions with software writers and users across the EU. IDC believes that the completed surveys fairly represent the common situations to be found among HPC users of parallel and highly parallel software. The variety of situations uncovered in this survey closely corresponded to those IDC found in prior worldwide studies with larger samplings.

3.1 Analysis of the Parallel Codes Landscape in Europe

The field research produced a consistent picture of Europe's strengths and weaknesses in parallel and highly parallel software.

European Strengths

- **Europe has some notably successful scientific and engineering software firms.** "Successful," in this context, primarily means that the Europe-headquartered firms' parallel software products and services are well recognized and are used not only within Europe, but across the world. Shining examples of software firms in this category include Allinea (UK), Dassault Systèmes (France), and Numerical Algorithms Group (UK). A number of other globally successful parallel software companies began in Europe, still maintain a strong presence in Europe, but moved their headquarters to the United States for business reasons.

  Here is a longer, though not exhaustive, list of European parallel software companies:

  - Allinea Software
  - Bright Computing
  - Bull
  - Cadfem

- Cluster Resources
- ClusterVision
- Dassault Systèmes
- ESI
- Eurotech
- GNS Systems
- Go Virtual
- Gridcore
- Integrated Engineering Software
- Kerlabs
- LSI
- MEGWARE Computer GmbH
- Mercury Visualization Science Group
- NAG
- Nema Labs
- NICE
- ParTec Cluster Competence Center
- PDE Solutions Inc.
- science + computing ag
- Scilab
- Ter@tec
- T-Platforms
- Transtec
- Ylichron

- **Europe is strong in important areas of parallel software development.** Field research conducted for this study produced a consensus on the following areas of European strength in parallel software:
  - **Scientific applications.** A large number of parallel scientific codes have been developed within Europe's universities and research institutes.
  - **Precedent.** The European Commission has successfully funded parallel software development before. A number of the HPC experts we interviewed

said that they personally had been involved in EC-supported software parallelization efforts in the 1990s.

- **Algorithms/solvers.** Europe is strong in developing parallel algorithms and solvers, although most of these have not been widely disseminated into application codes.

- **Mesh generators.** There are good skills within Europe for developing mesh generators. Many of these skills reside within small software vendors.

- **Numerical libraries.** This is an area in which Europe arguably leads the world (NAG and others).

- **Performance profiling and performance modeling.** Analyzing the behavior of application codes and other software on parallel hardware systems is crucial for improving software design and performance. Within Europe, capable performance profiling software has been developed (e.g., Vampir) and world-class performance modeling expertise exists (Universitat Autònoma de Barcelona, Regionales Rechenzentrum Erlangen, and others).

- **Open source software experience.** Within Europe there is also world-class experience with open source software, as evidenced by collaborations such as the Open File Systems Initiative.

European Weaknesses

- **Europe (and much of the rest of the world) has been overly focused on funding parallel hardware, to the detriment of parallel software.** Europe has been part of a long-term global trend that has focused HPC government funding support far more heavily on the acquisition of highly parallel hardware systems than on the development of parallel software capable of efficiently exploiting the hardware systems. Notable examples of this trend include the PRACE program in Europe, along with the U.S. Department of Energy's ASC and INCITE programs and the U.S. National Science Foundation's cyberinfrastructure program. Non-European nations have not made strong commitments to redressing this imbalance, and this presents Europe with a major opportunity to pursue global leadership in chosen areas of parallel software development. Because of the large and growing gap between hardware and software capabilities, IDC believes that software advances will be more important for future HPC leadership than hardware progress. This study points out useful pathways for parallel software development, and the work being done by the European Exascale Software Initiative (EESI) promises to provide a more detailed roadmap for European parallel software advances. But without appropriate funding and support, research studies and software roadmaps are likely to remain mere academic exercises.

- **European HPC programs have generally focused on "big science," to the near-exclusion of industry.** IDC research has shown that the PRACE program generally gets high marks for delivering on its core mission to provide Europe's scientific and research communities with access to world-class supercomputers. PRACE has held successful industry workshops for several years, but to date industry has had little access to PRACE HPC systems. The industrial workshops imply that industry will receive access to these systems in the future. And as an early realization of this potential, a PRACE Tier-0 petascale supercomputer is scheduled to be installed later this year at HRLS/University of Stuttgart, where it will be used to support engineering sciences and other work. Industrial access to large-scale supercomputers such as this will be useful, because even large industrial firms typically cannot justify purchasing the most powerful HPC systems on their own -- yet these firms can have computational challenges that are every bit as daunting as some large-scale scientific challenges. If European industry does not have meaningful access to leadership-class supercomputers, it will have no way to address the very problems that are most crucial for driving industrial innovation and maintaining global competitiveness. Access to highly parallel HPC systems also provides an opportunity for further parallelization of industrial software codes. Industrial software users in the present study stressed the value of this access -- and collaboration with domain scientists at academic HPC centers.

- **Many engineering software codes are outdated.** Respondents from both the European scientific and engineering communities noted that many important engineering codes have not been substantially parallelized and otherwise adapted to take greater advantage of contemporary parallel HPC systems. Software vendors often are too small to fund these parallelization efforts on their own or, even if larger, cannot make a business case for parallelization initiatives -- which may require substantial rewriting or even rethinking of the codes. Some respondents claimed that the European Commission had funded software development efforts in the past, and that they personally had been involved in these efforts.

- **Industrial software developers and users said there is a bias within Europe toward academic experts and away from experts in industry.** European industrial developers and users of parallel software complained that efforts to improve industrial software are too often "automatically directed" toward experts in academic computing centers. The industrial contingent stressed that important software expertise also exists in commercial firms and needs to be brought into collaborations with academics, so that the resultant software better meets the needs of industry. Not surprisingly, the industrial developers and users also wished for less "red tape" and less-restrictive licensing agreements for software developed with government funding support.

- **Funding is too short-term and the infrastructure is inadequate for developing robust, production-quality software.** There is a saying in the HPC community that "hardware is easy, software is difficult." Although this description is simplistic, it is true that an HPC hardware generation typically lasts three to four years, while a software generation may need to last 20 to 30 years and operate across many generations of changing hardware systems. The survey respondents noted that funding for parallel software development initiatives, in Europe and elsewhere, often lasts for only one or two years, and that this is not nearly enough time to develop software that is robust enough to operate reliably for multiple decades in real-world, production computing environments. Respondents said that at least 5-10 years of funding support is needed to create robust, production-quality parallel software. A major issue with the physical infrastructure is the difficulty some parallel software developers have in getting access to sufficiently large-scale HPC systems for development and testing purposes, along with expertise for helping to parallelize the codes so that they can exploit highly parallel HPC hardware systems.

3.2 Analysis of the Main Codes and Applications

Origins of the Most Important and Second Most Important Application Codes

Figure 1 and Table 2 show the origins for each of the surveyed organizations' most important scientific (or engineering) research code, and for the second-most-important scientific (or engineering) research code.

Combining the totals for both the first most important and second most important codes used at the sites, the largest number (32%) originated in academia. Almost as many (28%) stemmed from the surveyed organizations themselves. Note that some of these were academic organizations that also contributed to the aforementioned academic total. Seventeen of the most important codes (21%) had open source origins, nine (11%) came from industry, and six (7%) arose out of government.

Of the 40 first most important codes named by the sites, 35% originated in academia, 28% were home-grown, 18% had open source origins, 15% were commercially developed, and the remaining 5% came into being within government environments.

Of the 41 second most important codes, 29% each had academic and home-grown origins, 27% were open source, 10% originated in government, and 7% had commercial roots.

FIGURE 1

Origin of Codes Used at the Sites Surveyed

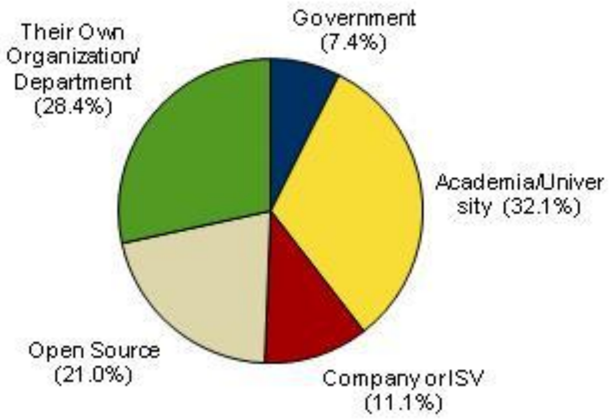

Source: IDC, 2011

TABLE 2

Origin Of Codes Used at The Sites Surveyed

	#1 Code	#2 Code	Combined	Percent of Sample
Government	2	4	6	7.4%
Academia/ University	14	12	26	32.1%
Company or ISV	6	3	9	11.1%
Open Source	7	10	17	21.0%
Their Own Organization/ Department	11	12	23	28.4%

Source: IDC, 2011

Top Application Codes Used At the Surveyed Sites

Table 3 lists the sites' most important and second most important applications codes by name. This collection of codes covers a fairly wide spectrum of scientific domains, along with some popular engineering applications.

TABLE 3

Top Application Codes Used At The Sites Surveyed

Code	Code	Code
ABINIT	GROMACS (6)	Overture
AIMPRO	GYSELA (2)	Pamcrash (ESI)
ANSYS (2)	IFS (2)	PARADIS
BQCD	INCA	ParaFEM
CASTEP	INHOUSE A	ParaGauss
CFD (STAR CD)	INHOUSE B	POEM
Code_Aster, Salome-Meca	IPSLCM	Python
Code_Saturne	JOREK	Quantum Espresso
COSMO / ECHAM	LAMMPS	R SPRINT
CP2K / CMPD (3)	Massively Parallel Quantum Computer Simulator	S12MPI

TABLE 3

Top Application Codes Used At The Sites Surveyed

Code	Code	Code
Deposit	Matlab	Sire
DLPOLY	Molecular Dynamics (AMBER)	STAMP
ECHAM - HAMMOZ	NAMD 93)	Telemac
Elmer	NASTRAN	TRIO-U
Eurad - IM	NEMO	TRIPOLI
FLOW-3D	Neptune_CFD	Unified Model
Fluent	Numpy/Scipy	Vampir
GaitSym	Octopus	VASP (3)
GENE	ONETEP	WRF
GPAW	Open File Systems Initiative	YALES2
GPREMS	OpenFOAM (OpenCFD) (2)	.

Source: IDC, 2011

Countries of Origins of the Application Codes

A large majority of the surveyed sites' most important applications were developed within Europe (Table 4A).

Respondents identified the national origins for 58 of the 81 most important codes (72%). Five-sixths (83%) of the identified applications were created within Europe, and only 17% originated in the United States. To repeat, the vast majority of the applications codes deemed most important by the surveyed European HPC sites were of European origin.

TABLE 4A

Code Creator's Country (Location)

Location	Location
Austria	Portugal, Spain, Germany and others
EU	Switzerland
Finland	UCL (Louvain la Neuve, Belgium)
France (15)	UK (13)

TABLE 4A

Code Creator's Country (Location)

Location	Location
France, UK	USA (10)
Germany (15)	Vienna, Austria (2)
Italy	Zurich + Italy
Netherlands (3)	.

Source: IDC, 2011

Primary Industries/Sectors Where the Code Is Used

Table 4B displays the sites' responses when asked about the primary sectors/industries in which the sites' most important ("first top code") and second most important ("second top code") applications. The responses make it clear that although scientific application codes predominate, applications used by key industries are fairly well represented within the sample population

FIGURE 2

Primary User Type for the Top codes

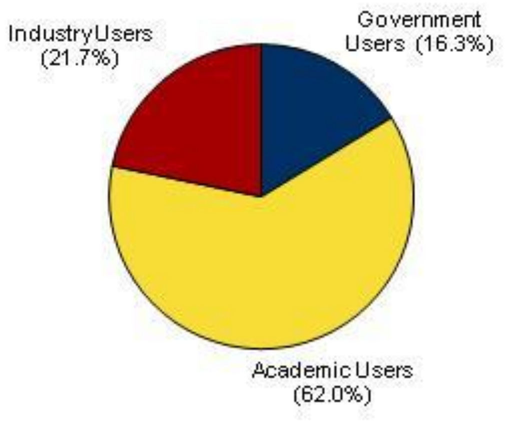

Source: IDC, 2011

TABLE 4B

Primary Industries/Sectors That The Top Two Codes Are Used In

Industry	Industry
Academic Research (4)	Hydraulics & environmental (2)
Aerospace, automotive (4)	Life Sciences
Biochemistry, Molecular Biology (4)	Magnetic Fusion Research (2)
CAE, Fluid-dynamics	Material science (3)
CFD, Aerospace (2)	Meteorology, Oceanography, Climatology (2)
Chemical , Semiconductor	Molecular dynamics
Chemistry (6)	Molecular simulations of bio and non-bio materials
Civil Engineering, Materials Science, Life Sciences	Multidisciplinary, Mathematics, Engineering
Climate modeling (3)	Pharmaceutical Industry (4)
Combustion , Aeronautics , Automotive	Plasma Physics research
Condensed Matter Physics, Material Modeling, Nanoscale	Power Generation (mainly nuclear pressurized water reactors)
Defense (20	Research, Atmosphere , Climate (2)
Electronic Structure Simulations of disordered or crystalline materials	Science, Signal processing (2)
Energy (4)	Shock physics, material properties
Engineering (3)	Solid state physics, Material properties
Fusion	Statistics & Biosciences
General Scientific Research (3)	Weather forecasts (2)

Source: IDC, 2011

Primary Country or Countries Where the Codes Are Used

Half (29) of the 58 "first top codes" and "second top codes" for which survey respondents provided this information (see Table 5A) were being used only within Europe, while the other half (29 codes) were employed both within Europe and elsewhere in the world.

Twenty-six of the 58 codes (45%) were also used in the U.S. Seventeen of the codes (29%) were being used in only one European country and nowhere else in the world.

TABLE 5A

Primary Countries That The Codes Are Used In (can be multiple per code)

Country	Country
All of Europe (12)	Netherlands, US, Germany
EU, USA (4)	UK (4)
France (7)	UK, Australia, Norway, S Korea
France, Great Britain	UK, US, Germany
France, UK, Deutch	UK, US, Japan
France, Belgium, Canada	UK, USA, China (2)
France, Germany	United States
Germany (6)	USA, Europe, Asia (3)
Germany, Finland	Worldwide (13)
Germany, UK, Japan, Russia	.

Source: IDC, 2011

Primary User Type for the Top Codes

As Table 5B indicates, academic researchers were the primary users for 62% of the top application codes (first and second most important codes combined), followed by industrial/commercial users (22%) and government users (16%). All 47 surveyed sites responded to this question.

For the first most important codes, academics made up 67% of the responses, compared with 20% for industrial/commercial users and 14% for government users.

For the second most important codes, academics represented 56% of the responses, compared with 24% for industrial/commercial users and 20% for government users.

TABLE 5B

Primary User Type For The Top Codes

	Top Code	#2 Code	Combined	Percent of Sample
Government Users	7	8	15	16.3%
Academic Users	34	23	57	62.0%
Industry Users	10	10	20	21.7%

N = 47, Note that multiple responses were allowed

Source: IDC, 2011

Average Use of Sockets and Cores for Single Jobs

The average number of processor sockets (physical processor parts) and processor cores the "first top codes" were exploiting for single jobs was 19,708 and 23,664, respectively, according to the surveyed sites (see Table 6A). For their second most important codes, the averages were lower but still impressive: 11,838 processor sockets and 13,631 cores. All 47 surveyed sites responded to this question.

TABLE 6A

Singe Job Size - MAXIMUM Number

(Note: Unlimited was set at 100,000)

	a. 1st code - SOCKETS	a. 1st code - CORES	b. 2nd code - SOCKETS	b. 2nd code - CORES
Average	19,708	23,664	11,838	13,631

N = 47

Source: IDC, 2011

Maximum Use of Sockets and Cores for Single Jobs

The median values for maximum scalability on single jobs were in the 512 to 2,000-socket range, with some values as high as 100,000 processor sockets . Where cores were concerned, the median was in the 512 to 1,000 range. Because quad-core and larger processors are most common in contemporary HPC systems, the fact that the median core scalability is lower than the median socket scalability indicates that the applications often cannot efficiently exploit all of the cores in the sockets. IDC knows from our other research that this is often the case, and in some cases users may exploit the memory and bandwidth associated with a core while not using the core's processing capability. This stems from the current imbalance between theoretical processing power (peak flops) and associated memory and bandwidth. The pattern of the responses for the "second top codes" was similar.

TABLE 6B

Singe Job Size - MAXIMUM Number

(Note: Unlimited was set at 100,000)

a. 1st code - SOCKETS	a. 1st code - CORES	b. 2nd code - SOCKETS	b. 2nd code - CORES
4	4	8	4
8	32	16	64
8	48	32	128
64	50	32	200
70	128	32	256
100	256	64	256
256	280	256	256
500	512	512	256
512	512	1,000	512
512	512	1,024	512
512	1,000	8,000	1,000
1,000	1,024	8,192	1,000
2,000	1,024	10,000	1,024
2,000	2,000	10,240	2,048
6,250	2,048	50,000	4,096
8,000	3,000	100,000	8,000
8,000	4,000	.	8,192
8,192	4,096	.	10,000
10,000	5,000	.	10,000
50,000	5,000	.	10,000
75,000	6,000	.	10,000
100,000	8,192	.	32,000
100,000	8,192	.	40,960
100,000	10,000	.	100,000
.	10,000	.	100,000
.	10,000	.	.

TABLE 6B

Singe Job Size - MAXIMUM Number

(Note: Unlimited was set at 100,000)

a. 1st code - SOCKETS	a. 1st code - CORES	b. 2nd code - SOCKETS	b. 2nd code - CORES
.	12,000	.	.
.	32,000	.	.
.	50,000	.	.
.	75,000	.	.
.	100,000	.	.
.	100,000	.	.
.	100,000	.	.
.	100,000	.	.
.	100,000	.	.
.	100,000	.	.

Source: IDC, 2011

Typical Use of Sockets and Cores for Single Jobs

Whereas Tables 6A and 6B described the maximum use of sockets and cores for single jobs, Tables 7A and 7B illustrate the typical use. As might be expected, the typical numbers were substantially less than the average numbers.

The average maximum number of processor sockets and processor cores the "first top codes" were exploiting for single jobs was 19,708 and 23,664, respectively (Table 6A). The typical numbers averaged 9,999 sockets and 7,990 cores (Table 7A).

For the HPC sites' second most important codes, the average maximum number of processor sockets and processor for single jobs was 11,838 sockets and 13,631 cores (Table 6A). The typical numbers averaged 2,410 sockets and 809 cores (Table 7A).

Typical job sizes for first most important codes were in the range of 16-32 sockets and 32-64 cores (Table 7B), compared with median values for maximum scalability on single jobs were in the 512 to 2,000-socket range and the 512-1,000 core range (Table 6B.

