

PRESIDENT'S INFORMATION TECHNOLOGY ADVISORY COMMITTEE
REPORT TO THE PRESIDENT



**Information Technology Research:
Investing in Our Future**

February 1999

PRESIDENT'S INFORMATION TECHNOLOGY ADVISORY COMMITTEE
REPORT TO THE PRESIDENT



**Information Technology Research:
Investing in Our Future**

February 1999

PRESIDENT'S INFORMATION TECHNOLOGY ADVISORY COMMITTEE

February 24, 1999

The President of the United States
The White House

Dear Mr. President:

We are pleased to present our final report, "Information Technology Research: Investing in Our Future," on future directions for Federal support of research and development for information technology. This report adds detail to the findings and recommendations in our interim report dated August 1998, and strengthens our previous recommendations regarding the importance of social and economic research on the impacts of information technology to inform key policy decisions.

PITAC members are strongly encouraged by and enthusiastically supportive of the Administration's Information Technology for the Twenty-First Century (IT2) initiative. This initiative is a vital first step in increasing funding for long-term, high-risk information technology research and development.

Increased Federal support is critical to meeting the challenge of capturing the opportunities available from information technology in the 21st Century through appropriate research and development. The economic and strategic importance of information technology and the unique role of the Federal Government in sponsoring information technology research make it necessary to increase Federal support over a period of years to ensure our Nation's future well-being.

We hope that our recommendations will be helpful as you consider the priorities for Federal investments. Thank you for your consistent support of our activities over the past year. We look forward to discussing this report with you, with members of your Administration and with members of Congress.

Respectfully,

Bill Joy
Co-Chairman

Ken Kennedy
Co-Chairman

President's Information Technology Advisory Committee

Co-Chairs

Bill Joy
Founder and Chief Scientist
Sun Microsystems

Ken Kennedy, Ph.D.
Ann and John Doerr Professor
of Computational Engineering
Rice University

Members

Eric A. Benhamou
CEO and Chairman
3Com Corporation

Sherrilynne S. Fuller, Ph.D.
Head, Division of Biomedical Informatics,
School of Medicine and Director,
Health Sciences Libraries and Information Center
University of Washington

Vinton Cerf, Ph.D.
Senior Vice President for Internet
Architecture and Technology
MCI WorldCom

Hector Garcia-Molina, Ph.D.
Leonard Bosack and Sandra Lerner
Professor, Departments of Computer Science and
Electrical Engineering
Stanford University

Ching-chih Chen, Ph.D.
Professor, Graduate School of Library
and Information Science
Simmons College

Susan L. Graham, Ph.D.
Chancellor's Professor of Computer Science,
Department of Electrical Engineering
and Computer Science
University of California, Berkeley

David M. Cooper, Ph.D.
Associate Director of Computation
Lawrence Livermore National Laboratory

James N. Gray, Ph.D.
Senior Researcher, Scalable Servers
Research Group
Microsoft Research

Steven D. Dorfman
Vice Chairman
Hughes Electronics Corporation

W. Daniel Hillis, Ph.D.
Vice President and Disney Fellow
Walt Disney Imagineering
Research and Development, Inc.

David W. Dorman
President and Chief Executive Officer
PointCast, Inc.

Robert H. Ewald
Former President and COO
Cray Research, Inc.

Robert E. Kahn, Ph.D.
President
Corporation for National Research
Initiatives

David J. Farber
Alfred Fitler Moore Professor of
Telecommunication Systems
University of Pennsylvania

PRESIDENT'S INFORMATION TECHNOLOGY ADVISORY COMMITTEE

John P. Miller, Ph.D.
*Director, Center for Computational Biology,
and Professor of Biology*
Montana State University

David C. Nagel, Ph.D.
*Chief Technology Officer, AT&T
and President AT&T Labs*

Raj Reddy, Ph.D.
*Dean, School of Computer Science, and
Herbert A. Simon University Professor of Computer
Science and Robotics*
Carnegie Mellon University

Edward H. Shortliffe, M.D., Ph.D.
*Associate Dean for Information Resources and Technology,
Professor of Computer Science, and Professor of Medicine*
Stanford University School of Medicine

Larry Smarr, Ph.D.
*Director, National Computational Science Alliance,
and Professor of Physics and Astronomy*
University of Illinois at Urbana-Champaign

Joe F. Thompson, Ph.D.
*William L. Giles Distinguished Professor of
Aerospace Engineering*
Mississippi State University

Leslie Vadasz
*Senior Vice President and Director of
Corporate Business Development*
Intel Corporation

Andrew J. Viterbi, Ph.D.
Vice Chairman, Board of Directors
QUALCOMM Incorporated

Steven J. Wallach
Advisor
CenterPoint Ventures

Irving Wladawsky-Berger, Ph.D.
General Manager, Internet Division
IBM Corporation

Table of Contents

Letter of Transmittal	
Members of the Advisory Committee	
Table of Contents	vii
Acknowledgments	ix
Executive Summary	1
Findings and Recommendation	3
Priorities for Research	4
The Government’s Essential Role	6
Conclusion	7
Sidebar : Rationale for Government Support of Long Term, Fundamental Research	8
1. Information Technology: Transforming our Society	11
1.1 Transforming the Way We Communicate	12
1.2 Transforming the Way We Deal With Information	13
1.3 Transforming the Way We Learn	13
1.4 Transforming the Practice of Health Care	14
1.5 Transforming the Nature of Commerce	14
1.6 Transforming the Nature of Work	15
1.7 Transforming How We Design and Build Things	15
1.8 Transforming How We Conduct Research	16
1.9 Transforming Our Understanding of the Environment	16
1.10 Transforming Government	17
Table of Transformations	18
2. Setting Federal Research Priorities: Findings and Recommendations	21
2.1 Findings	21
2.2 Recommendations for Research	24
3. Technical Research Priorities	27
3.1 Software Research	27
3.1.1 Findings	27
3.1.2 Recommendations	31
3.2 Scalable Information Infrastructure	35
3.2.1 Findings	35
3.2.2 Recommendations	38
3.3 High-End Computing	45
3.3.1 Findings	46
3.3.2 Recommendations	50
4. Socioeconomic Research and Policy Priorities	55
4.1 Findings	55
4.2 Recommendations	60

5. Creating an Effective Management Structure for Federal Information Technology R&D67
5.1 Agency Roles67
5.2 Policy and Coordination of Information Technology R&D69
5.3 Support and Implementation70
5.4 Annual Review75
5.5 Funding75
Sidebar: Can the Proposed Increases be Invested Effectively?76
Conclusion79

ACKNOWLEDGEMENTS

The following members of the community were especially helpful in working with PITAC as panel members or advisors after the interim report was released in August, 1998.

Victor Basili	Dave Patterson
George Campbell	Dan Reed
Jack Dongarra	Michael Skibo
Jim Flanagan	George Spix
Donna Hoffman	Tom Sterling
Jeff Johnson	Rick Stevens
Brewster Kahle	Michael S. Teitelbaum
Ed Lazowska	Mary Vernon
C. Dianne Martin	Stephen Wolff
Reagan Moore	Alexander Wong
Donald Norman	William Wulf
Craig Partridge	

PITAC would like to thank the following people for their contribution in writing white papers for PITAC's consideration for the interim report.

Kathie Blankenship	Rob Kling
Scott O. Bradner	Clifford Lynch
Michael L. Brodie	Daniel R. Masys
Michael Hawley	

PITAC would also like to acknowledge the contributions of the following people from the National Coordination Office for Computing, Information and Communications, without whom we would not have been able to produce this report:

Kay Howell	Yolanda L. Comedy
Sally Howe	Robert I. Winner

PITAC would also like to thank the numerous Federal employees who gave advice during the writing of this report, in particular, Henry Kelly, Tom Kalil and Lori Perine.

And finally, we are grateful to the entire staff at the National Coordination Office. Our meetings went smoothly because of their careful preparation.

Executive Summary

Information Technology will be one of the key factors driving progress in the 21st century—it will transform the way we live, learn, work, and play. Advances in computing and communications technology will create a new infrastructure for business, scientific research, and social interaction. This expanding infrastructure will provide us with new tools for communicating throughout the world and for acquiring knowledge and insight from information. Information technology will help us understand how we affect the natural environment and how best to protect it. It will provide a vehicle for economic growth. Information technology will make the workplace more rewarding, improve the quality of health care, and make government more responsive and accessible to the needs of our citizens.

Vigorous information technology research and development (R&D) is essential for achieving America's 21st century aspirations. The technical advances that led to today's information tools, such as electronic computers and the Internet, began with Federal Government support of research in partnership with industry and universities. These innovations depended on patient investment in fundamental and applied research.

We have had a spectacular return on that Federal research investment. Businesses that produce computers, semiconductors, software, and communications equipment have accounted for a third of the total growth in U.S. economic production since 1992, creating millions of high-paying new jobs. Government-sponsored University research programs have supported graduate education for many of the leaders and innovators in the field. As we approach the 21st century, the opportunities for innovation in information technology are larger than they have ever been—and more important. We have an essential national interest in ensuring a continued flow of good new ideas and trained professionals in information technology.

After careful review of the Federal programs this Committee has concluded that Federal support for research in information technology is seriously inadequate. Research programs intended to maintain the flow of new ideas in information technology and to train the next generation of researchers are funding only a small fraction of the research that is needed, turning away large numbers of excellent proposals. Compounding this problem, Federal agency managers are faced with insufficient resources to meet all research needs and have naturally favored research supporting the short-term goals of their missions over long-term high-risk investigations. While this is undoubtedly the correct local decision for each agency, the sum of such decisions threatens the long-term welfare of the nation.

The Nation needs significant new research on computing and communication systems. This research will help sustain the economic boom in information technology, address important societal problems such as education and crisis management, and protect us from catastrophic failures of the complex systems that now underpin our transportation, defense, business, finance, and health-care infrastructures. If the results are to be available when needed, we must act now to reinvigorate

the long-term IT research endeavor and to revitalize the computing infrastructure at university campuses and other civilian research facilities, which are rapidly falling behind the state of the art. If we do not take these steps, the flow of ideas that have fueled the information revolution over the past decades may slow to a trickle in the next.

To address these problems, the Committee estimated in its Interim report in August 1998 that the Federal government should increase its support for information technology research by a billion dollars per year by FY 2004. Since that time the Committee has sought comments from the community regarding its preliminary findings and recommendations, and convened several panels to review those recommendations. This effort produced a more detailed model for the costs of the research programs and other activities needed to address the problems identified in our report. As a result of these activities, the Committee has further refined the findings and recommendations presented in its Interim Report, and adjusted its funding recommendation accordingly. The Committee now recommends that the Federal government increase annual funding for information technology R&D over the five-year period from FY2000 to FY2004, as follows:

Recommending Funding Increases for Information Technology R&D (\$ millions)					
Area	FY 2000	FY 2001	FY 2002	FY 2003	FY 2004
Software	112	268	376	472	540
Scalable Information Infrastructure	60	120	180	240	300
High End Research	180	205	240	270	300
High End Acquisitions	90	100	110	120	130
Socioeconomic	30	40	70	90	100
Total	472	733	976	1192	1370

These increases are in addition to the programs in existence in FY99. The full report provides additional details on these budget recommendations, including discussions of the method used to produce them.