TABLE 7A

Singe Job Size - TYPICAL Number

	a. 1st code - SOCKETS	a. 1st code - CORES	b. 2nd code - SOCKETS	b. 2nd code - CORES
Average	9,999	7,990	2,410	809

N = 47

Source: IDC, 2011

TABLE 7B

Singe Job Size - TYPICAL Number

a. 1st code - SOCKETS	a. 1st code - CORES	b. 2nd code - SOCKETS	b. 2nd code - CORES
2	4	2	1
4	16	4	16
4	16	8	16
6	24	8	32
8	32	16	48
10	32	16	60
16	32	16	64
16	32	32	64
16	32	32	64
16	64	32	100
32	64	64	128
32	64	64	128
32	64	100	128
48	80	128	128
64	96	128	128
100	100	1,024	128
128	128	1,536	256
170	128	2,000	256
256	128	40,578	256
256	128	.	384

TABLE 7B

Singe Job Size - TYPICAL Number

a. 1st code - SOCKETS	a. 1st code - CORES	b. 2nd code - SOCKETS	b. 2nd code - CORES
256	500	.	500
512	500	.	512
666	512	.	512
1,360	512	.	1,000
1,536	1,000	.	1,024
4,096	1,000	.	1,536
8,192	1,000	.	3,000
262,144	1,024	.	5,000
.	1,024	.	8,000
.	1,024	.	.
.	1,024	.	.
.	1,536	.	.
.	4,096	.	.
.	4,096	.	.
.	5,000	.	.
.	8,000	.	.
.	8,160	.	.
.	8,192	.	.
.	262,144	.	.

Source: IDC, 2011

Potential for Raising Scalability to Hundreds of Sockets or Thousands of Cores

As Table 8 shows, the vast majority of the surveyed sites were confident in their ability to increase the scalability of the codes in order to exploit hundreds of sockets or thousands of cores.

For their "first top codes," 95% of the sites were confident they could exploit hundreds of sockets or thousands of cores. For the "second top codes," 93% of the sites responding to this question answered in the affirmative.

TABLE 8

Expand to 100s sockets or 1000s cores?

a. 1st code - YES	36
a. 1st code - NO	2
b. 2nd code - YES	25
b. 2nd code - NO	3

Source: IDC, 2011

Potential for Raising Scalability to Thousands of Sockets or Tens of Thousands of Cores

Table 9 shows that a smaller but still clear majority of the surveyed sites were confident in their ability to increase the scalability of the codes in order to exploit thousands of sockets or tens of thousands of cores.

For their "first top codes," 77% of the sites were confident they could exploit thousands of sockets or tens of thousands of cores. For the "second top codes," 59% of the sites responding to this question answered in the affirmative.

TABLE 9

Expand to 1000s Sockets or 10,000s Cores?

a. 1st code - YES	28
a. 1st code - NO	9
b. 2nd code - YES	16
b. 2nd code - NO	11

Source: IDC, 2011

Potential for Raising Scalability to Tens of Thousands of Sockets or Hundreds of Thousands of Cores

Table 10 shows that only a minority of the surveyed sites were confident in their ability to increase the scalability of the codes in order to exploit tens of thousands of sockets or hundreds of thousands of cores.

For their "first top codes," 36% of the sites were confident they could exploit tens of thousands of sockets or hundreds of thousands of cores. For the "second top codes," 25% of the sites responding to this question answered in the affirmative.

TABLE 10

Expand to 10,000s Sockets or 100,000s Cores?

a. 1st code - YES	9
a. 1st code - NO	16
b. 2nd code - YES	4
b. 2nd code - NO	12

Source: IDC, 2011

Number of Employees in the Support Organization, First Top Code

As Table 11 illustrates, the organizations supporting/maintaining the sites' most important application code tended to be relatively small, that is, having 25 or fewer employees (48%) or considerably larger, that is, having more than 100 employees (38%). Only one of the organizations had 50 to 100 employees.

TABLE 11

Number of Employees in Support Organization - FIRST CODE

a. <10	6
b. 10-25	4
c. 25-50	2
d. 50-100	1
e. >100	8

Source: IDC, 2011

Number of Employees in the Support Organization, Second Top Code

Table 12 essentially mirrors the bifurcated size pattern shown in Table 11. The organizations supporting/maintaining the sites' second most important application code also tended to be relatively small, that is, having 25 or fewer employees (33%) or considerably larger, that is, having more than 100 employees (58%). Only one of the organizations had 25 to 100 employees.

TABLE 12

Number of Employees in Support Organization - SECOND CODE

a. <10	2
b. 10-25	2
c. 25-50	1
d. 50-100	0
e. >100	7

Source: IDC, 2011

Number of Licenses/Clients Worldwide, First Top Code

The number of licenses or clients worldwide for the sites' first top codes varied greatly, from fewer than 10 (29% of responding sites) to more than 1,000 (also 29%). See Table 13. Respondents provided this information for 17 of the 40 (43%) of the first most important codes.

TABLE 13

Licenses or Clients Worldwide - FIRST CODE

a. <10	5
b. 10-25	3
c. 25-50	0
d. 50-100	2
e. >100	2
f. >1000	5

Source: IDC, 2011

Number of Licenses/Clients Worldwide, Second Top Code

The number of licenses or clients worldwide for the sites' second top codes also varied greatly, from fewer than 10 (33% of responding sites) to more than 1,000 (22%). See Table 14. Respondents provided this information for 9 of the 41 (22%) of the first most important codes, enough to show a pattern but not a detailed one.

TABLE 14

Licenses or Clients Worldwide - 2nd CODE

a. <10	3
b. 10-25	1
c. 25-50	1
d. 50-100	1
e. >100	1
f. >1000	2

Source: IDC, 2011

Age of the First Top Code

As Table 15 shows, half of the sites' first top codes (50%) were 10 or more years old, and 71% were at least five years old. One in seven codes (14%) was more than 20 years old. Respondents provided this information for 70% of the first most important codes,

TABLE 15

Age of FIRST CODE (PACKAGE A)

a. <1 year	2
b. 2-5 yrs	6
c. 5-10 yrs	6
d. 10-20 yrs	10
e. >20 yrs	4

Source: IDC, 2011

Age of the Second Top Code

Of the second top codes (Table 16), 58% were 10 or more years old, and 79% were at least five years old. About 11% were more than 20 years old. Respondents provided this information for 46% of the second most important codes,

TABLE 16

Age of Second CODE (PACKAGE B)

a. <1 year	2
b. 2-5 yrs	2
c. 5-10 yrs	4
d. 10-20 yrs	9
e. >20 yrs	2

Source: IDC, 2011

Last Major Technical Update, First Top Code

The good news is that 71% of the reported first top codes had undergone a technical update within the past year, and nearly all of these codes (97%) had been updated within the past five years. The term "major technical update" does not imply that the code has been substantially adapted or fundamentally rewritten, only that it has been modified for greater efficiency and/or compatibility with contemporary HPC systems. As noted elsewhere in this report, a separate recent IDC worldwide study revealed that only about 1% of HPC application codes scale to 10,000 or more cores, and most HPC codes are no more than eight-way parallel.

TABLE 17

Last Major Technology Update - FIRST CODE

a. <1 year	22
b. 2-5 yrs	8
c. 5-10 yrs	1
d. 10-20 yrs	0
e. >20 yrs	0

Source: IDC, 2011

Last Major Technical Update, Second Top Code

Similarly, 73% of the second top codes had undergone a major technical update within the past year, and nearly all of these codes (97%) had been updated within the past five years (Table 18). FROG

TABLE 18

Last Major Technology Update - 2nd CODE

a. <1 year	22
b. 2-5 yrs	7
c. 5-10 yrs	0
d. 10-20 yrs	1
e. >20 yrs	0

Source: IDC, 2011

Intellectual Property Ownership of the Application Codes

Earlier (Table 3), we saw that a large majority of the surveyed sites' most important applications were developed within Europe. Respondents identified the national origins for 58 of the 81 most important codes (72%). Five-sixths (83%) of the identified applications were created within Europe, and only 17% originated in the United States.

So, it should not be surprising that the intellectual property (IP) rights for most of the key application codes were also held by European organizations. As Table 19 shows, of the 62 codes whose IP ownership was identified, 40 (66%) were solely owned by European organizations, 2 (3%) had multinational ownership that included a European co-owner, 15 (24%) were open source, and 6 (10%) had non-European IP owners.

TABLE 19

Intellectual Property Owner Of The Top Codes

Owner	Owner
Cambridge Enterprise Ltd (2)	JSC(Juelich)
CEA (5)	Karlsruhe Institute of Technology (KIT)
CEA-EDF-AREVA	LLNL
CNRS+CEA+Universities	Mathworks
CORIA (U.S.)	Met Office
CSC	Myself (5)
Daresbury Laboratory	NAG (2)
ECMWF (2)	National Center for Atmospheric Research (NCAR)

TABLE 19

Intellectual Property Owner Of The Top Codes

Owner	Owner
EDF (4)	Open Source (9)
EDF, CEA, IRSN, Areva	Opencfd
Fachgebiet Theoretische Chemie, TU München	Research Group Prof. Jürg Hutter / IBM Research
German Weather Service / Max-Planck-Society	RWTH Aachen
GNU GPL: each developer	Technical University Munich
GPL (3) – Gnu Public License	The AIMPRO consortium
Hans De Raedt	University of Illinois
Hinnerk Stüben/ZIB Berlin, Yoshifumi Nakamura/RIKEN Japan	University of Manchester (2)
IBM	University of Stuttgart
INRIA, CNRS, Universities of Lyon, Le Havre, Poitiers	University of Vienna
IPSL and the NEMO Consortium	VASP group

Source: IDC, 2011

Willingness to Invest in Further Development of the Parallel Applications

Nearly all of the surveyed sites said they would be willing to invest in improving the software application codes that are most important for their users, that is, the "first top codes" and the "second top codes." For the first top codes, 36 of 38 responding sites (95%) expressed a willingness to invest (Table 20). For the second top codes, 27 of 28 sites responding to this question (96%) said they would be willing to invest in code improvement. 71% of the reported first top codes had undergone a technical update within the past year – most of the sites invest in improving their key codes as a matter of course.

As noted in the analysis preceding Table 17, the term "major technical update" does not imply that the code has been substantially adapted or fundamentally rewritten, only that it has been modified for greater efficiency and/or compatibility with contemporary HPC systems. As noted elsewhere in this report, a separate recent IDC worldwide study revealed that only about 1% of HPC application codes scale to 10,000 or more cores, and most HPC codes are no more than eight-way parallel. Hence, it is clear that the sites' "major technical updates" typically are not meant to achieve large advances in scalability or fundamental redesigns. In most cases, that work still needs to be done, and the sites are not planning to undertake it on their own.

In sum, their existing patterns of investment ("major technical updates") and expressed willingness to invest in response to the present question would seem to bode well for collaborative investments with more ambitious goals. The next table addresses the sites' willingness to partner with others in advancing their key software codes.

TABLE 20

Willing to INVEST?

Package A - YES	36
Package A -NO	2
Package B - YES	27
Package B - NO	1

Source: IDC, 2011

Willingness to Partner for Further Development of the Parallel Applications

As Table 21 indicates, the great majority of the surveyed HPC sites reported that they would be willing to partner with other organizations to improve their key software application codes. For the first top codes, 33 of 37 responding sites (89%) expressed a willingness to partner. For the second top codes, 24 of 26 responding sites (92%) said they would be willing to partner.

TABLE 21

Willing to PARTNER?

Package A - YES	33
Package A -NO	4
Package B - YES	24
Package B - NO	2

Source: IDC, 2011

Willingness to Work with the European Commission

As Table 22 shows, for the first top codes, 36 of 38 responding sites (95%) said they would be willing to work with the European Commission (i.e., within the European Commission framework) to improve these parallel application codes. For the second top codes, 24 of 26 sites (92%) replied in the affirmative. One of the two sites responding no said that, in its experience, "red tape" associated with European Commission contracts had been excessive "for the relatively small amounts of money involved."

TABLE 22

Willing to WORK WITH THE EU?

Package A - YES	36
Package A -NO	2
Package B - YES	24
Package B - NO	2

Source: IDC, 2011

Most Useful Partners for Improving the Parallel Application Codes

Responses are shown in Table 23 to the question, "What types of partners would be most useful in helping to improve these codes?" As the responses indicate, academic research centers and large HPC centers were popular preferences for partnering, but HPC vendors and industrial users were also frequently cited preferences, being named by 18 sites for their first top codes. Independent software vendors (ISVs) and hardware vendors (OEMs) were less-frequent but still important choices.

TABLE 23

Most Useful Partners (can name multiple partners)

Partners	Partners
Academia / University Experts (16)	Industry users (3)
Academic, Hardware manufacturer	ISVs (4)
Computational methods, scientists (2)	Local user groups
Computer Science, Numerical Mathematics	Open Source developers (2)
Computer scientists, HPC experts, INRIA (2)	Programmers, software engineers (2)
Data producers, data analysts	Research centers
Hardware providers (3)	Specialists in parallel computing from other research institutions (5)
HPC experts (3)	Support for open software through university
HPC experts, numerical gurus, industrial users (4)	Users, software developers (4)

TABLE 23

Most Useful Partners (can name multiple partners)

Partners	Partners
HPC facilities (4)	VASP group itself, local user groups
HPC optimization help	Weather experts in programming
HPC specialists, Computational chemists	.

Source: IDC, 2011

Primary Needs for Improving the First Top Codes

The most important needs for improving the first top codes, the sites reported (Table 24), were adequate internal expertise (cited by 41% of the responding sites) and adequate financial investment (23%). Internal expertise, in this context, means expertise in parallel programming and other skills needed to improve the codes. External expertise (12% of the sites) and partnerships (11%) were also frequently cited needs, as was the need to develop new approaches (11%) for improving their first top codes. The latter result is consistent with the earlier findings (see Tables 8, 9, and 10) in which many of the sites said they felt confident they could scale their first top codes (and second top codes) to exploit substantially larger numbers of processor sockets and cores.

Fewer sites (3%) said they would need to create business cases for improving their first top codes. The scope of the present study did not permit deeper probing into these responses. Based on our other recent research, however, IDC assumes that the academic and government organizations that constitute the majority of the surveyed sites are less accustomed than their industrial/commercial counterparts to needing to create business cases for software development initiatives. In some cases where sites said they would not need to create business cases, IDC assumes this is because the sites in question have already created business cases for improving their most important application codes.

TABLE 24

Primary Needs for Improving the Code - PACKAGE A (check all the apply)

	N	Percentage of Sample
d. internal staff/ experts	30	40.5%
a. money /investment	17	23.0%
f. external expertise	9	12.2%
c. partnerships	8	10.8%
e. new approach	8	10.8%

TABLE 24

Primary Needs for Improving the Code - PACKAGE A (check all the apply)

	N	Percentage of Sample
b. business case	2	2.7%

Source: IDC, 2013

Primary Needs for Improving the Second Top Codes

The responses for the second top codes (Table 25) closely mirror those for the first top codes (Table 24).

TABLE 25

Primary Needs for Improving the Code - PACKAGE B (check all the apply)

	N	Percentage of Sample
d. internal staff/ experts	19	38.0%
a. money /investment	11	22.0%
f. external expertise	7	14.0%
c. partnerships	6	12.0%
e. new approach	6	12.0%
b. business case	1	2.0%

Source: IDC, 2011

Programming Environments for the Sites' Most Important Application Codes

Table 26 shows the programming environments used in conjunction with the sites' first top codes and second top codes. Consistent with previous IDC HPC research studies, the most prominent programming languages are Fortran (still popular for scientific research), C, and C++. There were many mentions of the MPI application programming interface (API), and fewer mentions of the OpenMP API, for use with these programming languages. CUDA, Nvidia's parallel programming framework for use with the company's graphical processing units (GPUs), received a single mention.

TABLE 26

Describe Your Programming Environment (For your top two codes)

Programming Environment	Programming Environment
ANSI/C + MPI	FORTRAN, C, BLAS, FFTW, BLACS, MPI (OpenMP in development)
BLAS/LAPACK libraries, FFT libraries, MPI API, OpenMP	Fortran, MPI (14)
BOINC	Fortran, MPI, OpenMP, netcdf (2)
C, C++ compilers (7)	Fortran, MPI, OpenMP, scientific libraries (2)
C, MPI, Forge, Petsc	Fortran/Python - OS supported : Linux, port on windows performed by the open-source community
C/C++ compilers (GNU), CORBA engines	Fortran95 compiler, FFTw library, BLAS/LAPACK libraries, MPI library (2)
C/C++/Fortran compilers (GNU)	GNU and Intel Compilers, Linux
C++ compiler, BLAS/LAPACK libraries, FFTW library, MPI API	GNU compilers, MPI
C++, MPI, Forge, numerical libraries	GNU compilers, MPICH, Boost, python, Qt
C++, MPI, Python	GPL-based stream. GCC (perhaps Intel compilers). For Python, just an editor. (2)
C++/FORTRAN, MPI, OPENMP (6)	IMB xl compilers, BlueGene/P (2)
C++/Linux / gcc or Intel compilers / MPI/ OpenMP	MPI + OpenMP, f90 (2)
F95, MPI, OpenMP, fat node architectures	MPI, LAPACK, ScaLAPACK, BLAS, Compiler Gnu, Intel, IBM; Darcs, CVS, Perl, Hardware Linux Cluster
F95, MPI, OpenMP, fat node architectures	MPI, Linux
Fortran 90 / Perl / Python / MPI / OpenMP / HDF5 and soon GPU In option : MED / MUMPS / Salome Platform / Tecplot	MPI/OpenMP Fortran90 (or higher) + BLAS library + FFT library: MKL, ESSL or FFTW
Fortran 90, MPI CUDA, Autotools, Netcdf, blas lapack	OpenMPI, C on HECToR (UK supercomputer)
Fortran 90, NetCDF, HDF, CDO, Lapack, Blas, MPI	PETSc/SLEPc, SCALAPACK
Fortran 95 + MPI + Open MP + CUDA	Special supporting software
Fortran and C compilers (6)	Standard Linux programming environment, Fortran90, Fortran77, C, MPI, OpenMP (2)

Source: IDC, 2011

Other Things That Would Help Advance the Application Codes

The responses to this open-ended question varied widely (Table 27), but the most frequent theme was the need for more human expertise in parallel programming and in support for programming.

TABLE 27

What Else Would Help Advance These Codes?

Areas That Would Help	Areas That Would Help
Access to future systems prototypes, GPU programming, PGAS (2)	Leveraging volunteer developers, RCUK funding and industry sponsorship.
Better co-operation of the VASP group itself	Linking up with research organizations to quickly add new capabilities as they are developed.
Better scalability (2)	Machines with significantly more than 262,144 cores
Better support	Make them more robust, flexible, adaptable, portable, reusable
Broader support of open source programming model in general. Anything to break the hegemony of matlab	Manpower, low level programming expertise (e.g. assembler, SSE)
Close cooperation with computers centers who buy the machine	Memory optimization, modularisation
Encourage funding applications for software development	More human resources (more funding) (4)
EU network of developers of computational materials science software to share best practice	More permanent engineers and researchers (2)
Expertise in molecular simulation algorithms and HPC computing/programming (2)	New algorithms (2)
Extensive documentation	Optimization of long range electrostatic calculations and memory allocation for input and output operations.
External technical expertise (2)	Parallelization Know how (4)
Financial support	Patronization of development by industry and EU/RCUK funding for employing developers.
Funding to establish repository of benchmark reference results to aid development of new functionality	Reliable continuity of funding for contract researchers to avoid loss of key personnel
Funding to provide long-term support and development for established codes	Relieve senior developers of teaching duties
Growth of a user and developer community for open source contributions.	Restructuring the code, without breaking usability
Hybrid MPI-openMP coding	Scientific advances in scalable algorithms for current problem domain (2)
Improvements in automatic and scalable “high quality” (ideally hex-dominant) meshing.	Teaching computer science, math, and physics students the mores of open-source programming contributions

TABLE 27

What Else Would Help Advance These Codes?