Although there are unmet needs across the entire spectrum of research activities, priority for increased funding should be on long-term, high-risk investigations. In addition to increases for research itself, the Federal government must also ensure that the research community is equipped with the state-of-the-art facilities needed to carry out advanced projects. Finally, Federal budgets must continue to ensure that advances in information technology work to benefit all Americans and that all Americans have the education and training needed to prosper in a world that will increasingly depend on information technology.

To be successful, the expanded Federal research program we propose must be effectively managed. Current cross-agency coordination mechanisms are working well, but they suffer from the lack of well-defined responsibilities for ensuring that key areas are not overlooked. There has been no agency with the primary responsibility for ensuring that long-term, high-risk research is protected from the pressures that arise in mission agencies. Ideally there should be an agency charged with leading the organization of a fundamental information technology research program appropriate for the 21st century.

The Administration's proposed Federal budget for FY 2000 demonstrates a commitment to sustained growth in information technology research through its initiative, *Information Technology for the Twenty-First Century (IT²)*. This commitment is an important first step in what must be a continuing effort on the part of the Federal government to increase research dollars and to create a new management system designed to foster innovative research. But the effort cannot stop here. Further increases and continued oversight are needed to remedy the shortfall in long-term research investments that has accrued.

Findings and Recommendation

Federal information technology R&D investment is inadequate. Measured in constant (non-inflated) dollars, support in most critical areas has been flat or declining for nearly a decade, while the importance of information technology to our economy has increased dramatically. As a result, the Nation is gravely underinvesting in the long-term, high-risk research that can replenish the reservoir of ideas that will lead to innovations in information technology in generations to come.

Federal information technology R&D is too heavily focused on near-term problems. Much of the Federal investment in information technology R&D is being funded by mission agencies. In the face of the enormous increases in information technology problems to be addressed, funding agencies have had to prioritize their investments. Inevitably, priority has been given to short-term, mission-oriented goals over long-term research. This reflects the situation in the private sector as well. As a result, investment in long-term, high-risk research has been curtailed. This trend threatens to interrupt the flow of ideas that has driven the information economy in this decade and threatens efforts to solve nationally important problems.

Recommendation: Create a strategic initiative in long-term information technology R&D. To address these problems the Committee recommends that the President create a strategic initiative to support long-term research in fundamental issues in computing, information, and communications. The initiative should increase the total funding base by \$1.37 billion per year by FY 2004. The Federal funding agencies should use the resulting budget increases to encourage research that is visionary and high-risk. To do this, they will need to diversify modes of research support and increase the duration of projects. The goal should be to recapture in the universities and research labs much of the excitement that existed at top-rated departments in the past.

Priorities for Research

Four areas of the overall research agenda particularly need attention and must be a major part of a strategic initiative in long-term research and development:

Software – The demand for software has grown far faster than our ability to produce it. Furthermore, the Nation needs software that is far more usable, reliable, and powerful than what is being produced today. We have become dangerously dependent on large software systems whose behavior is not well understood and which often fail in unpredicted ways. Therefore, increases in research on software should be given a high priority. Special emphasis should be placed on developing software for managing large amounts of information, for making computers easier to use, for making software easier to create and maintain, and for improving the ways humans interact with computers. Specifically, the Federal program should:

- Fund fundamental research in software development methods and component technologies.
- Support fundamental research in human-computer interfaces and interaction.
- Support fundamental research in capturing, managing, analyzing, and explaining information and in making it available for its myriad of uses.
- Make software research a substantive component of every major IT research initiative.

Scalable Information Infrastructure – Our Nation's dependence on the Internet is increasing. While this is an exciting development, the Internet is growing well beyond the intent of its original designers and our ability to extend its use has created enormous challenges. As the size, capability, and complexity of the Internet grows, it is imperative that we do the necessary research to learn how to build and use large, complex, highly-reliable, and secure systems. It is therefore important that the Federal government:

- Fund research on understanding the behavior of the global-scale network and its associated information infrastructure. This should include collecting and analyzing performance data and modeling and simulating network behavior.
- Support research on the physics of the network, including optical and wireless technologies including satellites, cable, and bandwidth issues.
- Support research to anticipate and plan for scaling the Internet.
- Support research on middleware that enables large-scale systems.
- Support research on large-scale applications and the scalable services they require.
- Fund a balanced set of testbeds and research infrastructure that serve the needs of networking research, research in enabling information technologies and advanced applications.

High-End Computing – Extremely fast computing systems, with both rapid calculation and rapid data movement, are essential to provide accurate weather and climate forecasting, to support advanced manufacturing design, to design new pharmaceuticals, to conduct scientific research in a variety of different areas, and to support critical national interests. Although they achieve remark-

able performance in some cases, the current scalable, parallel, high-end computing systems are not well suited to many nationally important, strategic applications. To ensure that U.S. scientists continue to have access to computers of the highest possible power, funding should be focused on innovative architectures, hardware technologies, and software strategies that overcome the limitations of today's systems. Without major increases in funding in these areas, the realizable performance of new machines will fall far short of their potential. We specifically recommend that the Federal program should:

- Fund research into innovative computing technologies and architectures.
- Fund R&D on software for improving the performance of high-end computing.
- Drive high-end computing research by trying to attain a sustained petaops/petaflops on real applications by 2010 through a balance of software and hardware strategies.
- Fund the acquisition of the most powerful high-end computing systems to support science and engineering research.
- Expand the National Science and Technology Council (NSTC) Computing Information and Communications (CIC) High-End Computing and Computation (HECC) Working Group's coordination process to include all major elements of the government's investment in high-end computing.

Socioeconomic Impact – Information technology will significantly improve the flow of information to all people and institutions in the Nation, and could be a powerful tool for democratization. Our National well-being depends on understanding the potential social and economic benefits of on-going advances in information technology. However, problems are arising from the increasing pace of information technology-based transformations. To realize the promise of the new technologies, we must invest in research to identify, understand, anticipate, and address these problems. We must develop concrete objectives and comprehensive metrics through which to assess the ongoing transformations brought about by the integration of information technology into our lives. We must conduct careful research on the impact of the transformations against the objectives, and develop appropriate policies to deal with the knowledge we gain from the research. The Federal government should:

- Expand Federal initiatives and government/university/industry partnerships to increase information technology literacy, access and research capabilities.
- Expand Federal research into policy issues arising from information technology.
- Fund information technology research on socioeconomic issues.
- Expand the participation of underrepresented minorities and women in computer and information technology careers.
- Create programs to remove the barriers to high bandwidth connectivity posed by geographic location, size, and ethnic history of research, educational institutions, and communities.
- Accelerate and expand education in information technology at all levels—K-12, higher education, and lifelong learning.
- Strengthen the use of information technology in education.

Management and Implementation of Federal Information Technology Research –Building a Federal IT program suited to the needs of the Nation in the 21st century will require new management strategies, new modes of research support, and new implementation strategies. This new approach is demanded by the reality of Federal budget constraints, the need for more long-term cross-disciplinary team research, and the need to maintain a small, efficient, and coordinated research management process. It is essential that the Federal systems responsible for managing and implementing the new IT program be positioned to review the entire information technology research budget, to restore the balance between fundamental and applied research, to encourage long-term and high-risk collaborative research projects, and to employ a systematic review by participating Federal agencies and the private sector.

To achieve these goals, we recommend that the existing Federal information technology management and implementation structure be enhanced as follows:

- Strongly encourage NSF to assume a leadership role in basic information technology research. Provide NSF the necessary resources to play this role.
- Designate a Senior Policy Official for information technology R&D.
- Extend the HPCC program coordination model to all major Federal information technology R&D activities.
- Diversify the modes of research support to include more projects of broader scope and longer duration, placing a renewed emphasis on research carried out in teams.
- Fund collaborations with applications to drive information technology research, but take measures to ensure that research remains a primary goal.
- Fund centers for Expeditions into the 21st Century.
- Establish a program of Enabling Technology Centers that will drive research by examination of critical applications areas.
- Establish an annual review of research objectives and funding modes.

The Government's Essential Role

While the importance of information technology to the future of the economy and the government is clear, it may not be immediately obvious that government investment is needed to ensure continued progress. The PITAC members from industry were unanimous in their opinion that it is not feasible for the private sector to assume responsibility for long term, high-risk research, in spite of the success of the information technology industry. Their opinion is found in the attached sidebar.

We believe that the Federal Government must retain and expand its role in leading long-term fundamental research in information technology. Advanced Government services and national security depend on it. The benefits to our Nation and society will be huge. A loss of international leadership in information technology would be economically devastating.

We cannot rely on industry to fund the needed research because they necessarily focus, in view of economic realities, on the short term. Industry cannot and will not invest in solving problems of importance to society as a whole unless such investments make sense from a business perspective.

The rationale for funding long-term information technology research goes far beyond economic benefit and national security needs. Enormous societal gains can be reaped from advances in information technology. Only through research on a scale substantially greater than is being carried out today, can we build an infrastructure that will be available, affordable and usable by all citizens—one that can support the compelling “transformations” discussed in the next chapter of this report.

Federal investment in information technology directly supports the education and preparation of our young people for careers in information technology research, and the training of workers who need to upgrade their skills to keep pace with a changing marketplace. Trained people are not just a byproduct but rather a major product of publicly supported research. These trained professionals are critical national infrastructure, and will create and develop new ideas, form a talent pool for existing business, and launch new companies.

The benefits that the transformations described in the next chapter of this report can have for our Nation’s future are extraordinary. A networked society can reach out to all its citizens, and can bring our Nation closer together and address many societal issues. While it cannot resolve all these issues, information technology can give us leverage toward their solution.

But the realization of the positive transformations we will describe in the next chapter of this report is not guaranteed. The realization of each transformation will depend on the results of aggressive, well-managed Federal research programs. Long-term Federal investment in information technology research is necessary to incubate ideas to the point of clear commercial viability, and to develop methods of measuring and tracking our progress toward realizing our positive vision.

Conclusion

Information technology research is essential for the continued growth of the economy and for the solution of some of the most critical problems facing the Nation. Unless steps are taken now to reinvigorate Federal research in this critical area, we are very likely to see a significant reduction in the rate of progress over the coming decades. The cost to the Nation of such a reduction will be significantly greater than the investment needed to address the problem.

PITAC Industry Member's Rationale for Government Support of Long Term, Fundamental Research

Our PITAC report shows that as information technology has grown to be an extremely successful segment of the economy, and more important to the future of the economy and the government, it has outpaced the research budgets available to support it. At a time when we need more fundamental understanding of information technology, and more long term, high-risk projects, tighter overall budgets and near-term market pressures are squeezing research budgets.

Should we expect information technology research to be funded primarily by commercial interests? To understand this question it helps to review the history of Federal funding in this area.