Areas That Would Help	Areas That Would Help
Improvements in linear solver scalability, improvements in parallel IO,	.

Source: IDC, 2011

3.3 The Software Development Model: Suggested Ideas from the Field Research

This section, based on field research, summarizes ideas and suggestions from the survey respondents themselves. IDC's recommended approach appears later in this report. The study participants had varying suggestions for an effective development model for advancing parallel software needed by Europe's scientific and engineering research communities, but there was significant consensus around the following strategies:

3.3.1 Funding for Software Development Should Last At Least 5-10 Years

Multiple respondents pointed to the familiar IT industry saying, "Hardware is easy, software is difficult." Because fundamental adaptations and rewrites of parallel software require considerable time and expense, it makes sense to undertake them only infrequently. An HPC hardware generation typically lasts only 3-5 years, while a software generation must often last 20-30 years.

This means that parallel application codes and other components of the parallel software stack need to be designed with sufficient robustness, flexibility, and other attributes to operate efficiently across multiple generations of HPC hardware systems, including future systems whose architectures may only be vaguely characterized when the software is first developed.

Respondents stressed that at least 5-10 years are often needed to develop production-quality parallel software that has persistent (20-30 years) value. Most government funding for software development today, in Europe and elsewhere, is for only 1-2 years. This abbreviated timeframe was a major issue raised by software developers in IDC's government-funded study of the generally successful history of the U.S. Department of Energy's ASCI/ASC program, for example.

Respondents in the present study said that funding should be periodically renewable, based on progress toward stated goals.

3.3.2 A Suggested Model for Scientific Software Development

Discussions with study participants suggested ideas for a potential approach to developing a model that would be effective for the development/advancement of parallel software for scientific and related research use. Suggestions included:

- Convene a body of academic/government experts to create a software roadmap that includes milestones for sustained performance.

- Use the roadmap as the basis for procurements (RFPs) for targeted areas of parallel software development.

- Use the peer review process to award funding.

- Use the body of experts to monitor progress and make recommendations regarding renewed funding.

- Work closely with other organizations around the world, where appropriate, to avoid duplicated efforts and to quickly adopt useful advances.

IDC notes that the European Exascale Software Initiative (EESI) has already organized bodies of experts on various areas/aspects of parallel software and in the fall of 2011 will present a roadmap for the development of highly parallel scientific software in Europe. (Many of these experts have been drawn from other organizations with strong track records in parallel software development and advocacy.) In addition, the EESI actively participates in, and collaborates closely with, the International Exascale Software Project (IESP). The EESI and IESP, as their names suggest, focus heavily on highly parallel scientific software and do not have nearly as strong a focus on parallel software codes that are important for industrial/commercial use and competitiveness.

3.3.3 A Separate/Modified Model Is Needed for Parallel Software Used by Industry/Business

- Study participants from the world of industry/commerce stressed that they do not want decisions about the future of their key parallel software made exclusively by "bodies of academics." They said that all too often, initiatives for advancing industrial/commercial HPC software have been automatically referred to academic experts alone. The industrial/commercial respondents want to be integrally involved in developing/advancing key parallel software intended for their use, including contributing opinions, time, and money. This category of respondents includes HPC centers involved in industrial/commercial partnerships, as well as HPC hardware, software, and services vendors that typically serve government, academic, and industrial/commercial HPC users.

- For development initiatives involving key parallel software intended for their use, these respondents often prefer academic-industrial partnerships. Examples of successful partnerships of this kind include BMW and the University of Munich, Porsche, T-Systems, and the University of Stuttgart/HLRS, and others.

- For the development of pre-competitive parallel software whose use will be open to anyone, the industrial/commercial respondents would expect government funding to cover all or nearly all of the development cost. The respondents would expect to contribute expertise and, as appropriate, time on their HPC resources.

- For any part of the development that is not pre-competitive, the investing companies/organizations expect sole use of the software for at least 12-24 months.

- If one private company or party will benefit, the government (e.g., European Commission) and the company/party should co-fund the parallel software development.

- If multiple companies/parties will each gain unique IP benefits, a value-based formula should be created to determine the co-funding portion assigned to each company/party. Each company/party might pay a different amount.

- Several of the industrial/commercial respondents said that the European Commission has used a similar model before for parallelizing industrial/commercial software codes, adding that these many of these codes are now outdated, that is, unable to efficiently exploit today's parallel HPC hardware systems.

3.3.4 Summary of the Suggested Development Model Ideas for Scientific Parallel Software

- Target areas where Europe can lead
- Academic/government experts create a software roadmap
- RFPs are based on this roadmap
- Peer review determines who gets funded
- At least 5-10 years of funding, with periodic renewal eligibility
- Academic/government experts determine renewal
- 100% government (EC) funding for pre-competitive development, co-funding for any competitive portions

3.3.5 Summary of the Suggested Ideas for a Development Model for Industrial/Commercial Parallel Software

- Target areas where Europe can lead
- Academic-industry experts create goals, timeline
- RFPs are based on the goals, timeline
- Fixed-cost contracts are awarded, with fixed deadlines (usually non-renewable)
- Proportion funded by government (EC) varies by contract and IP terms

3.4 Governance Model Ideas for the Software Infrastructure: Suggestions from the Field Research

This section summarizes ideas and suggestions from the survey respondents. The final report will contain the IDC recommended approach.

The study participants had varying suggestions for an effective governance model for advancing parallel software needed by Europe's scientific and engineering research communities, but there was significant consensus around the following strategies:

3.4.1 Establish a Coordinating Body within the European Commission

The study participants generally thought that parallel software development for scientific and industrial/commercial use is important enough to establish a coordinating body within the European Commission to steer the development efforts. Most respondents said they understood that establishing a new entity within the European Commission might not be easy and might encounter opposition.

- The new EC body would coordinate Europe-wide R&D initiatives for the whole HPC "ecosystem," including hardware, software, networking, and other areas, to help ensure that R&D for Europe's HPC ecosystem proceeds in a consistent, integrated, efficient manner. This level of coordination would not affect the autonomy of EU Member States over their national HPC strategies; it would apply only to Europe-wide HPC R&D in which the European Commission plays an important role. Also, the respondents stressed that "coordination" is principally a mediating function and would not confer sole decision-making authority. The EC body would work with the Member States, HPC centers, HPC vendors, Europe's scientific and engineering communities, and others to plan and advance R&D initiatives that would include initial development and appropriate maintenance and advancement of the targeted technologies and capabilities.

- Where parallel software is concerned, the EC body would coordinate the end-to-end software development and funding process, including the bridge period to commercialization.

- In IDC's 2010 study for the European Commission, *A Strategic Agenda for Supercomputing in Europe: HPC 2020*, the European HPC community, including government officials overseeing HPC programs, generally agreed that the PRACE program could serve as the coordinating body for a European HPC strategy, but only if PRACE's mission were substantially expanded to embrace the whole HPC ecosystem, including access to hardware systems for industry, software development, and other areas. When directly presented with the challenges of developing/advancing highly parallel software benefiting all of Europe, respondents in the present study thought the appropriate path would be to establish a coordinating body within the European Commission that worked closely with organizations including PRACE as the lead hardware organization, EESI as a leading software development advisory group, and others. Among the goals would be co-development of hardware and software plans, to the extent feasible, and coordination of software and hardware initiatives to maximize the value of the large investments made in PRACE Tier-0 supercomputers.

- Ideally, co-design centers-of-excellence should be established to focus on the advancement of technologies, HPC resources, and methodologies for scientific and industrial/commercial domains that are most important for Europe (possibilities: climate/weather research, energy research, automotive research, physics research, and others). These centers-of-excellence (a different name might be chosen) would typically be located at existing Tier-0 or other appropriate HPC facilities.

- The EC coordinating body might report to the Directorate for Emerging Technologies and Infrastructures and should establish appropriate relationships with other Directorates whose areas will benefit from the coordinating body's mission and activities, such as:
  - Climate Action
  - Competition
  - Economic and Financial Affairs
  - Energy
  - Enterprise and Industry
  - Research and Innovation

3.5 Analysis of the Codes' Intellectual Property Rights

The intellectual property (IP) rights for most of the key application codes used by the surveyed sites were held by European organizations. Of the 62 codes whose IP ownership was identified, 40 (66%) were solely owned by European organizations, 2 (3%) had multinational ownership that included a European co-owner, 15 (24%) were open source, and 6 (10%) had non-European IP owners.

The study participants frequently said that satisfactory assignment of IP rights, and avoidance of excessive "red tape," would be important determinants for their participation in European Commission-directed software development initiatives. In the recommendations report, IDC has recommended an approach to address the IP rights issues in a more streamlined fashion.

4.0 RECOMMENDATIONS FOR ADVANCING PARALLEL SOFTWARE: SITUATION BACKGROUND

This section of the report contains the background information leading to IDC's recommended strategy for pursuing HPC parallel software development on a European basis. The section begins with an overview of the importance of pursuing a European strategy and continues with an assessment of Europe's circumstances and prospects for success in this endeavor. Following this are recommendations for the following:

- A methodology for creating an infrastructure for parallel software development
- A financial model and funding mechanisms
- A governance model and management scheme
- Actions to be taken at the European and Member States levels to develop an open, sustainable parallel software infrastructure for computational science and engineering in Europe

4.1 Use HPC Parallel Software Development to Help Close Europe's Innovation Gap

4.1.1 Europe's Economic Growth Is Threatened by an Innovation Gap

Commenting recently on the EU Innovation Union, launched in October 2010, Robert-Jan Smits, Director General for Research and Innovation of the European Commission, noted that:

> *"…research and innovation are key strands of the Europe 2020 strategy. Stark figures confront this ambition to use knowledge as a driver for sustainable growth. Albeit with large internal variations, Europe consistently spends less than 2 per cent of GDP on research and development, only two-thirds of that in the US and a little more than half the Japanese figure. Meanwhile, China's investment is growing year by year and will be on a par with Europe in a few years. The EU Innovation Union Scoreboard tells a similar story: a big innovation gap with Japan and the US, with China (not to mention India and Brazil) quickly catching up."*

4.1.2 HPC Is a Proven Tool for Accelerating Innovation

Why is it important to invest in HPC for innovation? HPC (also known as supercomputing) has been firmly linked to scientific and industrial innovation in many reports and studies and today major companies and countries recognize it as a key tool for future growth and innovation. In addition, HPC investments have greatly advanced the state of the art in general computer technology.

The role of HPC-based modeling and simulation as the third branch of the scientific method, complementing traditional theory and experimentation, has been well documented. The tight linkage between HPC and industrial/commercial innovation and competitiveness is also well established. In one worldwide IDC study, 97% of companies that had adopted supercomputing said they could no longer compete or survive without it.

To sustain economic growth within the gathering storm of global economic competition alluded to by Robert-Jan Smits, Europe should heavily exploit HPC's proven value as a powerful accelerator of scientific and industrial innovation. The adoption of HPC is critical not only for scientific research, but also in a growing array of industrial/commercial domains, including the automotive and aerospace, financial services, pharmaceutical/medical, energy, and consumer products industries. HPC modeling and simulation is used today to develop products ranging from cars and airplanes to new cancer-fighting drugs, electronic devices and even potato chips and diapers (nappies).

Political leaders are increasingly recognizing HPC's crucial value for driving innovation and competitiveness, for example:

- In his 2006 State of the Union address, U.S. President George W. Bush promised to trim the federal budget, yet urged more money for supercomputing. President Obama also mentioned supercomputing prominently in his 2011 State of the Union address.

- In 2009, Russian President Dmitry Medvedev warned that without more investment in supercomputer technology, Russian products "will not be competitive or of interest to potential buyers."

- In June 2010, Rep. Chung Doo-un of South Korea's Grand National Party echoed that warning: "If Korea is to survive in this increasingly competitive world, it must not neglect nurturing the supercomputer industry, which has emerged as a new growth driver in advanced countries."

4.1.3 Europe Is Well Positioned To Pursue Global HPC Leadership in Targeted Areas of Science and Engineering

In July 2010, IDC submitted a final report to the DG Information Society of the European Commission, entitled *A Strategic Agenda for European Leadership in Supercomputing: HPC 2020*. On the basis of extensive field research and market criteria – especially Europe's existing global strength in these areas – the report recommended the following domains as candidates Europe should target for world leadership by 2020:

- **Weather and Climate Research.** Europe already has world-class, and arguably world-leading, expertise in this increasingly important field. Weather and climate research affect many other scientific and engineering fields, including economic planning, agriculture, aviation, tourism, energy research (e.g., the carbon cycle), and even automotive engineering (e.g., emissions).

- **Clean And Sustainable Energy,** including nuclear reactor design/operation, oil and gas exploration, smart electrical grids, clean water, fusion energy, and other alternative energy research. This domain is a natural partner for weather and climate research. Europe already has world-class strengths in this domain, which covers a lot of scientific and economic ground and may become the most important domain of all in the not-distant future.

  - Europe is already one of the most "green-conscious" regions of the world.

- The first experimental fusion energy reactor (ITER project) is scheduled to go into operation in France.

- Within Europe there are large-scale wind turbine farms and other leading-edge alternative energy sites in operation or in the planning stages.

- Europe has amassed considerable scientific and engineering expertise in this field.

- **Automotive and Aerospace Design.** Europe today has world-class capabilities in the related fields of automotive and aerospace design engineering, including computational fluid dynamics (CFD) and computational structural analysis expertise and software creation. EU support for the European Aeronautic Defense and Space Company (EADS) is an existing initiative that will presumably continue.

- **Bio-Life Sciences**. Europe also has world-class strengths in the bio-life sciences sector, with leading pharmaceutical and medical technology firms, supported by deep academic expertise in the life sciences.

- **Particle Physics and Related Fields.** Europe's investment in the Large Hadron Collider (LHC) at CERN has already set the stage for global leadership in particle physics, and in HPC-based model and simulation to support this work. The LHC is already attracting top particle physicists from throughout the world – a good example of how the "brain drain" can be reversed in Europe's favor. Even before the availability of the LHC, European particle physicists at CERN and other institutions were on a par with the best in the world.

- **Modeling of Materials/Molecular Dynamics.** The Car-Parrinello method that is central to molecular dynamics research was invented in Europe and molecular dynamics is used in many scientifically and economically important fields, including materials science (development of new materials, aging of materials), alternative energy (improved design for solar cells, wind turbines, etc.), drug discovery and other biomedical research, nanotechnology, product engineering, and more.

4.1.4 Some Of Europe's Best Firms Are Ahead in Exploiting HPC for Innovation

An IDC study for the Washington, DC-based Council on Competitiveness showed that European best-in-class automotive and aerospace companies have pushed HPC usage much deeper into their organizations on average than U.S. tier 1 firms — more frequently extending its use from traditional upstream applications in R&D and design engineering into high-value downstream uses, such as manufacturing and production. The study also found that some of the European automotive, aerospace and bio-life sciences firms more often require their suppliers to use HPC than occurs in North America. The European and North American energy firms are on a par in this regard. (See the report: "*Advance: Benchmarking Industrial Use of High Performance Computing for Innovation".* A Worldwide study conducted by IDC for the Council on Competitiveness and the U.S. Department of Energy. 2008. Downloadable without charge at http://www.compete.org/publications/detail/486/advance/)

4.1.5 Parallel Software Development Is the Key to HPC Leadership

The goal of HPC investments should not be HPC leadership for its own sake, but to exploit HPC's proven value as a powerful accelerator of scientific and engineering innovation, which will help to sustain economic growth. HPC is ultimately about what happens when the world's most creative scientific and engineering minds are accelerated by using the most powerful class of computational tools. It is for this reason that political leaders, including U.S. President Obama, Russian Federation President Medvedev, China, and others have publicly stressed the need to invest more heavily in HPC.

The EU and the Member States already make substantial investments in large-scale supercomputers through the PRACE program, and IDC's 2010 study, *A Strategic Agenda for European Leadership in Supercomputing: HPC 2020*, recommended additional investments to ensure that Europe becomes and remains globally competitive in its strategic investments in HPC. But the scale of the investments needs to be appropriate given the amount being spent/invested in HPC by other regions of the world.

But in the present era of petascale computing and in the exascale era that will begin before 2020, sustained performance leadership on real-world HPC applications and workloads will be determined far more by software advances than by hardware progress. For multiple decades the HPC community, in Europe and elsewhere, has focused primarily on highly parallel hardware systems, to the detriment of more advanced software. A 2011 IDC worldwide study of HPC data centers revealed that 52% of their software codes are running on one node or less of an HPC system. Only 12% of the codes could exploit 1,000 or more cores, and only 1% could exploit 10,000 or more cores (*IDC HPC End-User Study of Industry/Application Trends in Technical Computing*.) The largest HPC hardware systems today contain more than 200,000 cores, and million-core supercomputers will begin to arrive before the end of this decade.

The underlying problem is that modern HPC hardware, with large numbers of CPU cores whose memory capacity and bandwidth typically decrease with each new CPU generation, is causing a mismatch with existing application software. This drives a need to fundamentally redesign and rewrite HPC application software for greater parallelism, in order to perform well on contemporary and future HPC systems.

Historically the bulk of sustained performance progress on real-world applications has come not from hardware parallelism but from Moore's law–governed, generational jumps in single-threaded processor speeds. This evolutionary "gravy train" hit a heat-and-power wall half a dozen years ago. Since then, several alternative processor strategies (notably GPGPUs and FPGAs) have emerged to help get the train moving forward again, but these too will need empowerment from software advances. Nations and regions that invest in moving their parallel application and systems software forward will see tremendous competitive and scientific advances.

4.2 Europe Can Become A Leader

4.2.1 Europe Can Become the Global Innovation Leader with Appropriate Investments in Parallel Software Development

If Europe quickly begins to provides appropriate, sustained investment and support for parallel software development in the targeted areas, IDC believes that by the year 2020 or soon afterward, Europe can become the global innovation leader in these scientifically and economically important domains in which Europe already has substantial strengths, and in which competition is already based heavily on HPC usage. In addition, Europe could become a major source for the next generation of HPC application software codes.

The field research conducted for this study shows that Europe would start this pursuit from a position of considerable strength. A high proportion of the key application codes that are used today within the sample population of European HPC sites were created in Europe, and European ownership of these codes (intellectual property rights) is also substantial. And with wise, focused investments, Europe could increase its share of key application code ownership. Europe has a number of globally successful scientific and engineering software firms, a larger number of nationally and regionally successful software firms, and is strong in many important areas of parallel software development.

Software Areas To Invest In And Those To Avoid

The EU should avoid investing in application codes that are already being parallelized today and those at the other extreme that are very hard to scale beyond their current limits due to the nature of the underlying science. This leaves a large set of codes that could be advanced and provide Europe a great strength.