Since World War II, the Federal government has funded advanced information technology research to meet its own requirements, which have ranged from critical national-defense applications to weather forecasting and medical sciences. Federal funding has seeded high-risk research and yielded an impressive list of billion-dollar industries (such as the Internet, high performance computers, RAID disks, multiprocessors, local area networks, graphic displays). Federally funded university research has trained most of our leading IT researchers. Information technology industries provide hundreds of thousands of jobs and much of the nation's recent economic growth. The Federal investment to date has had tremendous benefits for our Government, our Nation and our economy.

The information technology industry expends the bulk of its resources, both financial and human, on rapidly bringing products to market. The U.S. information technology industry has created an awesome and continuous growth of capabilities based on the most intensely competitive marketplace the world has ever seen. Nearly every available person and dollar in this industry is focused on bringing the next version or the next product to market. Delivery product cycles are as short as every three to six months. The company that fails here misses the next short-term cycle and will not be successful.

The information technology sector has continued to invest extensively in research and development, and has more than doubled its annual R&D investment over the past 10 years, to reach the current level of about \$30B.¹ For information technology companies, the fraction of the total budget devoted to R&D is roughly twice the U.S. industry average.² Over 90 percent of information technology R&D expenditures are allocated to product development, with the major portion of the remaining expenditures going toward near term, applied research.

Data from the top 50 information technology companies between 1990 and 1997 show that while both overall revenues and overall R&D spending grew, the portion of revenue invested in R&D declined. The ratio of R&D expense to overall revenue declined from about 9 percent to about 7 percent at the end of the period.³ The declining percentage of investment in R&D reflects the intense pace of the information technology marketplace, and the equally intense competitive pressures on prices and profit margins. Today, industry's efforts are mostly focused on projects that will generate near-term revenues. The reasons for this short-term focus can be found in the changes that have taken place in the information technology industry through the 1990s.

The decade of the 90s has been characterized by a major change in the nature of information technology industry sales. There has been a shift from complex products with large margins to more commodity-like products, such as PCs, with lower margins. Services is the fastest growing segment in the industry, and profit margins in the people-intensive services segment are often lower than those in software and hardware products. Such price and resulting profit declines have required widespread cost control programs. These programs impacted both the expenditures and the nature of the R&D programs.

Another inhibitor to long-term research spending is intense global competition, which puts great pressure on prices, profits, and expenses. Such pressures escalate as new companies are formed at a rapid pace and try to dislodge each other and existing companies, often by competing on price. Further, the incentives associated with long term research, such as prospects for higher P/E ratios, opportunities to receive long term R&D tax credits or to write-off the investments of "in-progress R&D" are either non-existent, or weak and getting weaker.

The fundamental research investments in university research by the Federal government have served to train the majority of our information technology professionals. In recent years, however, the shortage of workers with adequate skills makes it difficult for companies to grow both near and long-term research, even if budgets allowed. This alone suggests the critical need for additional government support of university research.

Past Government investments in high-risk research have helped fuel the intense pace of the information technology marketplace. The U.S. has the most energetic, viable and productive technology transfer mechanism in the world. Ideas freely flow from universities and national labs to existing and new companies. In 1998, over \$12 billion dollars were invested by venture funds in new companies.⁴ The basic feedstock for these investments has been Government support of basic information technology research. If this feedstock is allowed to deplete, this economic growth engine could slow or disappear.

The United States must not only continue, but also substantially increase, long-term fundamental information technology research programs. Once innovative research ideas have been explored, American companies are well positioned to seize the viable ideas, commercialize and distribute them. The government must increase its investment in the pipeline that generates these ideas and the researchers to work on them. This will create a better society and bring even greater prosperity to our nation.

¹ *Endless Frontier, Limited Resources: U.S. R&D Policy for Competitiveness*, Council on Competitiveness, Washington, D.C., 1996: 110.

² *Ibid*, 111.

³ Data from Compustat, an S&P subscription database. Data was compiled by Compustat from SEC 10-K and 10Q report filings, 1990-1997, then summarized by IBM.

⁴ "Venture Capital Investments Hit Another Record High Despite Stock Market Turmoil." *Money Tree Highlights/Q3: 1998*. Price Waterhouse Coopers. http://209.67.194.61/reportsq398_highlights.asp

Industry Members of PITAC

Bill Joy, PITAC Co-Chair
Sun Microsystems

Eric Benhamou
3Com Corporation

Vinton Cerf
MCI Communications

David W. Dorman
PointCast, Inc.

Steve D. Dorfman
Hughes Electronics Corp.

Robert Ewald
Former Cray Research, Inc.

James N. Gray
Microsoft Research

W. Daniel Hillis
Walt Disney Imagineering

David C. Nagel
AT&T Labs

Leslie Vadasz
Intel Corporation

Andrew J. Viterbi
QUALCOMM

Steven J. Wallach
CenterPoint Ventures

Irving Wladawsky-Berger
IBM Corporation

1. Information Technology: Transforming Our Society

Information technologies are an integral part of people's lives, businesses, and society. Advances in microprocessors, memories, storage, software, and communication technologies make it possible to build computers and computing devices that are increasingly affordable, as well as to enable the development of increasingly powerful systems at reasonable costs. The wide acceptance of Internet standards and technologies is helping us build global computer networks capable of connecting everything and reaching everyone.

Since ancient times, networks have offered opportunities for growth and innovation and have supplied structure to our economic and social systems. From the roads and aqueducts of the Roman Empire, to nineteenth century continental railroad systems, to the telecommunications, broadcast, and satellite networks of the twentieth century, networked capabilities have allowed us to overcome barriers of time and space, and to access and open new frontiers for human interaction and ingenuity.