HPC Technology Leadership Areas

In support of leadership in the targeted scientific and engineering domains, IDC recommends that Europe should also target global leadership in the following HPC technology areas. Europe has world-class strengths in all of these areas today, and each of these areas will become far more crucial for leadership by the year 2020.

- **Making Large-Scale HPC Systems Highly Productive As Innovation Tools for Users.** A major part of this productivity challenge is enabling parallel software applications to exploit large-scale, complex HPC systems more efficiently than they do today. Another part is helping scientific and engineering researchers, as well as students, to learn how to run their application codes on the large-scale HPC systems.

- **Developing Highly Scalable Algorithms and Software Applications,** especially related to the targeted scientific and engineering domains, but by no means limited to these domains. Most HPC codes were originally written to run on one single-threaded processor. Many of these codes have never been fundamentally re-written. Instead, these "dusty deck" codes have been laboriously modified over time to enable modest scalability. In other cases, the lack of more scalable algorithms curbs sustained performance. This is a solvable problem, although there may not be enough people on Planet Earth with the right combination of brainpower and experience to solve it for every deserving application.

- **Facilitating Parallel Programming for Future-Generation, Heterogeneous Supercomputers.** Future-generation supercomputers will employ up to one million or more processors cores and multiple types of processors/co-processors. The ability to create programs that efficiently exploit these highly parallel, heterogeneous supercomputers will be extremely valuable. This challenge will tap into European expertise in performance modeling and performance prediction.

4.2.2 Exascale Software Is Not Enough

The European Exascale Software Initiative (EESI) was created to help organize Europe's efforts in these areas, in part by bringing together islands of software talent within Europe to identify the requirements and daunting challenges associated with developing exascale (i.e., extremely parallel) software. The EESI final report scheduled for submission later this year will no doubt provide useful new perspectives for exascale software development in Europe. (As we indicate later in this report, IDC believes the impressive organizing effort that has gone into the EESI should not be abandoned, but needs to be followed by a longer term action-results oriented approach. This group of European software and related experts will continue to be of great value for parallel software development in Europe.)

Many important parallel software applications in the targeted areas, especially the engineering sciences, do not have the potential to exploit substantial fractions of exascale supercomputers in this decade. This is okay, as it is important to keep in mind that investing to enable a key automotive application to advance from 8-core scalability to 64-core scalability, for example, could have great competitive and economic value for Europe's automotive industry. And in fact, EESI and other organizations have also been assessing the scalability requirements of industrial applications.

The Value of Less Highly Parallel, Iterative Software Solutions

More and more often, scalability is constrained by the limitations of the underlying, known science. Where the science doesn't support large single runs of the problem, smaller iterative runs allow users to home in on productive solutions. Long-established examples include stochastic modeling in the financial services sector and parametric modeling in the design engineering realm (e.g., automotive and aerospace design). Newer examples are the increasingly complex ensemble models in the weather/climate domain, and a whole host of genomics and proteomics applications.

All of this is to say that iterative solutions are no less valuable simply because they are less scalable than the highest-performing HPC codes. Iterative solutions will continue to proliferate as the algorithms and known science in more and more domains fail to keep pace with the rampant parallelism and peak performance of ultrahigh-end supercomputers.

4.3 A Framework for Establishing an Infrastructure for Highly Parallel Codes

This section of the report assesses the European environment for parallel software development and recommends a methodological framework for creating a European infrastructure to support this development.

4.3.1 Assessment of the Current Level of Cooperation, Competition and Potential Synergies for Parallel Software Development in Europe

Cooperation and Competition

Where HPC in general and parallel software development in particular are concerned, Europe today is a vibrant patchwork of national, regional, local, and Europe-wide initiatives and activities. This is hardly surprising, given the diverse histories, cultures, and geopolitical circumstances of Europe's nations and regions. On one hand, this uneven diversity poses challenges for the leveling processes of negotiation and cooperation. On the other hand, Europe's diversity contributes to a richness of perspective that is sometimes lacking in more homogenous, politically unified nations and regions.

There is a steadily growing willingness to cooperate across Europe. This cooperation is evident in an array of successful HPC research programs, including PRACE, DEISA, CEA-DAM, CERN, HPC Europa, Blue Brain, EGEE, ESA, Forschungzentrum Juelich, HWW, GEANT, and others too numerous to name here. It is also evident in initiatives such as EESI and EGI.

Competition will remain within Europe's HPC community and the scientific and engineering communities it serves. There will be competition for funding and recognition within and among nations, between Europe's large and smaller nations, between scientific and industrial HPC users, and between HPC-supported research disciplines. (One survey respondent said, "The computational physicists and chemists are so well organized, they get the lion's share of the funding, and everyone else lives on crumbs.") IDC believes that competition, when it is properly channeled and when the rules are fair, can be as important as collaboration for spurring innovation within Europe.

In summary, Europe's HPC community today is characterized by "co-opetition," a generally healthy mix of cooperation and competition. This spirit of co-opetition extends to the European HPC community's relationships with non-European HPC and scientific/engineering research communities. Europe's HPC centers, and the scientific and engineering researchers they serve, are involved in many productive, mutually beneficial collaborations with their non-European counterparts. Hence, Europe needs a parallel software development strategy and infrastructure that accelerate European innovation without discouraging continued open, beneficial collaborations with non-European partners. But these current co-opetitions need more investment and direction in order to more directly advance application software and better apply HPC for innovation.

Potential Synergies

Today, Europe has many of the important elements needed to create a sustainable European ecosystem for parallel software development benefiting the European Union and the Member States, although these elements have not been brought together as closely as needed yet. The following multinational and Europe-wide initiatives stand out as successful examples (among others) of synergy that might be more fully exploited ("synergized") in support of a European parallel software development strategy and infrastructure.

- **PRACE.** The PRACE program has begun to populate Tier-0 HPC centers with petascale supercomputers to support advanced scientific research, and in at least one case also to support advanced engineering research (e.g., HLRS).

Parallel software developers need access to large-scale supercomputers, and to expertise in using these systems. The Tier-0 centers already possess substantial expertise in helping leading computational scientists and engineers to exploit high-end supercomputers. And while most of this activity has been on the side of scientific research, PRACE has held three successful industry workshops, implying that greater access to PRACE Tier-0 systems for advanced industrial research may follow sometime soon.

- **DEISA**, the Distributed European Infrastructure for Supercomputing Applications, is a consortium of leading national Supercomputing centers that aims at fostering the pan-European world-leading computational science research. The European Commission-funded DEISA infrastructure is based on a tight coupling of eleven national supercomputing centers from seven European countries, using dedicated network interconnections of GĖANT2 and the NRENS.

- **EGI** is a foundation established under Dutch law to create and maintain a pan-European Grid Infrastructure (EGI) in collaboration with National Grid Initiatives (NGIs) and European International Research Organisations (EIROs), to guarantee the long-term availability of a generic e-infrastructure for all European research communities and their international collaborators. Its mission is to enable access to computing resources for European researchers from all fields of science, from High Energy Physics to Humanities.

- **EESI.** The European Exascale Software Initiative is scheduled to submit a final report later this year on the requirements and challenges for highly parallel software development in Europe – with a strong focus on software with the potential to exploit substantial fractions of future exascale supercomputers, and a lesser but still important focus on industrial software. Software vendors from multiple continents are involved in EESI, as are Europe-based HPC systems vendors Bull (France) and Eurotech (Italy), along with many systems vendors based outside of Europe.

- **National and Regional Initiatives within Europe.** In addition to Europe-wide initiatives, there are many outstanding national research efforts within Europe that could be invited to contribute perspectives and expertise on behalf of the European parallel software development. Cited as successful by Europe's HPC community in IDC's 2010 study for the European Commission were the following: CEA-DAM, CERN, HPC Europa, Blue Brain (EPFL), EGEE, ESA, Forschungszentrum Juelich, and GĖANT – and IDC is aware of many other important national and regional talent centers for highly parallel software development. Successful academic-industrial partnerships, such as HLRS/University of Stuttgart's HWW joint venture with Porsche and T-Systems; the 21-member, HLRS-led Automotive Simulation Center Stuttgart (http://www.asc-s.de/); the Forschungszentrum Juelich-BMW partnership; HPC industrial outreach programs at SARA (Netherlands), the Barcelona Supercomputing Center (Spain), and elsewhere; the weather and climate research initiatives at ECMWF and the UK Metereological Offiice, Deutscher Wetterdienst, Meteo France, Meteo Swiss (with the help of the CSCS supercomputer center), and others; and the EADS and CERN initiatives have provided models for the future.

- **Exascale Research Initiatives Within and Involving Europe.** In addition to EESI, numerous exascale research initiatives based in Europe, or involving

Europe, could be invited to help drive the European parallel software development strategy forward. Here is a partial list of these initiatives:

- **Europe: Exascale Research Initiative.** In September 2009, Cray announced the Exascale Research Initiative, a collaboration with the Swiss National Supercomputing Center (CSCS), the University of Edinburgh (EPCC), and the University of Lugano to explore new ideas and technologies needed to deliver a supercomputing system capable of sustaining exaflop application performance.

- **Global: International Exascale Software Project (IESP).** The original impetus for this project came from Dr. Jack Dongarra, University of Tennessee, and the project quickly became international in scope, including Europe. IESP's goal, as its name suggests, is to explore and address the challenges associated with developing software that can efficiently exploit future exascale supercomputers. IESP has held workshops in North America, Europe, and Asia.

- **Europe: The Bull-INRIA Exascale Partnership**. INRIA (the French National Institute for Research in Computer Science and Control), and Bull have a partnership agreement to carry out joint research to help overcome the challenges posed by the coming generation of 'exascale' supercomputers. A roadmap has already been set out, enabling the teams involved to work together on supercomputer architecture, software environments and services.

- **Europe: Exascale Innovation Center.** In March 2010, IBM and the Forschungszentrum Juelich announced the Exascale Innovation Center, with the goal of developing hardware and software for an exascale supercomputer by the year 2020. The plan is to develop a prototype computer by 2015 for deployment in 2019.

- **Europe: ExaCluster Laboratory.** On May 31, 2010, Forschungszentrum Jülich, Intel, and ParTec signed a multi-year agreement to create a new ExaCluster Laboratory (ECL) located on the campus of Forschungszentrum Jülich. The Lab will explore the key challenges of building Exascale computing systems; it will initially employ about a dozen researchers expercting to triple the staff over time. It will innvestigate in systems management software for large heterogeneous supercomputers with a focus on scaling, including open exascale runtime system software, software tools, and simulation software.

- **Europe: Flanders ExaScale Lab.** Based at the imec research facilities in Leuven, Belgium, this collaboration between Intel and five Flemish universities aims to develop software to run on future Intel-based exascale computer systems. This lab will be part of the Intel Labs Europe network.

- **Europe: The Exatec HPC Lab.** Created by Intel with French partners CEA, GENCI, and the University of Versailles, this initiative is focused on exascale hardware and software optimization.

- **Global: Hewlett-Packard Exascale Computing Lab.** The goal of this initiative, which takes place within HP Labs, is to build compute fabrics for next-generation IT solutions utilizing a cross-layer, inter-disciplinary approach across conventional divisions – of CPU, system, and data center; compute, storage, networking, and packaging; hardware, firmware, systems, and software; over performance, power, availability, and manageability; with cost-effectiveness and agility.

4.4 Strengths, Weaknesses, Opportunities and Threats (SWOT)

This section is an assessment of the strengths, weaknesses, opportunities and threats of the current scenario in terms of capability to respond to the challenge of emerging parallel software demand. The preceding section lays the foundation for this SWOT analysis by assessing Europe's position regarding cooperation, competition, and potential synergies.

Strengths for European Parallel Software Development

- **Europe has a long and strong history of parallel software experience and capabilities.** Europe possesses world-class expertise in the development of mathematical models and algorithms, HPC system performance prediction and modeling, parallel systems software, and parallel programming for code development and advancement.

  - Separate IDC research studies have consistently shown that one of the major impediments to advancing parallel software development is a serious worldwide shortage of parallel programmers and parallel algorithm developers. This shortage affects Europe as well, but Europe is as strongly positioned as any global region – and stronger than most.

- **There is a substantial, vibrant market within Europe for parallel software used on HPC systems.** Appendix 6.3 lists a number of the software suppliers and organizations in Europe. A few European parallel software vendors, including Dassault Systèmes (France) and NAG (UK) have achieved global visibility and success. Highlights about this market include the following:

  - The worldwide market for HPC software – including application software and the wide variety of system software, libraries and tools collectively described as middleware – amounted to about €3 billion in 2009 and is projected by IDC to exceed €4 billion in 2013. Europe's share of this market in 2009 was about €800 million, or about 26% of the worldwide market.

  - The worldwide software market associated with high-end supercomputers (priced at $3 million (€2.25 million) and above) amounted to about €813 million in 2009 and is expected to exceed €1.25 billion in 2013. Europe's share of the 2009 market was about €105 million, or about 13%. IDC predicts that Europe's share in 2013 will grow to about 20% as Europe invests more in high-end HPC supercomputers and software required to make them productive.

- **The EESI initiative, under the direction of Jean-Yves Berthou (Électricité de France), has already identified and enlisted the cooperation of a large number of European experts in parallel software development for scientific and industrial research.** IDC strongly recommends that this group not be wholly disbanded after submitting its final report later this year. See the section, "Steps and Procedures for Creating the Infrastructure," below. EESI is not the only useful, Europe-wide parallel software initiative, but it is the most inclusive to date.

- **There is momentum toward providing greater access for industrial and engineering research.** EESI and PRACE so far have focused primarily on European HPC requirements for scientific research. In IDC's 2010 study for the European Commission, 100% of the more than 200 survey respondents from Europe's HPC community said that HPC access is also important for industrial innovation and competitiveness, and 66% said HPC access is extremely important for industry. It is heartening that PRACE has held three successful industry workshops, presumably with the intent to provide industry with greater access to PRACE systems in the future, and that EESI will include assessments of industrial/commercial parallel software requirements in its report due to be submitted later this year. It is also encouraging that a PRACE Tier-0 supercomputer will be installed soon at the Gauss Centre for Supercomputing (GCS) based at HLRS (Stuttgart), where the system will be used for industrial and engineering simulations on behalf of Germany's three national supercomputer centers.

Weaknesses for European Parallel Software Development

- **Europe is a kaleidoscope of 27 nations.** Achieving consensus is inherently more challenging in Europe, where 27 nations are involved, than in single nations such as the U.S., Japan, or China. But as noted earlier, in recent years IDC has witnessed a strong and steadily rising willingness to cooperate within Europe.

- **The European HPC software market is limited in size.** Europe accounts for only about one-quarter of the worldwide market for HPC hardware systems and parallel software. This means that Europe-based HPC vendors generally cannot thrive and continually fund world-class innovation unless they can match the investments of competitors who have access to the larger worldwide market. A number of Europe-based software vendors have been successful at this, but they are exceptions. In any case, the limited size of Europe's HPC market constitutes a barrier for Europe-based parallel software vendors in comparison with U.S.-based vendors that have access to a considerably larger domestic market as a launching pad for worldwide marketing, sales, and distribution. It would help greatly if some or most of the new software can be targeted for both Europe and worldwide markets.

  - **Most key European parallel applications have very limited distribution.** The good news is that field research for this study revealed that a substantial portion of the key software applications used at the surveyed European HPC sites were created in Europe and are owned or otherwise controlled by European organizations. The not-so-good news is that the usage of many of these applications is limited to no more than a few sites. There are relatively few examples, such as Dassault Systèmes and NAG, of parallel software codes created in Europe that have achieved widespread popularity within the European or global HPC communities.

- **Europe has a strong bias toward scientific research, as opposed to engineering research**. Given that HPC-based modeling and simulation have established themselves as the third branch of scientific inquiry, complementing traditional theory and experimentation, it is highly appropriate for Europe to provide scientific researchers with access to most powerful computational resources possible.

- IDC's field research for the 2010 European Commission study revealed that a substantial majority of the surveyed European scientific researchers (not to mention industrial researchers) believed it is also important to provide European industry with greater access to these resources. It is not feasible to pursue the twin goals of accelerating innovation and economic competitiveness without industry on board. This is an evolutionary lesson learned by the INCITE program in the U.S., which opened up access to industry starting in 2005, mainly for advanced, pre-competitive research (the majority of INCITE resources remain allocated for science). In the field research for the present study, industrial respondents consistently remarked on Europe's excessive bias toward science in HPC initiatives, noting that industrial representatives were often excluded even when they would be the end-buyers and end-users of the parallel software under consideration.

- **Many European software initiatives are targeted at collaborations and developing reports.** As opposed to longer term programs to actually develop new software.

Opportunities for European Parallel Software Development

- **Capture global leadership in scientific and industrial areas where Europe is already strong.** As described earlier, these areas include weather and climate research, clean and sustainable energy, automotive and aerospace design engineering, bio-life sciences, particle physics, cloud computing, molecular dynamics/modeling the properties of materials, and exascale applications, across a wide spectrum of disciplines. These are all scientifically and economically important areas in which Europe has substantial strengths today that provide the basis for pursuing global leadership.

- **Parallel software development is the key to future HPC leadership.** Relatively few scientific application codes as written today, and even fewer industrial applications, will be able to exploit substantial fractions (say, 20% or more) of the largest supercomputers in the near-term future. HPC hardware has raced out ahead of most parallel software, most of which was originally written 10 to 20 years ago to run on far less parallel and less heterogeneous hardware systems. IDC believes that in the next five to 10 years, many HPC applications will need to be substantially adapted or wholly rewritten to exploit HPC hardware systems more efficiently. The nation or region that undertakes this task soonest, and with the greatest determination and support, will be well positioned to gain worldwide HPC leadership – which is the key to innovation leadership in targeted scientific and engineering domains.

- **No other nation or region is accepting this challenge today.** Although parallel software development initiatives exist in many areas of the world, no nation or region to date has committed itself, politically and financially, to a strategic program of parallel software development. Most are funding major hardware development and acquisition programs. Programs that originally included software as part of the strategic goals have typically downshifted, usually because of inadequate funding, to the more traditional course of pursuing HPC hardware advances without attendant parallel software advances. This gap creates a strong opportunity for Europe to seize the most valuable form of HPC leadership for innovation and competitiveness – superior sustained performance running real-world parallel software applications.

Threats for European Parallel Software Development

- **Over-enthusiasm for Europe's position in the global "Petascale Race."** It is important for European researchers to have access to the most powerful supercomputers Europe can afford, and it is genuinely exciting that the PRACE program has already begun to deploy petascale supercomputers at Tier-O centers. If Europe is to use its HPC resources effectively to help close Europe's scientific-industrial innovation gap, it will be important not to get excessively caught up in the prestige of the international "Petascale Race" that will morph into an "Exascale Race" before the end of this decade. Europe and the Member States will need to strike a sober balance between investments in extreme-scale supercomputers and investments in the parallel software and other resources needed to make these systems highly productive for significant numbers of scientific and industrial/engineering researchers.