The free flow of information is essential to a democratic society. Advances in information technology have the potential to significantly enhance this flow of information, and thus strengthen the institutions of our society, from financial markets to government agencies. The flow of information must, however, not only be “free,” but “fair.” Financial markets, for example, have learned that they must guard against abuses, such as insider trading. Businesses and healthcare institutions must guard against the misuse of personal information put in their trust by their customers. As we have the opportunity to use information technology to strengthen our societal institutions, we must understand the potential pitfalls, and the safeguards we must put in place to achieve both a free and fair flow of information.

As we approach the new millennium, it is clear that the “information infrastructure”—the interconnected networks of computers, devices, and software—may have a greater impact on worldwide social and economic structures than all networks that have preceded them. The advances in computing and communications technologies of the last decade have already transformed our society in many ways. These advances have transformed the ways in which we view ourselves, our relationships with each other and with other communities, and the ways in which we obtain services, ranging from entertainment and commerce to education and health care. Even so, we have only just begun to grasp the opportunities and experience the transformations that will occur as these technologies mature.

Major technical advances are needed to build a smoothly functioning information infrastructure that links together all people, institutions, and relevant devices (e.g., cars, gas meters, home thermostats, air conditioners) in our Nation and beyond. Only vigorous information technology research and development programs will enable us to achieve our objectives.

But, hard as the technical challenges might be, we must keep in mind the great socioeconomic issues ahead of us on the road to becoming a fully networked society. Thus, we complement the call for research to support the required technical advances with a call for research programs to help us understand and enhance the positive effects of information technology on our economy, society, culture and political system.

The information revolution puts a premium on basic knowledge, not just information technology literacy, but basic skills in reading, writing, communications, and teamwork. Education and training have become lifelong pursuits for our workforce, as new jobs requiring new skills are created, and older jobs and skills become obsolete. The Nation must ensure that access to the benefits of the information infrastructure are available to everyone in our Nation: to those living in small towns and rural areas as well as in big cities, to those living in poor inner city neighborhoods and tribal reservations, as well as in well-to-do suburbs and those who face daily challenges from disabilities. We should use information technology to bridge the gaps in our society, not to create new ones.

A significant portion of our national progress in computing and communications over the past decade has been leveraged from the Federal research programs established by the High Performance Computing Act of 1991 (P.L. 102-194). These programs comprised the High Performance Computing and Communications (HPCC) initiative, which was responsible for moving the U.S. into an era of teraflop computers and gigabit networks. A major focus of the HPCC initiative was a set of Grand Challenges, difficult scientific and engineering problems whose solution were advanced by applying high performance computing and communications technologies and resources. Our recommendations build upon the solid foundation of the HPCC program and expand the HPCC vision to meet the challenges facing us in the 21st century.

To ensure a rapid, smooth, and extendible transition into the 21st century, the President's Information Technology Advisory Committee has identified ten critical "National Challenge Transformations." These information technology transformations will affect how we communicate, how we store and access information, how we become healthier and receive proper medical care, how we learn, how we conduct business, how we work, how we design and build things, how we conduct research, how we sustain a livable environment, and how we manage our government in the next millennium. Exploring these dynamic transformations enables us to identify common information technology challenges critical to our Nation's future and provides a framework for our recommendations for Federal research investments.

1.1 Transforming the Way We Communicate

Vision: At least one billion people worldwide can access the Internet simultaneously and engage in real-time electronic meetings, download the daily news, conduct secure financial transactions, or talk to friends and relatives around the world. This can be done regardless of the language the participants speak, since language translation can be done instantaneously, and regardless of physical limitations of the individual, because devices can accept and provide input and output in many different ways.

The Internet lies at the heart of our communications revolution. But, the current Internet must be expanded in scale to accommodate anticipated growth in use and demands for reliability comparable to that of the modern telephone system. New and improved modes of human interaction with computers must be developed to enrich and simplify the way we communicate. We must understand the behavior of extremely large and complex systems and address the potential fragility of large numbers of autonomously interacting software systems. Global networking raises a host of international issues and even poses questions about the nature of national boundaries as information flows across them invisibly and multinational corporations use worldwide networks to pursue their global interests. Perhaps the biggest challenge of all is to understand how human beings can best take advantage of the new electronic communication possibilities, both individually and in groups.

1.2 Transforming the Way We Deal With Information

Vision: An individual can access, query, or print any book, magazine, newspaper, video, data item, or reference document in any language by simply clicking a mouse, touching a computer screen, talking to a computer, or blinking an eye. Individuals can easily select among modes of presentation: data, text, images, or audio. Information can be referenced and derivations can be incorporated in many new ways, adding value and revealing insights through networked and software-enabled tools. Entertainment can be richer and more personalized, enabling individuals to access music, videos or live events that appeal to them.

This transformation requires significant improvements in data access methods, including high performance information systems and tools to help individuals locate information and present, integrate, and transform the information in meaningful ways. Systems will require interfaces accessible both to experts and novice or infrequent users regardless of physical ability, education, or culture. Multi-modal human-computer interaction technologies are needed including speech, touch, and gesture recognition and synthesis. There are research requirements for topics ranging from network reliability and bandwidth, to scalable software support, database structure and retrieval algorithms, high-performance computing, and robust, reliable, secure ways to deliver and to protect critical information. Providing high-quality entertainment over the Internet requires considerable more bandwidth, advances in audio and video streaming, and better tools to attract good content creators. Challenging issues regarding dissemination of information in electronic form including copyright, intellectual property rights, and the development of realistic business models remain important policy and research topics.

1.3 Transforming the Way We Learn

Vision: Any individual can participate in on-line education programs regardless of geographic location, age, physical limitation, or personal schedule. Everyone can access repositories of educational materials, easily recalling past lessons, updating skills, or selecting from among different teaching methods in order to discover the most effective ways of learning. Educational programs can be customized to each individual's needs, so that the information revolution reaches everyone and personal digital libraries provide a mechanism for managing ones accumulated knowledge resources. Learning involves all our senses, to help focus each student's attention and better communicate educational material.

Information technology is already changing how we teach, learn, and conduct research, but important research challenges in the field of education remain. In addition to research to meet the scalability and reliability requirements for information infrastructure, improvements are needed in the software technologies to enable development of educational materials quickly and easily and to support their modification and maintenance. We know too little about the best ways to use computing and communications technology for effective teaching and learning, in particular, how to effectively use multimedia capabilities to create a richer, and more appealing learning experience. We need to better understand what aspects of learning can be effectively facilitated by technology and which aspects require traditional classroom interactions with the accompanying social and interactive contexts. We also need to determine the best ways to teach our citizens the powers and limitations of the new technologies and how to use these technologies effectively in their personal and professional lives.

1.4 Transforming the Practice of Health Care

Vision: Telemedicine applications are commonplace. Specialists use videoconferencing and telesensing methods to interview and even to examine patients who may be hundreds of miles away. Computer-aided surgery with Internet-based video is used to demonstrate surgical procedures to others. Powerful high-end systems provide expert advice based on sophisticated analysis of huge amounts of medical information. Patients are empowered in making decisions about their own care through new models of interaction with their physicians and ever-increasing access to biomedical information via digital medical libraries and the Internet. New communications and monitoring technologies support treatment of patients comfortably from their own homes.

Future requirements for electronic medical records and health-system intranets will lead to increased reliance on the national infrastructure for communications, data sharing, and direct provision of care at a distance. Privacy and knowledge repositories are important research topics. Research in user interfaces is needed to understand how to make telemedicine applications not only efficient, but also satisfy the more human needs of both patient and physician. Robotics and remote visualization methods, supported by high-reliability, low-latency communications, are needed to support applications such as telepresence surgery.

1.5 Transforming the Nature of Commerce

Vision: Any company can be easily reached by its customers, regardless of location. It can receive immediate customer feedback, and rapidly adjust marketing strategies, prices, or product inventories based on that feedback. Consumers can shop for the best products, services, and prices from the convenience of their hotel room, home, or office. Electronic purchases can be made securely and with total privacy, providing suppliers and retailers with immediate access to cash generated by sales and consumers with automated statements detailing spending and purchases that allow for improved personal financial management.

Electronic communication is already dramatically changing how commercial transactions between companies are conducted, how digitally based goods and services are distributed, and how retail sales are made. Companies are using information technology to get closer to their customers and suppliers. Technology is also helping to reduce paper work and purchasing costs by streamlining the acquisition process and allowing companies to more efficiently find the best suppliers. Privacy and security are critical research topics if electronic commerce is ultimately to earn the full confidence of consumers. The reliability of the communication networks, computers, and business applications are vital to the success of U.S. companies. As the marketplace becomes increasingly global, understanding electronic commerce in the context of international trade relations will also become increasingly important.

1.6 Transforming the Nature of Work

Vision: The workplace is no longer confined to a specific geographic location, as workers can easily access their tasks and colleagues from alternate locations or while en route. Workers can do their jobs without regard to physical proximity to major metropolitan areas. They can choose where they live based on nearness to family or lifestyle preference rather than job market opportunities. A highly flexible workplace is able to accommodate each individual's needs, from working parents to workers with disabilities.

By some projections, as many as 15 million U.S. workers will become telecommuters over the next decade. This should enhance productivity and organizational flexibility as well as provide environmental benefits. To support large numbers of workers in non-traditional office settings, including the rapidly growing number of home businesses, we will need high-speed networking capability, equally available to many workers regardless of location or disability. Software technologies that allow work teams to collaborate effectively will be needed, and the privacy and reliability of the information infrastructure that permits this collaboration will be critical to success. The social and economic implications of telecommuting need to be studied. Computing and communications are also dramatically altering the skill base that workers need to perform their jobs. We need to determine how both employers, employees, and the self-employed can respond effectively to these changes.

1.7 Transforming How We Design and Build Things

Vision: Products and structures, from the highly complex, such as automobiles and buildings, to those used in everyday life, such as consumer appliances and fashion, can be designed with computer simulations that accurately represent the physical properties of the systems being built. Designers, manufacturers/builders, suppliers, and end-users participate in the design process, providing one another immediate feedback. Multiple designs and manufacturing processes can be rapidly explored yielding safer products, higher quality, and lower costs.

Global competition continues to press United States manufacturers to attract new customers and retain current customers by increasing productivity, reducing cost, improving quality, maintaining

maximum flexibility, and reducing design cycle time. Information technology has revolutionized the entire product development design cycle and will continue to do so. High-end computing technologies are needed for concept design, simulation, analysis with interactive control and computation steering, the mining of archived data, and the rendering of data for display and analysis. There is a critical need to link engineering development processes with business processes like planning, purchasing, scheduling, investment, and cost management. There is also a need for networked computers that allow simultaneous modification of a standard product to meet customers' needs. As computer prices continue to decrease, these design, engineering and manufacturing capabilities can be applied to many more products.