- **Failure to adequately recognize and finance parallel software development within Europe's HPC strategy.** Although Europe's Gross Domestic Product (GDP) is about the same as that of the U.S., Europe spent only about half as much on HPC in 2009 at the U.S. did. As noted at the start of the recommendations section of this report (Section 4.1), Europe also suffers from an "innovation gap" in relation to other nations and regions – a gap IDC believes increased HPC investments can cost-effectively help to close. Europe has already begun ramping up its investments in HPC hardware systems. Without appropriate investments in parallel software development, the hardware investments will have a lower return-on-investment and Europe will miss a unique opportunity to achieve worldwide innovation leadership in targeted scientific and industrial domains.

- **Failure to provide a central authority at the European level to carry out the HPC and parallel software strategy.** In the extensive field research IDC conducted for the 2010 study for the European Commission, respondents made it clear that the EU needed to drive a European HPC strategy because no member state acting alone could afford to compete for leadership with the U.S., Japan, or China. In response, IDC recommended for pragmatic reasons that the generally well-respected PRACE program could drive Europe's HPC strategy, but only if the mission of PRACE well greatly expanded to embrace the development of software and other resources as strongly as PRACE has embraced HPC hardware systems. Field research for the present study produced a consensus that a new body should be created within the European Commission to coordinate Europe's HPC strategy on a holistic basis, working in close collaboration with PRACE and other existing programs and initiatives. Respondents stressed that this should be a coordinating body and should not be given sole decision-making authority. This shift did not indicate any lessening of admiration for PRACE, but rather a desire to move ahead more quickly than might be feasible if PRACE were required to expand its mission to embrace multiple new functions.

- **The continuing economic downturn.** In the midst of the continuing European and global economic crisis, funding for HPC and parallel software development must compete for funding with other important priorities. It is important for both the EU and the national governments to recognize – as has already happened in the U.S., Japan, China and Russia – that HPC and related parallel software development can be a crucial part of the solution to the economic crisis, by boosting European innovation and economic competitiveness. It is with these goals in mind that the U.S., for example, increased HPC funding 38.5% in 2009, the most difficult year of the global economic recession.

- **Other nations and regions could become more focused on parallel software development at any time.** The greatest potential threat comes from the United States, which has not made a strong financial commitment to parallel software development yet, but which could do so in the future. Lesser threats are from China, which is likely to continue focusing on expanding its domestic HPC market for at least a few years, and Japan, India, and Russia, which historically have not been major developers of parallel software for the global HPC market place.

4.5 Drivers and Barriers Towards the Forming of a Unified European Community

This section addresses the main drivers and barriers, economic, technological, practical, organizational, cultural, towards the forming of a unified European community of computational scientists on e-infrastructures.

Main Drivers

- **Economic:**

  - Europe has been under-spending on HPC and from 2007 to 2009, Europe lost 10% market share in the worldwide HPC supercomputer market space, a very significant decline. To catch up and keep pace with competing nations and regions, Europe has been working to increase its HPC investments.

  - As Robert-Jan Smits, Director General for Research and Innovation of the European Commission noted recently, Europe is suffering from an "innovation gap" in relation to other nations and regions. This gap needs to be addressed for Europe to sustain economic growth in the future. IDC believes that HPC, as a proven innovation accelerator, can help to close the innovation gap.

- **Technological:**

  - In Europe and elsewhere, HPC hardware system technology has raced out ahead of the parallel software development needed to efficiently exploit the hardware systems. Few HPC software applications today can exploit a substantial fraction (say, 20% or more) of the largest supercomputers. IDC believes that in the next five to 10 years, many HPC parallel software applications will need to be heavily adapted or wholly rewritten to take better advantage of large-scale supercomputers.

  - Europe, especially through the PRACE program, has already begun deploying petascale supercomputers at Tier-0 HPC centers. Substantial parallel software development will be needed to increase the returns-on-investment from today's petascale supercomputers and the exascale systems that will begin arriving before the year 2020.

- **Practical:**

  - Europe has a long and strong history of parallel software experience and capabilities. Europe possesses world-class expertise in the development of mathematical models and algorithms, HPC system performance prediction and modeling, parallel systems software, and parallel programming for code development and advancement.

- **Organizational:**

  - Europe has established strong momentum in establishing cooperative HPC initiatives, including PRACE, DEISA, EGI, and others. These initiatives, along with others at the national and regional levels, could become important contributors to a European infrastructure and strategy for parallel software development.

  - The EESI initiative has already identified and enlisted the cooperation of a large number of European experts in parallel software development for scientific and industrial research.

- **Cultural:**

  - Although national identities remain strong within Europe, identification with Europe has been increasing, especially in the context of global economics and competition.

  - Among scientific and industrial researchers, cooperation across national boundaries has become much more common. Their world has evolved from the former era of lone scientists working in isolation, to teams of scientists collaborating across distances. This bodes well for further unifying Europe's scientific and engineering research communities, primarily around domain-specific challenges.

Main Barriers

- **Economic:**

  - Funding for HPC and parallel software development must compete with other important priorities within Europe and the Member States.

  - The continuing economic downturn makes requests for additional funding especially challenging.

  - The economic divide between Europe's richer and less-wealthy nations (i.e., those that contribute most and least to the common funding pool) makes it difficult to apply European HPC resources on a fully egalitarian basis.

- **Technological:**

  - For communications-intensive and data-intensive HPC problems, solution times are better when the problems are run on local HPC resources. This can hamper collaboration by multi-national scientific and engineering research teams. It also limits the number of HPC parallel software codes that can be used effectively today in grid and cloud environments.

- **Practical:**

- Europe is affected by the global shortage of parallel programmers and algorithm developers.

- Europe might not fund the HPC e-Infrastructure resources sufficiently to accommodate the unified European community of computational scientists.

Organizational:

- The elements needed to support a European HPC and parallel software infrastructure and strategy are scattered today as separate "islands," including PRACE, DEISA, EGI, EESI, and others. No central body exists to coordinate these important initiatives to advance a holistic strategy (hardware, software, grid/cloud environments, etc.).

- EESI, in particular, has performed a crucial service by identifying and enlisting the cooperation of a large number of European experts in parallel software development for scientific and industrial research. If this initiative is allowed to disband after it submits its final report to the European Commission later this year, an important organizing effort might need to be undertaken from the start again.

- PRACE and EESI have primarily been focusing on the requirements for advanced scientific research, with less attention to industrial research. A European strategy to use parallel software on HPC systems in order to address Europe's innovation gap cannot succeed without both scientific and engineering/industrial researchers on board. At least one PRACE Tier-0 system, at the HLRS-based Gauss Supercomputing Centre, will host industrial and engineering simulations. European industry needs more access than this to large-scale HPC systems and parallel software that can exploit these systems more effectively than is possible today. And the "islands" of domain-specific engineering researchers need to be better organized, especially to pursue pre-competitive software development that will benefit Europe and the Member States.

- A practical path doesn't yet exist, in Europe or elsewhere, to support parallel software from the development stage to the brink of commercialization. Existing initiatives often provide only a year or two of funding, which is inadequate for developing parallel software that is robust and flexible enough to provide persistent value over the period of 20-30 years that is required. At least five to 10 years of funding eligibility is needed for major codes.

Cultural:

- Europe's strong support for scientific researchers must continue, but industrial/engineering researchers need to receive more support than they do today, including funding to enable key industrial software applications to make more effective use of existing and future HPC systems.

- Nationalism, and the economic divide between the richer and less-wealthy European nations, are additional barriers to unifying Europe's scientific and engineering research communities – although this is far less true than a decade ago.

5.0 RECOMMENDATIONS FOR ADVANCING PARALLEL SOFTWARE: KEY ACTIONS

IDC recommends the following steps and procedures for creating a sustainable, parallel software infrastructure for Europe. These recommendations take the new field research conducted for this study strongly into consideration, along with IDC's own perspective on European and global HPC developments and trends. In all cases, these recommendations are IDC's and not the EU or other individuals.

5.1 Recommendation 1: Create a Governing Body

- **The European Commission should establish a new body to coordinate Europe-wide R&D initiatives for the whole HPC "ecosystem,"** including hardware, parallel software, networking, and other areas, to help ensure that R&D for Europe's HPC ecosystem and HPC-enabled innovation proceeds in a consistent, integrated, and efficient manner. This level of coordination would not affect the autonomy of EU Member States over their national HPC strategies; it would apply only to Europe-wide HPC R&D in which the European Commission plays an important role. This coordination would principally be a mediating function and would not confer sole decision-making authority. The EC body would work closely with the Member States, critical established programs such as PRACE, DEISA, EGI and others, HPC centers, HPC vendors, Europe's scientific and engineering communities, and others to plan and advance R&D initiatives that would include initial development and appropriate maintenance and advancement of the targeted technologies and capabilities. Like the survey respondents, IDC realizes that establishing a new body within the European Commission may not be easy and may encounter resistance; but if the EC wishes to employ the proven value of HPC parallel software, hardware and other resources to help close Europe's innovation gap, we believe a central coordinating body such as this is indispensable.
  - Where parallel software is concerned, the EC body would coordinate the end-to-end software development and funding process, including the bridge period to commercialization.
- Before establishing the coordinating body, **the EC should convene the interested parties** (listed above) **for a two to three day meeting** to hear perspectives, build support, and begin more detailed planning and logistics (it would be presumptuous of IDC to offer detailed recommendations prior to discussions involving the interested parties).
  - As part of this meeting, **convene DEISA and EGI representatives to begin mapping out the communications network needed to support scientific and engineering research access and collaborations across distances** – that is, to extend communications networking in order to support the unification of Europe's scientific and engineering research communities. This would include extending grid/cloud access to small and medium-size enterprises (SMEs) and organizations and nations pursuing small and

medium-size science (SMSs). See next section (Section 4.2.4) regarding the e-Infrastructure requirements for SMEs and SMSs.

- As part of this meeting, **map out software/hardware co-design centers-of-excellence corresponding to the targeted scientific and engineering research domains (see next section of the report)**.

- **The EC should prevent EESI from disbanding** and should attempt to extend its lifetime and the commitment of its members, at least for the important purpose of representing the HPC parallel software perspective on an ongoing basis for Europe and the Member States. The EESI is an invaluable sounding-board asset for a creating the European parallel software development strategy and infrastructure.

- **Representatives of the EC Directorate for Emerging Technologies and Infrastructures should begin to discuss the rationale for the proposed EC coordinating body with appropriate EC officials**, which might include other Directorates whose areas will benefit from the coordinating body's mission and activities, such as:

  - Climate Action
  - Competition
  - Economic and Financial Affairs
  - Energy
  - Enterprise and Industry
  - Research and Innovation

5.2 Recommendation 2: Establish Centers of Excellence for Parallel Software Development

These centers-of-excellence will be needed to help unify Europe's scientific and engineering research communities, and the most practical way to do this is by key domains. Users in the same domain will be interested in the same set of computational challenges and parallel software applications. Each center of excellence would assume primary responsibility for EC-supported parallel software development in its domain, including the creation of plans recommending which parallel codes should be advanced, how and when this would be done, who would lead the effort, and how much funding would be needed.

The centers-of-excellence would be hosted by large (or, in some cases, perhaps medium-size) HPC centers with strong expertise in the scientific or engineering research domain to which the center-of-excellence is primarily dedicated. They would also fit well in the HPC development test-bed recommendation from IDC's 2010 recommendations to the EU. The prestige and benefits of hosting a European center-of-excellence for domain-specific parallel software development would need to be attractive enough to persuade the most qualified HPC centers (or laboratories) to apply for the hosting contracts.

Some of the centers-of-excellence would cover multiple software domains, perhaps as many as 8 to 10. While other centers-of-excellence may only cover a single domain. Each domain would be a very specific domain with both experts in that domains as well as expert programmers.

- The proposed initial domains are weather and climate, clean and sustainable energy, automotive/aerospace/manufacturing, bio-life sciences, particle physics and related fields, and materials science/nanotechnology.
  - For example, CERN should qualify as a center-of-excellence for European particle physics research involving HPC parallel software, hardware, and cloud computing resources.
- In some domains, such as the highly competitive automotive and aerospace research areas, climate, and energy, multiple large HPC centers should be used to improve HPC software.
  - Centers already heavily involved in these domains could agree on which aspect of research each already excels in and which center should take the lead on each aspect on behalf of Europe and the Member States. This would not affect the centers' primary missions and activities that are not associated with EC funding.
  - It is critical that Europe's scientific and research industrial initiatives in the targeted domains are better identified, coordinated, and funded for the purpose of accelerating European innovation and competitiveness.

5.3 Recommendation 3: Create a Parallel Software Clearinghouse

On way to dramatically improve Europe's standing in HPC software is to create a world-class parallel software clearinghouse. European innovation and competitiveness will be more strongly served if HPC's proven value is more pervasively exploited by all organizations including small and medium-size enterprises (SMEs), and by nations and organizations pursuing small and medium-size science (SMSs). Prior IDC studies have shown that the major barriers to HPC adoption, especially for SMEs, include knowing what software to use; the high licensing cost for commercial parallel software; along with access to expertise in using this software on HPC systems.

The parallel software clearinghouse could carry out its functions most efficiently as a single organization via a website that incorporates all the software developers, HPC centers and HPC experts into one central site. It could include major sections like: the ability to purchase any software package (including free software); a section that describes what software to use in which situations (including sub-parts like libraries); a section for buying and selling software IP; a section of obtaining services and training; etc. It should be very easy to use and it could also have access to major systems via direct, grid or cloud models.

For example it could be used to:

- Negotiate a large number of contracts for parallel software use with appropriate ISVs. The ISVs would benefit from access to new customer sets – SMEs and SMSs, many of whom would not have used the ISVs' parallel software before. Pricing for ISV parallel software should be no higher than if purchased directly from the ISVs, and could be lower, based on the potential benefit to the ISVs of attracting new customers.

- Provide access to the ISV parallel software on a pay-as-you-go, software-as-a-service (SaaS) basis. This could become a major advantage if Europe can lead the way to transitioning to a pay-as-you-need-it software pricing model. The EU could work closely with software companies and help them try this approach and help them figure out new business models. It may require significant upfront investments by the EU.

- The clearinghouses would also offer access to compute cycles, human expertise, and other resources at the PRACE/other large HPC centers, as well as all HPC sites that wish to participate in the program. The program should try to get as many sites a possible under the clearinghouse umbrella. This access would be offered in many ways, including new models like on an infrastructure-as-a-service (IaaS) basis and could be delivered locally or via a grid or cloud network.

- Clearinghouse services could even include the full outsourcing of HPC jobs and workloads. It could provide an infrastructure for developing/expanding a new industry – companies providing total outsourcing for HPC workloads.

- Developers of new parallel software would benefit from having an easier path to commercialization – a commercial outlet in which their new parallel software can have an opportunity to succeed or fail on its own merits.

- New HPC users that use the clearinghouse(s) would benefit primarily by gaining access to HPC parallel software, hardware resources, services, and human expertise without the necessity for, and commitment to, capital expenditures and the internal hiring of HPC-proficient IT personnel.

    - IDC believes that this model would substantially address the chief barriers to HPC adoption that SMEs and SMSs have consistently cited in prior worldwide studies.

- Venture capitalists and other investors would benefit from one or multiple new "storefronts" (within the clearinghouse) where they could "shop" for promising IP and parallel software companies to invest in. This investment would in turn contribute to the health of Europe's parallel software development community and overall HPC market – ultimately serving the cause of European innovation and competitiveness.

5.4 Recommendation 4: Establish "Tiger Teams" to Improve HPC Software Across Europe

In order to quickly and fully address the fundamental problem in brining Europe's HPC software forward, IDC recommends that a large number of "Tiger Teams' are created with the goals of:

1. Locating the most important software codes that can be improved.

    - This requires a quick evaluation based on attributes like:

        - How important is the code to Europe?, e.g. how many sites use the code, it is in an important domain, etc.

        - Is the code already being parallelized?

        - Does it appear possible to improve the code?

        - Is the site or code owner willing to work with a Tiger Team?

2. Going to the selected sites and working with them to assessed their codes and make initial improvements.

    - The teams would visit the site and look at the code in detail to determine what can be done quickly (within 2 days effort) and what can be done with major changes

3. Demonstrating the value of these early improvements.

    - The team would then make the "easy" changes to the code and create a demo or test of the new improved code

    - In addition, as part of its primary mission the team will conduct direct training of the code users, e.g., train the on-site personnel in the use of the new scaled-up application

4. Create a plan for the next steps in improving their codes.

    - The team would work closely with the code owner and create an action plan for improving their code. The plan would include all the steps and who the site should work with. The team would make the initial contacts between the code owner and the EU resources.

    - The team would represent the "clearinghouse" and the EU software initiatives as well as the major HPC sites and resources. The team in one view is the sales arm for the software initiative.

    - Return later, if appropriate, to provide further help.

The Tiger Teams would consist of:

1. Two member teams, where one is an expert in the scientific or engineering domain and the other is an expert at programming. The programmer needs to be an expert at looking at a new code and quickly figuring out how to make it scale better. Similar to the role of a benchmark expert within a vendor.

2. A home office support group that would handle: 1) the initial sorting of the potential codes; 2) the making of the appointments and convincing users to use the teams and centers, including reaching agreement for a visit; and 3) the scheduling of the team visits.

The Tiger Teams would typically spend 3 to 5 days at a site, and return once a year if needed.

Tiger Team members would carry prestige as trusted resources acting in the interests of Europe and the Member States. Tiger Team members would be drawn from academia and industry.

The goal of the Tiger Teams is to create the ability to address thousands of HPC application codes, including many of the subparts of the various software programs. As a second phase, the teams can be used to train in thousands of different companies, university departments and groups on how to use more advanced software in their work. The recommendation would be to build to about 100 tiger teams within two years. This would create a capability to address many thousands of codes a year once everything is in place and fully staffed.

These Tiger Teams would greatly help to encourage HPC adoption by SMEs/SMSs in order to help accelerate European innovation and competitiveness. These users are often called the "missing middle," because many of them use parallel software on their desktop computers but have not yet made the move to HPC. The teams would also help medium and large concerns, including those with highly scalable software that need further advancement to reach petascale and exascale levels.(Wikipedia defines a tiger team as "a group of experts assigned to investigate and/or solve technical or systemic problems.")

Potential Tiger Team Structure

The Tiger Teams need to be in a single organization structure for a number of reasons, including:

- They need a single point of coordination to prioritize and for selecting the codes to be address
- To avoid competition between teams
- To see that they are used for improving Europe's software and for addressing other tasks or issues

They should be a separate entity for the teams and home office functions, with a multi-year funding structure. The EU will need to fully fund them for the first few years, but later on they could become funded by the clearing house organization.

The Tiger Teams need to be structured in a way that they can help private companies including ISVs and manufacturing concerns. A number of options exist, including doing the software development work as R&D, publishing the results of the work, etc.

5.5 Recommendation 5: Fully Fund the Recommendations

In the next section the funding levels and options are addressed. A critical requirement for success is that the funds can be located and applied to these areas in a timely manner.

6.0 RECOMMENDATIONS FOR ADVANCING PARALLEL SOFTWARE: FUNDING

6.1 Financial Model and Analysis of Funding Mechanisms

6.1.1 Alternative Scenarios and Recommended Funding Level

IDC recommends that the EU fully fund the software investments recommended in the 2010 EU report, which would cover the centers-of-excellence initiative, and in addition fund two key new initiatives:

1. The software clearinghouse

2. The Tiger Teams

Table 4 (below) shows the required level of HPC software funding from the IDC's 2010 report, *A Strategic Agenda for Supercomputing in Europe: HPC 2020,* conducted on behalf of the DG Information Society and Media of the European Commission (SMART 2009/0055, contract number 2009/S99-142914). Field research conducted for the present study has not substantially altered the base funding recommendations shown in this table. For the development of HPC parallel software, IDC continues to recommend that the EU follows the Full Leadership funding level in the table, requiring $600 million a year increase. In addition, IDC recommends funding the two new initiatives as part of the increased funding.

TABLE 28

Alternative HPC Funding Scenarios: Yearly Funding ADDITIONS By Year Five

(Millions of euro added per year)

	Full Leadership Funding Level	Funding To Reach Major Goals Level	Partial Funding Level	Minimal Increase Funding Level
HPC System Funding Increases (Tier 0 systems)	**150**	120	70	40
HPC System Funding Increases (Tier 1 Systems)	**100**	90	50	10
HPC Development Test-beds (H/W)	**50**	40	25	0
HPC Development Test-beds (People)	**75**	60	35	0
Exascale/Parallel Software Development	**150**	125	55	25
Scientific Talent Magnet Program	**75**	60	25	0
Total Yearly Funding Increase	**600**	495	260	75

Note: These figures include HPC funding paid by the EU, by Member States and contributions by vendors.

Source: IDC, 2010

6.1.2 Funding Requirements

The following elements of the European parallel software infrastructure require major incremental additional funding. The total corresponds to IDC's previous recommended level of €150 million. IDC recommends that this amount needs to be reached within 3 to 4 years and then needs to grow by around 10% each year.

- **EC HPC coordinating body.** Because this body would coordinate the Europe's entire HPC strategy, not just parallel software development, IDC assumes that its operating costs would not be borne solely by the parallel software infrastructure budget. IDC further assumes the following:

    - Existing EC officials with HPC-related responsibilities would, without the need for additional compensation, devote time to the work of the HPC coordinating body.

    - The HPC coordinating body would require only one or two key administrators and perhaps a "team' of 8 to 10 experts, plus 2-3 others to support the group. Hopefully many, if not all of the expert team could be volunteers from HPC centers across Europe.

    - Non-salary expenses would be mainly for meeting expenses and related travel, with meetings held in EC facilities whenever feasible. There would also be costs for creating materials, printing and all of the usual business costs.

    - IDC estimates that overall funding needed by the HPC coordinating body would be in the range of €400,000 to €500,000 per year, and that the parallel software infrastructure's "share" of this total could be less, perhaps less than half.

- **Centers-of-excellence for parallel software development.** IDC estimates that once they have been established, annual funding needed to enable each of the six European centers-of-excellence for parallel software development to perform the functions described in Section 5, including dedicated work on developing or heavily adapting 2-3 key codes per center, would average about €12,500,000 per year, or €75,000,000 annually for all six centers. These costs are already part of the previous recommended spending level.

- **The web-based Clearinghouse for HPC software, hardware access, training and services.** IDC estimates that the EU should plan on heavy funding in the first few years to get the clearinghouse quickly created. Costs include getting all of the many organizations involved, many contracts, payout systems, the web design, etc. For the first 2 to 3 years, the EU should plan on investing close to €10,000,000 a year, and then plan on divesting it to the private sector. The actual operating costs will be much higher, but there are many ways that the clearinghouse would generate some revenues to offset the costs. The goal should be to make it a separate company in the future, so that the EU only has to fund it for a few early years (or to operate it as a revenue generating venture).

- **Tiger teams to quickly improve HPC software across Europe.** IDC estimates that the tiger team program, once it has been established and built up to about 100 teams in its second year, would need funding of about €20,000,000 to €25,000,000 per year. This cost would be on-going, but could ramp down after 4 to 5 years, depending on how much progress is made.

6.1.3 Funding Sources and Mechanisms

The funds should come from a number of sources, including the EU, the Member States, HPC suppliers (vendors), industrial users of HPC-level parallel software, and revenues from the clearinghouse operations. Table 29 provides a possible split-out of funding responsibilities for the European parallel software e-infrastructure.

- We are proposing a 44/44/12 split between the European Commission and the Member States and vendors for the basic funding. But there are variations by category:
  - The clearinghouse should be primarily funded by the EU for the first two years and then either sold or operated as in a revenue generation model (it may take 4 to 5 years to get to a net positive cash flow).
  - The Tiger Teams should be fully funded by the EU in the first year or two and then move towards a split funding model.
- We believe suppliers/users will be willing to contribute more significantly to some elements of the initiative than to others. The figures in Table 29 reflect our assumptions. Supplier/user funds should come from the following sources:
  - A combination of financial and in-kind (staffing) contributions for the parallel software development efforts within the centers-of-excellence.
  - ISV contributions for inclusion of their parallel software codes in the clearinghouse program in the form of advertisements and other fees.
  - Payments to help compensate Tiger Teams, whose responsibilities include making assessments that can lead to the use of vendors' HPC products. Perhaps in the form of a commission payment.

TABLE 29

Potential Funding Sources For The Recommended Funding Level: By Year 3

(Millions of euros per year in new funding)

	EU	Member States	HPC Suppliers	Total
Investments in parallel software development across Europe	4.8	20	8	32.8
EC HPC coordinating body operations	0.2	0	0	0.2
Centers-of-excellence	20	45	10	75
Clearinghouse program	20	0	2	22
Tiger Teams program	20	0	0	20
Total Yearly Funding Increase	65	65	20	150

TABLE 29

Potential Funding Sources For The Recommended Funding Level: By Year 3

(Millions of euros per year in new funding)

	EU	Member States	HPC Suppliers	Total

Note: These figures include HPC funding paid by the EU, by Member States and contributions by vendors.

Source: IDC, 2011

6.2 The Software Development Model

This section covers IDC's recommendations for using a longer period of funding support, and the recommended development models for scientific and engineering software.

6.2.1 Funding Eligibility for Software Development Should Extend to 5-10 Years

Because fundamental adaptations and rewrites of parallel software require considerable time and expense, it makes sense to undertake them only infrequently. An HPC hardware generation typically lasts only 3-4 years, while a software generation must often last 20-30 years.

This means that parallel application codes and other components of the parallel software stack need to be designed with sufficient robustness, flexibility, and other attributes to operate efficiently across multiple generations of HPC hardware systems, including future systems whose architectures may only be vaguely characterized when the software is first developed. Funding needs to be periodically renewable, based on progress toward stated goals.

At least 5-10 years are typically needed to develop production-quality parallel software that has persistent value (lasting 15 to 30 years). Most government funding for software development today, in Europe and elsewhere, is for only 1-2 years. This abbreviated timeframe was a major issue raised by software developers in IDC's government-funded study of the generally successful history of the U.S. Department of Energy's ASCI/ASC program, for example.

6.2.2 A Suggested Model for Scientific Software Development

For the development and advancement of parallel software for use by scientific researchers, IDC recommends the following procedures:

- Convene a body of academic/government experts to create a software roadmap that includes milestones for sustained performance. Take the EESI report due out this fall into careful consideration (see below).
- Use the roadmap as the basis for procurements (RFPs) for targeted areas of parallel software development.
- Strongly encourage the centers-of-excellence to respond to RFPs. But others may bid on equal terms for the parallel software development contracts.

- Use the peer review process to award funding.

- Use the body of experts to monitor progress and make recommendations regarding renewed funding.

- Work closely with other organizations around the world (e.g., IESP), where appropriate, to avoid duplicated efforts and to quickly adopt useful advances.

The European Exascale Software Initiative (EESI) has already organized experts on various areas/aspects of parallel software and in the fall of 2011 will present a roadmap for the development of highly parallel scientific software in Europe. Many of these experts have been drawn from other organizations with strong track records in parallel software development and advocacy. In addition, the EESI actively participates in the International Exascale Software Project (IESP). The EESI and IESP, as their names suggest, focus heavily on highly parallel scientific software and do not have nearly as strong a focus on parallel software codes that are important for industrial/commercial use and competitiveness.

6.2.3 A Separate/Modified Model Is Needed for Parallel Software Used by Industry/Business

Study participants from the world of industry/commerce stressed that they do not want decisions about the future of their key parallel software made exclusively by "bodies of academics." Industrial/commercial respondents should be integrally involved in developing/advancing key parallel software intended for their use, including contributing opinions, time, and money. This category of respondents includes HPC centers involved in industrial/commercial partnerships, as well as HPC hardware, software, and services vendors that typically serve government, academic, and industrial/commercial HPC users.

For development initiatives involving key parallel software intended for their use, these respondents often prefer academic-industrial partnerships. Examples of successful partnerships of this kind include BMW and the University of Munich, Porsche, T-Systems, and the University of Stuttgart/HLRS, and others.

For the development of pre-competitive parallel software whose use will be open to anyone, the industrial/commercial respondents would expect government funding to cover all or nearly all of the development cost. The respondents would expect to contribute expertise and, as appropriate, time on their HPC resources.

For parts of the development that is not pre-competitive, the investing companies/organizations expect sole use of the software for at least 12-24 months.

- If only one private company or party will benefit, the government (e.g., European Commission) and the company/party should co-fund the parallel software development.

- If multiple companies/parties will each gain unique IP benefits, a value-based formula should be created to determine the co-funding portion assigned to each company/party. Each company/party might pay a different amount.

Procurements and RFPs for the development and advancement of parallel software for use by industrial and engineering researchers should be disseminated to the recommended centers-of-excellence for industrial/engineering domains. Other parties should be able to compete on equal terms for these contracts. In cases where other parties win the contracts, IDC expects that they will often want to use the resources of the centers-of-excellence, where they can obtain access to large-scale supercomputers and expertise in using them.

6.2.4 Summary of the Suggested Development Model Ideas for ***Scientific Parallel Software***

- Target areas where Europe can lead
- Academic/government experts create a software roadmap
- RFPs are based on this roadmap
- Peer review determines who gets funded
- At least 5-10 years of funding, with periodic renewal eligibility
- Academic/government experts determine renewal
- 100% government (EC) funding for pre-competitive development, co-funding for any competitive portions

6.2.5 Summary of the Suggested Ideas for a Development Model for ***Industrial/Commercial Parallel Software***

- Target areas where Europe can lead
- Academic-industry experts create goals, timeline
- RFPs are based on the goals, timeline
- Fixed-cost contracts are awarded, with fixed deadlines (usually non-renewable)
- Proportion funded by government (EC) varies by contract and IP terms

6.3 Governance Model Ideas for the Software Infrastructure

6.3.1 Establish a Coordinating Body within the European Commission

The new EC body would coordinate Europe-wide R&D initiatives for the whole HPC "ecosystem," including hardware, software, networking, and other areas, to help ensure that R&D for Europe's HPC ecosystem proceeds in a consistent, integrated, efficient manner. This level of coordination would not affect the autonomy of EU Member States over their national HPC strategies; it would apply only to Europe-wide HPC R&D in which the European Commission plays an important role. The EC body would work with the Member States, HPC centers, HPC vendors, Europe's scientific and engineering communities, and others to plan and advance R&D initiatives that would include initial development and appropriate maintenance and advancement of the targeted technologies and capabilities.

Where parallel software is concerned, the EC body would coordinate the end-to-end software development and funding process, including the bridge period to commercialization. The recommended clearinghouse program would provide a commercial outlet for existing and new software developers as part of the European parallel software e-infrastructure.

In IDC's 2010 study for the European Commission, *A Strategic Agenda for Supercomputing in Europe: HPC 2020*, the European HPC community, including government officials overseeing HPC programs, generally agreed that the PRACE program could serve as the coordinating body for a European HPC strategy, but only if PRACE's mission were substantially expanded to embrace the whole HPC ecosystem, including access to hardware systems for industry, software development, and other areas. When directly presented with the challenges of developing/advancing highly parallel software benefiting all of Europe, respondents in the present study thought, and IDC agrees, that the appropriate path would be to establish a coordinating body within the European Commission that worked closely with crucially important organizations including PRACE as the lead hardware organization, EESI as a leading software development advisory group, DEISA and EGI as the lead networking organizations, and others. Among the goals would be co-development of hardware and software plans, to the extent feasible, and coordination of software and hardware initiatives to maximize the value of the large investments made in PRACE Tier-0 supercomputers.

The EC coordinating body might report to the Directorate for Emerging Technologies and Infrastructures and should establish appropriate relationships with other Directorates whose areas will benefit from the coordinating body's mission and activities, such as:

- Climate Action
- Competition
- Economic and Financial Affairs
- Energy
- Enterprise and Industry
- Research and Innovation

6.4 Actions to be Taken by the European Commission and the Member States

6.4.1 Actions Needed by the European Commission

Determine the receptivity within the Commission to establishing the European parallel software e-infrastructure, including a new body to coordinate the European HPC strategy, including the parallel software e-infrastructure. (The field research conducted for this study produced a favorable consensus for an EC-based coordinating body.) And determine which internal funding mechanism(s) are available and begin the process needed to secure EC funding.

- **Pros:** the e-infrastructure is needed to maximize the ROI on recently deployed PRACE Tier-0 supercomputers, and to enable HPC to deliver its proven value as an accelerator for scientific and industrial innovation on behalf of Europe and the Member States. Because parallel software development is a key to leadership in HPC and in the scientific and engineering research domains that depend on HPC, and no other nation or region has significantly committed itself to tackling the parallel software challenge, this constitutes a unique opportunity for Europe to pursue and seize global leadership in the targeted domains that are of great scientific and economic importance.

- **Cons:** This will need to compete with other European priorities during a difficult economic period.
- **Cost:** IDC recommends new funding for the entire program that builds to a level of €150 million/year in the fifth year of the initiative.
- **Duration:** Ongoing, in order for Europe to remain globally competitive in HPC-dependent scientific and engineering research.
- **Likelihood of success:** IDC believes that if Europe funds and implements the parallel software e-infrastructure as recommended, Europe's chances of seizing global leadership in the targeted domains are excellent.
- **Potential benefits:** Close Europe's current innovation gap vis-à-vis other leading regions and nations of the world. Provide European scientific and engineering researchers with the powerful computational resources they need to out-innovate and out-compete the rest of the world.

Prevent the EESI participants from disbanding after they submit their final report in fall 2011. Let the EESI director (Jean-Yves Berthou, EDF) know that this group will be needed as a sounding board for the parallel development e-infrastructure planning and implementation process.

- **Pros:** Much effort went into identifying and organizing the experts who have been participating in the EESI initiative. They are needed as sounding boards for planning the parallel software e-infrastructure. Enlisting their help would save the time and expense needed to organize this group again.
- **Cons:** EESI has focused primarily on the parallel software requirements for advanced scientific research and only to a lesser extent on the requirements for advanced engineering research.
- **Cost:** Essentially none. The European Commission only needs to ask the EESI Director not to disband the group after they submit their final report in fall 2011.
- **Duration:** A quick, single communication.
- **Likelihood of success:** Very good
- **Potential benefits:** The EESI would be very useful as a sounding board for planning the parallel software e-infrastructure.

Convene a 2-3 day meeting, to take place by December 2011, where the interested parties will discuss the findings and recommendations of this study and begin organizing the parallel software e-infrastructure program, as recommended. This meeting would need a detailed agenda.

- **Pros:** A meeting like this is needed so the interested parties can gain a better understanding of the recommended plan, express their interest in participating in the plan, and begin to map out roles and responsibilities.
- **Cons:** The meeting has to be organized and promoted several months in advance.
- **Cost:** IDC estimates that costs for a meeting of 50-75 key people (5-6 of them EC officials) will be less than €10,000 if the meeting is held at an EC facility, or 2-3 times that amount if it is held elsewhere.
- **Duration:** 2-3 days
- **Likelihood of success:** Good
- **Potential benefits:** Gain understanding and support for the recommended parallel software e-infrastructure; identify objections needing to be addressed.

Establish the European HPC/parallel software coordinating body.

- **Pros:** As the field research said, a central coordinating body is needed to ensure European parallel software development is well coordinated with European planning for hardware, networking, and other HPC-related resources. No existing body has all of the competencies to carry this out.
- **Cons:** Establishing a new body within the European Commission may be difficult and may encounter resistance.
- **Cost:** See Table 29.
- **Duration:** Ongoing
- **Likelihood of success:** Good
- **Potential benefits:** A better return on the investments Europe has already made and plans to make in large-scale HPC systems. Better coordination of the effort to close Europe's innovation gap vis-à-vis other regions and nations.

Establish the domain-specific, European centers-of-excellence for parallel software development.

- **Pros:** The centers-of-excellence will be needed to help unify Europe's scientific and engineering research communities, and the most practical way to do this is by domain. Users in the same domain will be interested in the same set of computational challenges and parallel software applications.
- **Cons:** Requires creating new organizations, even though they will be hosted by existing HPC centers.
- **Cost:** See Table 29.
- **Duration:** Ongoing

- **Likelihood of success:** Good to excellent, depending on the challenges associated with advancing parallel software capabilities in each domain.

- **Potential benefits:** Parallel software development is the key to HPC leadership, and to leadership in the scientific and engineering research domains that depend on HPC. Europe has a unique opportunity to seize HPC leadership and help close Europe's current innovation gap.

Establish the web-based clearinghouse program for parallel software and related HPC resources and services.

- **Pros:** European innovation and competitiveness will be more strongly served if HPC's proven value is more pervasively exploited by all sizes of companies and organizations including small and medium-size enterprises (SMEs), and by nations and organizations pursuing small and medium-size science (SMSs).

- **Cons:** Requires establishing a program and recruiting a strong leader for it.

- **Cost:** See Table 29.

- **Duration:** Ongoing

- **Likelihood of success:** IDC expects the clearinghouse to achieve successes for the most part, but there are also likely to be a few partial successes and outright failures.

- **Potential benefits:** Establish a better commercialization path for European parallel software developers. Accelerate European innovation and competitiveness.

Establish the Tiger Teams program and recruit the first (10-20) tiger teams.

- **Pros:** Encourage HPC adoption by European companies, universities and other organizations by helping them directly to improve their software. This is the fastest path for moving the software across Europe forward to a more scalable world.

- **Cons:** Requires starting the program and an increasing effort to recruit tiger team experts.

- **Cost:** See Table 29.

- **Duration:** A two-year ramp-up to about 100 tiger team. Ongoing at that level afterward.

- **Likelihood of success:** Very good.

- **Potential benefits:** More pervasive European use of HPC, resulting in increased innovation and competitiveness.

6.4.2 Actions Needed by the Member States

Provide the level of support for the recommended European parallel software e-infrastructure, including the recommended new coordinating body within the European Commission and the other elements of the initiative.

- **Pros:** No Member State acting alone can afford to compete in HPC or parallel software development with leading nations such as the U.S., China, and Japan. The parallel software e-infrastructure is a natural extension of existing European HPC collaborations such as PRACE, DEISA, and others.
- **Cons:** This initiative will need to compete with other funding priorities within the Member States during a difficult economic period.
- **Cost:** See Table 29.
- **Duration:** Ongoing
- **Likelihood of success:** Very good
- **Potential benefits:** The program aims to enable the Member States' scientific and engineering research communities to become even more innovative and productive, by providing them with access to parallel software that can more effectively exploit HPC systems.

Agree to participate in the 2-3 day meeting the European Commission will convene by December 2011.

- **Pros:** A meeting like this is needed so the interested parties can gain a better understanding of the recommended plan, express their interest in participating in the plan, and begin to map out roles and responsibilities.
- **Cons:** The meeting will take time away from participants' jobs.
- **Cost:** The main direct costs would be travel expenses to attend the meeting.
- **Duration:** 2-3 days
- **Likelihood of success:** Good
- **Potential benefits:** Gain understanding and support for the recommended parallel software e-infrastructure; identify objections needing to be addressed.

6.5 Potential Risks If No Action Is Taken

IDC believes that the potential consequences of not implementing the European parallel software e-infrastructure are as follows:

- **Another Region or Nation Could Seize This Opportunity and Further Widen Europe's Innovation Gap.** Other areas of the world, particularly the United States, have strong credential in HPC-level parallel software development. To date, no nation or region has committed itself to seriously tackling the parallel software challenge that, in the view of IDC and many of the external experts we talk with, is the key to future HPC leadership. Given the politico-economic climate in the U.S., IDC believes that Europe has a limited window of opportunity to fund and implement a program to advance and support parallel software development. At some point, the U.S. or another capable nation is likely to take up this banner and run with it.

- **Europe Could Become Inferior To The U.S. and Asia in Scientific and Engineering Research.** Failure to expand the EU HPC strategy to keep pace with the rest of the world could cause Europe to become a second-class region for science and advanced engineering. It could also leave each EU member state to compete on its own, resulting in less ambitious, more fragmented scientific and engineering initiatives and outcomes. Smaller EU Member States would likely suffer most from this scenario.

- **Europe and The EU Member States Could Lose Industrial Competitiveness and Jobs.** The existing EU HPC strategy already lags the U.S. and Japan in providing industry with access to world-class HPC resources, including parallel software. If nothing is done to remedy this situation, Europe and the Member States could fall seriously behind these and other nations in industrial innovation and economic competitiveness. In a nutshell, other nations and global regions could have a major advantage in bringing more advanced, higher-quality products to market in shorter timeframes than European companies. And because HPC is a major weapon against countries with lower labor costs, failure to provide adequate access could result in the permanent loss of European industrial firms and jobs.

- **Europe and The EU Member States Could Experience An Escalating Brain Drain To The U.S. And Asia, Along With Greater Difficulty In Attracting Talented Scientists And Engineers.** The Europe-to-U.S. brain drain is already occurring, including scientists relocating to work at U.S. national laboratories with the best HPC resources. This brain drain would likely escalate if Europe failed to keep its HPC resources on a par with those of the U.S. and Asia. If Europe's current deep pool of scientific and engineering talent is allowed to drain away, it might take a decade or more to restore it even with the best of efforts and funding. It makes far more sense to address the brain drain now.

- **Europe and The Member States Could Become Increasingly Reliant On The U.S. and Asia For Scientific, Industrial, and Technological Advances.** If Europe fails to keep pace in HPC with the U.S., Japan, and others, Europe might be forced to import scientific, industrial, and technological advances from other areas of the world – without having much except money and political favors to offer in exchange.

- **The Smaller and Less Affluent EU Member States Could Lose The Ability To Access and Benefit From Large HPC Systems. This Could Widen The Digital Divide In Europe, To The Detriment Of The Smaller Countries.** If Europe failed to keep pace in HPC with the U.S., Japan, and others, the smaller and less affluent EU Member States would likely suffer the most. Unlike Europe's wealthier states, the smaller countries typically cannot afford to fund world-class HPC systems on their own and rely to a greater extent on access to tier 0 EU HPC systems. It's easy to see how failure to expand the EU HPC strategy could widen the digital divide between richer and less-affluent Member States.

- **Europe's Existing Strengths in Hardware, Software and Other HPC-Related Technologies Could Diminish or Disappear From Lack Of Advancement.** It almost goes without saying that Europe's existing world-class skills in HPC-related technologies would also erode without continuing access to not only to world-class HPC systems, but to the world-class parallel software needed to exploit these systems efficiently.

6.6 Software Areas to Focus Investments On

IDC recommends that the EU selects HPC software that matches these key criteria:

1. The application code or supporting software is important in some way to Europe. For example:

   - Its use in one of the domain areas of focus.
   - There are many users of the software package or tool set.
   - The code could provide a major breakthrough, scientific advancement, innovation or economic return.
   - Improving subparts of codes, e.g., libraries, drivers and various utilities if used broadly can have a major impact on a larger set of application codes.
   - Improving the efficiency of the scaling of codes that have fees related to CPU hours can save users a large amount in software license fees, but requires the EU to work closely with ISVs.

2. The software code should be evaluated based on:

   - The code has potential to scale.
   - No one is currently investing in making it scale better.
   - The path to improving the code makes technical sense and a plan can be created to improve its scaling.

3. Codes that should be avoided include:

   - Software that has major sections that are singled thread or scalar in nature, and it doesn't appear to be possible to change its underlying nature.
   - Codes that are only used by a few users (unless they are for very critical and frequently run problems).
   - Codes that already scale well and don't have major issues moving into a petascale era.

4. The Tiger Teams and the Centers of Excellence need to focus on codes that can show the most success, in the shortest amount of time:

   - Applications that can be scaled with minimum changes should be addressed first.
   - Applications that require a moderate level of effort should be more carefully evaluated and weighted against the potential returns.
   - Applications that require extensive redesign need to have the most extensive analysis of the work required compared to the potential returns. Unfortunately many of the most useful codes to be addressed require very

extensive redesigns. But this also creates a major opportunity for the code's future once it is redesigned.

5. In addition to improving HPC application parallel scaling, some codes can be improved in other ways, and the program should also address them:

   - Codes that scale to only 1 core today could show great results if they could be made to scale 4-way.

   - Redesigning codes so that large problem sizes can be addressed, e.g. larger memory sizes.

   - Improving I/O to storage and between nodes can often lead to strong performance improvements.

Table 30 shows a sample rating of a number of important application codes based on the need for scaling improvement (based on scaling of core counts). Table 31 shows an example of the importance of the codes based on the number of users. The EU program recommendations are to create tables like these for hundreds of application codes in Europe and use them to select the codes that both need the most help and that have the most value to Europe.

TABLE 30

Example Ranking of Key Codes by: THE NEED FOR IMPROVING SCALING

Weak Scaling Codes:		Moderate Scaling Codes (Great Candidates):		Strong Scaling Codes (May already be okay):	
Code Name	Typical Job Scaling In Cores	Code Name	Typical Job Scaling In Cores	Code Name	Typical Job Scaling In Cores
Deposit	1	ANSYS	16	DLPOLY	1,000
Code_Aster, mechanical solver	4	CFD (STAR CD)	16	NAMD	1,000
NASTRAN	4	Fluent	16	ParaFEM	1,000
.	.	INTEGRATED ENGINEERING CODE SUITE	16	Quantum Espresso	1,000
.	.	Numpy/Scipy	16	GENE	1,024
.	.	Python	16	GPAW	1,024
.	.	Pamcrash (ESI)	24	GROMACS	1,024
.	.	AIMPRO	32	GYSELA	1,024
.	.	ANSYS	32	INTEGRATED PHYSICS CODE SUITE	1,024

TABLE 30

Example Ranking of Key Codes by: THE NEED FOR IMPROVING SCALING

Weak Scaling Codes:		Moderate Scaling Codes (Great Candidates):		Strong Scaling Codes (May already be okay):	
.	.	FLOW-3D	32	IFS	1,536
.	.	IPSLCM	32	IFS	1,536
.	.	ParaGauss	32	TRIO-U	3,000
.	.	VASP	32	BQCD	4,096
.	.	WRF	48	NAMD	4,096
.	.	GPREMS	60	POEM	5,000
.	.	CASTEP	64	STAMP	5,000
.	.	ECHAM - HAMMOZ	64	GYSELA	8,000
.	.	GROMACS	64	YALES2	8,000
.	.	ONETEP	64	CP2K / CMPD	8,160
.	.	Openfoam	64	S12MPI	8,192
.	.	Overture	64	Massively Parallel Quantum Computer Simulator	262,144
.	.	VASP	64	.	.
.	.	Octopus	80	.	.
.	.	Molecular Dynamics (AMBER)	96	.	.
.	.	GROMACS	100	.	.
.	.	PARADIS	100	.	.
.	.	Code_Saturne	128	.	.
.	.	Elmer	128	.	.
.	.	Eurad - IM	128	.	.
.	.	GROMACS	128	.	.
.	.	INHOUSEA	128	.	.
.	.	INHOUSEB	128	.	.
.	.	Matlab	128	.	.
.	.	Neptune_CFD	128	.	.

TABLE 30

Example Ranking of Key Codes by: THE NEED FOR IMPROVING SCALING

Weak Scaling Codes:		Moderate Scaling Codes (Great Candidates):		Strong Scaling Codes (May already be okay):	
.	.	OpenFOAM (OpenCFD)	128	.	.
.	.	VASP	128	.	.
.	.	CP2K	256	.	.
.	.	JOREK	256	.	.
.	.	Telemac	256	.	.
.	.	COSMO / ECHAM	384	.	.
.	.	ABINIT	500	.	.
.	.	GaitSym	500	.	.
.	.	TRIPOLI	500	.	.
.	.	Gromacs	512	.	.
.	.	INCA	512	.	.
.	.	NAMD	512	.	.
.	.	R SPRINT	512	.	.

Source: IDC, 2011

TABLE 31

Example Ranking of Key Codes by: IMPORTANCE BASED ON THE NUMBER OF USERS

Codes With Only a Few Users:	Codes With 50 to 100's of Users (Great Candidates):	Codes With Many Users , Over 1,000 (Great Candidates):
BQCD	AIMPRO	ABINIT
Code_Saturne	Eurad - IM	ANSYS
Deposit	GENE	CASTEP
GPREMS	IFS	Code_Aster, mechanical solver
INHOUSEA	R SPRINT	CP2K / CMPD
INHOUSEB	TRIPOLI	DLPOLY
Massively Parallel Quantum Computer Simulator	Unified Model	FLOW-3D

TABLE 31

Example Ranking of Key Codes by: IMPORTANCE BASED ON THE NUMBER OF USERS

Codes With Only a Few Users:	Codes With 50 to 100's of Users (Great Candidates):	Codes With Many Users , Over 1,000 (Great Candidates):
PARADIS	COSMO / ECHAM	GROMACS
POEM	Elmer	GROMACS
S12MPI	NEMO	Matlab
STAMP	Octopus	NAMD
Codes with 10 to 50 users:	ONETEP	Openfoam
ECHAM - HAMMOZ	ParaFEM	Quantum Espresso
GaitSym	Telemac	ANSYS
INCA	GROMACS	CFD (STAR CD)
Neptune_CFD	IFS	Fluent
ParaGauss	IPSLCM	Molecular Dynamics (AMBER)
Overture	NAMD	NASTRAN
TRIO-U	NAMD	Python
.	OpenFOAM (OpenCFD)	Vampir
.	Pamcrash (ESI)	VASP
.	.	WRF

Source: IDC, 2011

7.0 SOME OUT-OF-THE-BOX IDEAS / RECOMMENDATIONS

The previous recommendations form the base IDC recommendations for advancing Europe's parallel software. This section adds ideas that are a bit different, but could provide a major improvement.

7.1 Outsource the Tiger Teams to an Existing Company like TeraTec or NAG

By outsourcing the Tiger team organization to a single concern, the EU can obtain the full benefits of the teams, without having to create a new organization or new agency. The EU will need to carefully monitor and clearly set the tasks for the organization.

7.2 Setup a Structure Where ISVs Can Help Fund Some of These Initiatives

Private companies are willing to fund various activities that help to support their product's future, but they want to first see the business and see the revenues before committing money.

The process could work as follows:

- The EU invests in R&D along side the ISVs to help make their scale better
- The ISVs then start selling these new codes to see if they are welcomed in the marketplace
- If the code sells well, the ISV repays the EU over a period of time, e.g. over 5 years, based on the sales success

8.0 CONCLUSIONS

8.1 Key Findings

The field research IDC conducted for this study produced the following key findings:

- Europe has some notably successful scientific and engineering software firms and is strong in important areas of parallel software development.

- Europe and the world have been overly focused on funding parallel hardware to the detriment of parallel software, and on "big science" to the near-exclusion of industry.

- Funding today is too short-term and the infrastructure is inadequate for developing robust, production-quality software.

- The vast majority (83%) of the most important parallel software applications in use at the surveyed European HPC sites were created in Europe. Intellectual property rights for a substantial majority of the sites' most important application codes (66%) were exclusively owned by European organizations. This finding runs counter to the frequent assumption within the worldwide HPC community that European HPC sites rely primarily on the U.S. for their most important software applications.

- Only a few of the site's single most important application codes scale beyond 1,000-2,000 processor cores today, but 77% of the sites thought they could exploit tens of thousands of cores and 36% said they believed they could scale to hundreds of thousands of cores. The largest HPC systems today contain more than 200,000 cores, and million-core systems are expected to begin arriving before the end of this decade.

- Half of the sites' first top codes were 10 or more years old, and 14% were more than 20 years old. Although most (71%) of the first top codes have undergone a technical update within the past year, and nearly all (97%) have been updated within the past five years, this does not imply that the codes have been substantially adapted or fundamentally rewritten to efficiently exploit today's and tomorrow's large-scale high performance computing (HPC) hardware systems.

- Nearly all of the surveyed sites said they would be willing to invest in improving the parallel software application codes that are most important for their users, and to partner with outside organizations, including the European Commission, in these development/advancement efforts.

- The most important needs for improving the first top codes were adequate internal expertise and financial investment.

- An HPC parallel software generation must be robust and flexible enough to be useful for 20-30 years and over multiple generations of HPC hardware systems. At least 5-10 years of funding are needed to develop production-quality parallel software with this persistent value.

- The development models for scientific and industrial/commercial parallel software need to be different.

- Parallel software development for scientific and industrial/commercial use is important enough to establish a coordinating body within the European Commission. Working closely with PRACE, EESI, and others, this body should coordinate R&D for the European HPC ecosystem.

8.2 Benefits of an Expanded Parallel Software Strategy

The main benefits of adopting a European strategy for the development/advancement of server-based parallel software for scientific and engineering research are as follows:

- Parallel software advances will be more important for future HPC leadership than progress in designing parallel hardware systems.
- A parallel software development strategy gives Europe the opportunity to create the world's best computational resources to enable scientific and industrial/commercial leadership.
- Europe has already made substantial investments in HPC hardware systems.

8.3 IDC Recommendations

- Use HPC Parallel Software Development to Help Close Europe's Innovation Gap
  - Europe has a limited window-of-opportunity to become a global innovation leader in targeted domains of scientific and engineering research
- Establish an e-Infrastructure for Advancing Parallel Software
  - Create a new EC body to coordinate the holistic parallel software and holistic HPC strategy for Europe
  - Establish European centers-of-excellence for parallel software development
  - Create a European Web-based parallel software clearinghouse
  - Establish "Tiger Teams" to improve HPC access across Europe
  - Establish appropriate parallel software financing, development, and governance models for the e-Infrastructure

9.0 APPENDIX

9.1 Survey Guide

IDC is conducting a study on behalf of the European Union of the requirements, needs, and experiences of HPC users related to parallel software codes and applications in support of making the use of HPC stronger, more productive and more innovative across Europe.

All of your responses to the questions will be kept confidential and only summary averages will be reported.

By HPC we mean technical compute servers used on computationally intensive or data intensive problems as used in many industries/sectors including automotive, government labs, finance, oil/gas, universities, the life sciences, etc.

Name: ___________________

Email: _____________________________

Phone Number: _____________________

Fax Number: _______________________

Company: _____________________________

QUESTIONS ABOUT YOUR MOST IMPORTANT HPC SOFTWARE CODES

NOTE: If you have more than 2 key HPC application codes, please complete an additional survey form.

1) What are the names of your top 2 most important HPC Software Application Packages / codes?
   a. _______________________________
   b. _______________________________
2) Where does each of these packages or codes originally come from?
   a. Package "a": _____________ a. Government
      b. University
      c. A Company or ISV
      d. Open source
      e. Your own organization or department
   b. Package "b": _____________ a. Government
      b. University
      c. A Company or ISV
      d. Open source
      e. Your own organization or department
3) Who was the Parallel Software Application Creator -- Company or Organization name?
   a. Package "a": _______________________________
   b. Package "b": _______________________________
4) Where is the provider located -- Organization or Company Location -- Country?
   a. Package "a": _______________________________
   b. Package "b": _______________________________
5) What are the Primary Industries or sectors that the Software Application Package Codes are Used In?
   a. Package "a": _____________ , _____________ , _____________
   b. Package "b": _____________ , _____________ , _____________
6) What are the Primary Countries the Software Application Package Codes are Used In?
   a. Package "a": _____________ , _____________ , _____________
   b. Package "b": _____________ , _____________ , _____________
7) What type of users mostly use each package?
   a. Package "a": _____________ (Government, Academic or Industry)
   b. Package "b": _____________ (Government, Academic or Industry)

QUESTIONS ABOUT THE SCALING OF YOUR MOST IMPORTANT APPLICATION PACKAGES OR CODES

8) Scaling: For A Single Job -- What Is The **Maximum Number** Of CPU Sockets and Cores that a single job can use?
   a. Package "a": # of Sockets _____________ , # of Cores ________________
   b. Package "b": # of Sockets _____________ , # of Cores ________________

9) Scaling: For A Typical single Job -- What Is **The Typical Number** Of CPU Sockets and cores Used?
   a. Package "a": # of Sockets ______________ , # of Cores __________________
   b. Package "b": # of Sockets ______________ , # of Cores __________________
10) Scaling: Do you or your provider **technically know how to expand this code** to run on **100's** of CPU sockets (or 1,000 cores)?
   a. Package "a": ______________ (Yes or No)
   b. Package "b": ______________ (Yes or No)
11) Scaling: Do you or your provider **technically know how to expand this code** to run on **1,000's** of CPU sockets (or 10,000 cores)?
   a. Package "a": ______________ (Yes or No)
   b. Package "b": ______________ (Yes or No)
12) Scaling: Do you or your provider **technically know how to expand this code** to run on **10,000** Plus CPU sockets (or 100,000 plus cores)?
   a. Package "a": ______________ (Yes or No)
   b. Package "b": ______________ (Yes or No)
13) Who are the primary Competitors for each Software Application Package?
   a. Package "a": ______________ , ______________ , ______________
   b. Package "b": ______________ , ______________ , ______________
14) How large is the organization that provides, supports or sells each application package -- in number of employees?
   Package "a": ______________
   - a. under 10 employees
   - b. 10 to 25 employees
   - c. 25 to 50 employees
   - d. 50 to 100 employees
   - e. over 100 employees

   Package "b": ______________
   - a. under 10 employees
   - b. 10 to 25 employees
   - c. 25 to 50 employees
   - d. 50 to 100 employees
   - e. over 100 employees
15) Number of Licenses or Clients Worldwide for each package:
   a. Package "a": ______________
   - a. under 10 licenses /clients
   - b. 10 to 25 licenses /clients
   - c. 25 to 50 licenses /clients
   - d. 50 to 100 licenses /clients
   - e. 100 to 1,000 licenses /clients
   - f. over 1,000 licenses /clients

   b. Package "b": ______________
   - a. under 10 licenses /clients
   - b. 10 to 25 licenses /clients
   - c. 25 to 50 licenses /clients
   - d. 50 to 100 licenses /clients
   - e. 100 to 1,000 licenses /clients
   - f. over 1,000 licenses /clients
16) What Is the Age Of Each Software Application Package?
   a. Package "a": ______________
   - a. under 1 year
   - b. 2 to 5 years
   - c. 5 to 10 years
   - d. 10 to 20 years
   - e. over 20 years

   b. Package "b": ______________
   - a. under 1 year
   - b. 2 to 5 years
   - c. 5 to 10 years
   - d. 10 to 20 years
   - e. over 20 years
17) How long ago was the Last major Technology Update Made to Each Software Application Package?
   a. Package "a": ______________
   - a. under 1 year
   - b. 2 to 5 years
   - c. 5 to 10 years
   - d. 10 to 20 years

e. over 20 years

b. Package "b": ______________ a. under 1 year
b. 2 to 5 years
c. 5 to 10 years
d. 10 to 20 years
e. over 20 years

18) Who Owns the Intellectual Property Rights for each of These Codes?
   a. Package "a": ______________
   b. Package "b": ______________

ABILITY AND WILLINGNESS TO IMPROVE THESE HPC CODES

19) Overall for these codes, are you or your software code provider willing to take these actions to improve these codes?
   a. **Willing to INVEST** in Improving these Software Application Package Codes):
   Package “a” ___ Yes/No
   Package “b” ___Yes/No
   b. **Willingness to Develop New Partnerships** to Improve these Software Application Package Codes) :
   Package “a” ___ Yes/No
   Package “b” ___Yes/No
   c. **Willingness to Work with the EU** to Improve these Software Application Package Codes) :
   Package “a” ___ Yes/No
   Package “b” ___Yes/No
20) What type of **partners** Would Be Most Useful in Helping to Improve these codes?
   Package “a” a __________ b______________ c______________ d_____________
   Package “b” a __________ b______________ c______________ d_____________
21) What are the Primary Needs for Improving This Code -- Indicate all that apply?
   a. Package "a": ______________ a=money / investments
   b=business case-many more customers
   c=Partnerships-share costs & risks
   d=Internal People or experts
   e=A whole new approach to their code
   f=external technical expertise
   b. Package "b": ______________ a=money / investments
   b=business case-many more customers
   c=Partnerships-share costs & risks
   d=Internal People or experts
   e=A whole new approach to their code
   f=external technical expertise
22) Describe the programming environment used in developing the codes (compilers, special supporting software and/or hardware, etc.)?:
   a. Package "a": ____
   b. Package "b": ____
23) What Else Would it Take to Help Advance these Codes?
   Package "a":
   Package "b":

9.2 EU-Based HPC Parallel Software Vendors

HPC software is created by a variety of sources, including universities and government organizations as well as commercial enterprises. Europe has considerable strengths in each of these domains and across many organizations. For reference, a partial list of EU-based software vendors is given here:

Allinea Software is a leading supplier of tools for multicore and high performance computing (HPC). Its Distributed Debugging Tool (DDT) and Optimization and Profiling Tool (OPT) are used by universities, government research institutes, and commercial organizations. With new product features aimed at novel computing architectures, and the announcement of new GPGPU features, Allinea is a leader in the tool arena.

Bull (les Clayes Sous Bois, France) is one of the leading European IT companies. Bull has a large team of HPC experts and designs and delivers integrated solutions, from departmental clusters to world-class supercomputers. Bull has a growing number of HPC customers at research centers, universities, and industrial companies.

Since 1985, the **CADFEM** (Grafing, Germany) has been a leader in ANSYS and competence in CAE in Germany, Austria, and Switzerland.

Cluster Resources (Cambridge, U.K.) is a leader in unified intelligent automation software for high performance computing (HPC) datacenters and cloud computing environments. With more than a decade of experience deploying adaptive operating environments powered by Moab in the world's most advanced datacenters, Cluster Resources delivers software and services that enable organizations to obtain a unified perspective of their resources and optimize service levels through intelligent policy-based governance.

ClusterVision (Amsterdam, Netherlands) is a specialist in the design, implementation, and support of small- and large-scale computer clusters. ClusterVision's team of experts has designed and built some of the largest and most complex computational, storage and database clusters in Europe. With a background in applied scientific research and practical experience with a wide range of HPC technologies, the team provides tailor-made solutions. ClusterVision has offices in most major European countries.

ESI (Paris, France) is a world leading software vendor for digital simulation of prototypes and manufacturing process engineering in applied mechanics. The key to ESI's success is the use of realistic material physics, providing "as good as real" virtual solutions, in order to replace the lengthy trial and error processes on real prototypes.

Eurotech (Amaro, Italy) develops and produces HPC systems and miniaturized computers for pervasive computing applications (NanoPCs) in the defense, security, transport, aerospace, industrial, medical, and research sectors.

GNS Systems (Braunschweig, Germany) offers information technology services for product development and engineering, including the planning, implementation, and operation of complex systems and applications infrastructures.

Go Virtual (Askim, Sweden) provides HP HPC-systems, servers, and workstations as well as software from partners including Metacomp Technologies (CFD++ and CAA++), ThermoAnalytics (RadTherm, WinTherm),and Intelligent Light (FieldView).

Gridcore is a privately held company based in Göteborg, Sweden. Gridcore is a worldwide HPC solutions provider that has formed long-term partnerships with companies such as IBM, ANSYS, and Microsoft. Gridcore also owns and operates Gompute, an HPC-on-demand service that operates worldwide.

Kerlabs (Rennes, France) is a spin-off of INRIA, created in 2006 to develop the Kerrighed.One operating system for clusters.

LSI (Bracknell, Berkshire, U.K.) is a leading provider of innovative storage, silicon, systems, and software technologies for some of the world's leading OEMs in the storage and networking markets.

MEGWARE Computer GmbH (Chemnitz-Roehrsdorf, Germany) distributes high-performance computing systems and IT-equipment to industry, banks, trade, universities, schools, and public facilities. The company's solutions are customized to the specific requirements of a broad range of customers.

Mercury Visualization Sciences Group (Mérignac, France) is a leading provider of high-performance 3D visualization toolkits and application software for demanding industrial and scientific applications. Mercury VSG customers are in the geosciences, materials science, oil and gas, manufacturing, and engineering industries.

NAG (Oxford, England) is a worldwide leader in numerical software and high performance computing services. NAG serves leading HPC customers such as the HECToR program in the U.K.

Nema Labs (Göteborg, Sweden) offers technology to migrate software to multicore platforms. The FASThread product line is designed to parallelize code automatically and reliably.

NICE (Cortanze, Italy), with its EnginFrame grid portal, aims to increase user productivity through highly customizable, intuitive access to grid-enabled applications and infrastructures.

The **ParTec Cluster Competence Center** (Munich, Germany) specializes in cluster operating software and support services designed to deliver ground-breaking performance for large scale supercomputing clusters. Parastation5 is the current release of ParTec's cluster operating and management software.

science + computing ag (Tübingen, Germany), a subsidiary of the Bull Group, is a service and software company for technical/scientific design and simulation environments (CAD/CAE/CAT). s+c focuses on the operation of complex Unix, Linux and Windows environments.

Scilab (Le Chesney Cedex, France) is the free open source software for numerical computation. It is distributed worldwide with more than 70,000 downloads every month coming from more than 80 countries. A new company, Scilab Enterprises, has been created to offer support, migration, development and integration of specific applications in the companies technological environments.

Ter@tec is a European initiative based in France whose objective is to build the first technopole in Europe dedicated to Simulation and High Performance Computing. One characteristic of Ter@tec is that it regroups a large number of industrial companies, users and providers, and the major R&D centers, public and private. Teratec is building the Ter@tec Campus (15 000 m2 of offices and labs), close to the Very Large Computing Center (several petaflops, including the future French PRACE system).

T-Platforms (Moscow, Russia) provides HPC systems, software, storage and services.

transtec (Tübingen, Germany) provides consultation and integration services for HPC projects, including computer systems, storage, backup, and desktop requirements.

Ylichron (S. Maria di Galeria, Italy) was founded in 2005 as a spinoff company of ENEA (the Italian Agency for the New Technologies, the Energy and the Environment). The company's main product is the HCE compiler.

9.3 Summary of Relevant IDC ISV Data and Information

The Origins and Growth of HPC Software

Following its initial growth spurt in the late 1970s and 1980s, the market for HPC systems has expanded over time by adapting to the requirements of successive waves of new users — in large part through advances in software. Each new wave of users has expected HPC system vendors to do more for them, by providing software to make these systems easier to deploy and use.

In 1976, the CRAY-1 supercomputer was delivered to its first customer, Los Alamos National Laboratory (LANL), as a blazingly fast hardware platform with no operating system. Not to worry: LANL and others in the first wave of HPC users, primarily government and university researchers, typically had enough in-house technical savvy and personnel to write software themselves when the need was critical. The second wave of adoption carried HPC into industry, initially the automotive and aerospace sectors, starting in the late 1970's. These users required HPC vendors to provide not only an operating system and other system software, but to port the key third-party ISV applications needed to run the users' industry-specific problems — and to run the applications with the reliability expected in production computing environments.

HPC Software in the Era of Clusters

The third important wave of HPC market growth began with the birth of the cluster in the late 1990s and took off in earnest in 2002. Since that year, clusters have almost single-handedly driven the rapid growth in the HPC market and are now the dominant species of technical servers. But especially since the advent of clusters, HPC hardware has increasingly outdistanced software development, such that today it is not uncommon for software to be capable of exploiting only a small fraction — often not even 5% to 10% — of a contemporary hardware system's theoretical ("peak") processing power. The peak speeds of the largest HPC hardware systems have jumped about one thousand-fold per decade during the past two decades, primarily by greatly increasing the number of processing elements. The largest HPC systems today include more than 200,000 processor cores, and systems with more than one million cores are just a few years away. Yet, relatively few HPC software applications ("codes") today are able to exploit more than 128 cores.

HPC systems have become more difficult for HPC software to exploit efficiently not only because of their skyrocketing parallelism and theoretical peak speeds, but also because with few exceptions they employ standard x86 microprocessors from Intel and AMD that were designed for broader computer markets — especially personal computers — rather than to meet the specific demands of HPC users. Adding to the software challenges is the growing trend toward heterogeneous processing in HPC systems, that is, the use of more than one type of processor.

Because x86-based HPC hardware systems are firmly entrenched and widely available from multiple vendors based in the U.S., Japan, and Europe (Bull, EuroTech), future HPC leadership will likely depend far more heavily on advances in software than on initiatives to develop novel hardware systems.

- **Programming Languages.** Fortran, C, and C++ are the still-important legacy programming languages for HPC, but the rise of clusters to market dominance in recent years has made the Message-Passing Interface (MPI) standard the preferred protocol for programming high performance computers. MPI can be very labor-intensive and is not particularly efficient even for today's large HPC systems, not to mention the million-core systems that are on the near horizon. More efficient alternatives have been available for some time, notably Co-Array Fortran (CAF) and Unified Parallel C (UPC), and other so-called PGAS (Partitioned Global Address Space) languages are under development by Cray, IBM, and Oracle-Sun, but research has repeatedly shown that few HPC users are ready to leap to a new, more efficient programming language that would require new learning and the rewriting of applications.

- **Applications Software.** Software applications, also known as "programs" or "codes", enable users to carry out specific tasks, such as word processing on personal computers or climate modeling on HPC systems. Government and academic users of large HPC systems typically employ application software that they have created themselves ("in-house codes") or that has been created by some other government or academic organization for common use ("community codes"). Industry is typically far more reliant on application software that is purchased from and maintained by commercial software firms called independent software vendors (ISVs) — although ISV software may also be used in government and academia. IDC studies show that open source application software does not yet play a major role at most government, academic or industrial HPC sites, nor is it expected to do so in the foreseeable future. Only about 3% of HPC applications are "open source" codes. In addition, most open source software is middleware and not application software.

Studies of HPC Applications Software

In 2005 and 2006, IDC conducted two related, pioneering studies on HPC applications software. Although these studies were commissioned by two U.S. entities, the Defense Advanced Research Projects Agency (DARPA) and the Council on Competitiveness, their scope was global. The first study looked at HPC applications software from the standpoint of end users, while the second study explored the same phenomenon from the perspective of the independent software vendor (ISV) community.

The *Study of ISVs Serving the HPC Market: The Need for Better Application Software* (July 2005) showed that 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. As for-profit companies in most cases, ISVs 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 5% of their overall revenues, and in some cases this figure is less than 1%. Even if they could afford this investment, the motivation for major rewrites is generally inadequate because the HPC market is too small to reward this investment. For business reasons, the needs of HPC users are often an important but secondary concern. The *Study of ISVs Serving the HPC Market: Part B — End User Perspectives* (February 2006) surveyed a select group of well-known U.S. businesses that are highly experienced HPC users. IDC asked them about their requirements for HPC-specific application software and related resources. The HPC end users IDC interviewed for Part B represented a wide range of industries, from defense contractors to an entertainment company and a consumer products supplier. Most of the firms (83%) said they have unsolvable problems that are 5–100 times larger than the problems they can solve today. Also, it is important to note that when industrial HPC end users talk about solving larger problems, they typically don't mean simply doing more of the same thing. In most cases, they mean solving problems with greater resolution that can lead to new insights and superior new products.

IDC Top HPC Predictions for 2011 And Beyond

IDC annually provides its clients with HPC predictions for the coming year. Here are the predictions for 2011 and beyond.

- More Real-World Applications Will Be Run at Trans-Petaflop Speeds
- There May Be More Emphasis on Software … Finally
- The Alternative Processor Wars Will Keep Heating Up
- The HPC Staffing Shortage Will Grow More Acute
- Cloud Computing Will Ramp Up Slowly in HPC But Will Find a Number of Strong Niches
- The HPC Storage Market Will Continue to Grow Faster than the HPC Server Market
- InfiniBand Will Continue to Take Market Share from Proprietary Interconnects While Ethernet Remains the Leader
- Power and Cooling Will Become an Even Greater Concern

9.4 Terminology and Definitions

Terminology and Definitions

Application Workload

IDC uses this term to refer to a set of related scientific or engineering applications, plus the associated infrastructure or support applications (e.g., visualization, data management). Each application workload is characteristic of a market segment. Hence, the application workload for biological sciences differs from that used in chemical engineering or other segments.

Cluster

IDC defines an HPC cluster 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 systems integration strategies. In addition, HPC clusters must use base technologies that are readily available in the open marketplace, such as x86 microprocessors, commercial interconnects (in cases where a cluster includes an interconnect), and a commercial operating system (e.g., Linux or Windows). A system using a base technology that is available to only one or two vendors, such as an IBM POWER processor or the Cray SeaStar interconnect, does not qualify as a cluster in IDC's usage; nor do products employing SMP or MPP architectures that are designed as single systems. Clusters vary greatly in size and range in price from about $25,000 (€19,000) to tens or hundreds of millions of dollars each.

Supercomputer

In this report, IDC uses the term *supercomputer* in two senses, depending on the context. First, IDC uses the term *supercomputer* to refer to HPC systems priced at $500,000 (€375,000) and above. This use of the term is always associated with one or more numerical values, such as a figure in a numerical table that represents the value of all *supercomputers* sold within the EU in a given year. When the term is not associated with numerical values, it is used as a synonym for an *HPC system* of any type and price. This corresponds to the popular and most pervasive use of the term *supercomputer*.

Technical Computing and High-Performance Computing (HPC)

IDC uses the term *technical computing and high-performance computing (HPC)* to encompass the entire market for computer servers used by scientists, engineers, analysts, and other groups using computationally intensive modeling and simulation applications. Technical servers range from small servers costing less than $5,000 to the large-capability machines valued in hundreds of millions of dollars. In addition to scientific and engineering applications, technical computing includes related markets/applications areas including economic analysis, financial analysis, animation, server-based gaming, digital content creation and management, business intelligence modeling, and homeland security database applications. These areas are included in the technical computing market based on a combination of historical development, applications type, computational intensity, and associations with traditional technical markets.

9.5 Background Documents

Turning the European Union into an Innovation Union. Talk by Robert-Jan Smits, Director General for Research and Innovation of the European Commission. June 13, 2011. http://www.publicservice.co.uk/feature_story.asp?id=16729

Investigation Report on Existing HPC Initiatives. Berthou, JY, et al. Interim report of the European Exascale Software Initiative (EESI). 2010. This and other valuable information resources are available at the EESI website: www.eesi-project.eu

High Performance Computing: A Key Technology for Europe. Sawyer, M., Parsons, University of Edinburgh.2011. http://www.planethpc.eu/images/stories/hpc_a_key_technology_for_europe.pdf

EGI-InSPIRE Project Presentation. Steven Newhouse. European Grid Initiative. Updated July 2011. https://documents.egi.eu/public/ShowDocument?docid=506

PRACE D1.5 Final Report. Thomas Eickermann, FCJ. June 27, 2010. http://www.prace-project.eu/documents/public-deliverables/d1-5.pdf

The International Exascale Software Roadmap. Dongarra, J., Beckman, P. et al., Volume 25, Number 1, 2011, International Journal of High Performance Computer Applications, ISSN 1094-3420. http://www.exascale.org/mediawiki/images/2/20/IESP-roadmap.pdf

IDC HPC End-User Study of the Industry/Application Trends in Technical Computing. A multi-client, demand-side study, worldwide in scope. December 2010.

Development of a Supercomputing Strategy in Europe. Interim report of an IDC study on behalf of the DG Information Society and Media of the European Commission (SMART 2009/0055, contract number 2009/S99-142914). April 2010. http://www.hpcuserforum.com/EU/downloads/IDCWP12S35.10.2010version1.3.pdf

A Strategic Agenda for Supercomputing in Europe: HPC 2020. Final report of the IDC study, *Development of a Supercomputing Strategy in Europe,* conducted on behalf of the DG Information Society and Media of the European Commission (SMART 2009/0055, contract number 2009/S99-142914). September 2010. http://www.hpcuserforum.com/eu/downloads/SR03S10.15.2010.pdf

Advance: Benchmarking Industrial Use of High Performance Computing for Innovation. Worldwide study conducted by IDC for the Council on Competitiveness and the U.S. Department of Energy. 2008. Downloadable without charge at http://www.compete.org/publications/detail/486/advance/

Study of ISVs Serving the High Performance Computing Market: Part A, Independent Software Vendors. IDC study for the Council on Competitiveness. 2005. http://www.compete.org/publications/detail/392/hpc-software-study-part-a-current-market-dynamics/

Study of ISVs Serving the High Performance Computing Market: Part B, HPC End Users. IDC study for the Council on Competitiveness. 2006. http://www.compete.org/publications/detail/393/hpc-software-study-part-b-end-user-perspectives/