1.8 Transforming How We Conduct Research

Vision: Research is conducted in virtual laboratories in which scientists and engineers can routinely perform their work without regard to physical location—interacting with colleagues, accessing instrumentation, sharing data and computational resources, and accessing information in digital libraries. All scientific and technical journals are available on-line, allowing readers to download equations and databases and manipulate variables to interactively explore the published research.

High-speed computers and networks are enabling scientific discovery across a broad spectrum—from mapping the human brain to modeling climatic change. Research problems are becoming more complex and interdisciplinary in nature. As a result, researchers are finding innovative ways to collaborate with their colleagues across the globe. Key research technologies include high-end computing to allow higher fidelity models of complex physical phenomena, advances in collaborative environments, visualization of complex datasets, data mining techniques and management of very large datasets and databases.

1.9 Transforming Our Understanding of the Environment

Vision: Information technology can help us to improve a variety of problems from water and air quality to controlling the effects of toxic material. For example, reliable climate models permit us to determine the rate and regional distribution of climate change to support accurate projections by sector and region. Sophisticated models accurately predict the response of ecosystems to changes in temperature, water availability, and atmospheric composition. Fully integrated models allow scientists and policy makers to consider information on climate trends, population trends, resource utilization, and the value of natural and economic resources when making decisions regarding technically feasible and cost-effective options to reduce environmental impacts or adapt to climate change.

To better support national and international energy and environmental policy, the United States requires an unprecedented acceleration and extension of research into climate modeling in order to improve the accuracy of local and regional forecasting. Progress in this area depends on improvements in computational methods. This will require orders of magnitude increases in computing capability to deal with the immense size of these problems in both time and space. We also need

other advanced information technologies such as improved numerical methods and algorithms, tools for data storage, management, analysis and visualization, software development and testing techniques, and advanced networks for distributed computing.

1.10 Transforming Government

Vision: Government services and information are easily accessible to citizens, regardless of their physical location, level of computer literacy, or physical capacity. Intelligent systems guide citizens by providing a one-stop shopping experience for locating requested information. Documents and forms can be accessed, completed, and submitted electronically. Automated business processes allow nearly instantaneous response to citizens' requests. In times of natural emergencies, emergency crews have instant access to three-dimensional building models, risk analysis and assessment, high-resolution local weather predictions, stress analyses of damaged structures, rapid evacuation planning tools, and emergency agency coordination.

There is a huge potential to make all government institutions both more efficient and more responsive through information technologies. Technical challenges include significant improvements in systems and methods for accessing data, including high performance data storage and tools to locate and present information. Robust, reliable, and secure networks and software to deliver and protect critical information are important research topics. It is imperative that improvements in government be available to all citizens, so we must understand and surmount barriers to access.

Transforming the Way We...	Challenges	Benefits
1. Communicate	<ul style="list-style-type: none"> • Scaling for growth and reliability a la the telephone system. • Improving human interaction with computers. • Fragility of systems. • Global networking issues. • Finding best use of new communication possibilities, 1-on-1 and in groups. 	<ul style="list-style-type: none"> • One billion users can access the Internet simultaneously, regardless of language and physical limitations.
2. Deal with Information	<ul style="list-style-type: none"> • Improving data access methods. • Multi-modal human-computer interaction technologies. • Reliability and bandwidth, better audio and video streaming. • Scalable software support. • High-performance computing. • Delivering and protecting critical information. • Policy for electronic dissemination of information. 	<ul style="list-style-type: none"> • Everyone can access, query, and print any book, magazine, newspaper, video, data item, or reference document, regardless of language, using mouse, touch screen, speech, or eye blink. • Value is added to information through networked and software-enabled tools.
3. Learn	<ul style="list-style-type: none"> • Scalability and reliability of the information infrastructure. • Improving software technologies for development of education materials and support of their modifications and maintenance. • Determining the best use of computing and communication technology for effective teaching and learning. • Learning which traditional teaching methods to leave alone. • Learning how to teach citizens best use of these new technologies. 	<ul style="list-style-type: none"> • Regardless of location, age, handicaps, or schedule, anyone can participate in on-line education programs. • Everyone can access educational materials to discover the best learning style for them. • Customized educational programs exist for everyone, so no one is left behind.
4. Conduct Commerce	<ul style="list-style-type: none"> • Having sufficient privacy and security to ensure consumer confidence. • Reliability of communication networks, computers, and business applications needs to be high. 	<ul style="list-style-type: none"> • Customers can reach any company regardless of location. • Immediate feedback facilitates fast adjustment of marketing strategies and inventories. • Consumers shop at their convenience. • Companies can immediately access funds from sales. • Consumers have automated statements permitting improved financial management.

Transforming the Way We...	Challenges	Benefits
5. Work	<ul style="list-style-type: none"> • Developing high-speed networking for all, regardless of location or handicap. • Developing software to allow effective collaboration. • Ensuring privacy and reliability of the information infrastructure. • Determining how employers, employees, and the self-employed can respond to changes. 	<ul style="list-style-type: none"> • Workers have access to jobs regardless of proximity to population centers. • Workers can live where they want, not needing to be near jobs. • Workplace can better accommodate individual needs.
6. Practice Health Care	<ul style="list-style-type: none"> • Ensuring privacy of information repositories. • Developing robotics and remote visualization methods to support applications such as telepresent surgery. 	<ul style="list-style-type: none"> • Doctors use teleconferencing and telesensing to interview and examine patients. • Surgical procedures can be demonstrated with Internet-based video. • High-end systems provide expert advice. • Patients access biomedical information, gaining empowerment to make decisions.
7. Design and Build Things	<ul style="list-style-type: none"> • High-end computing technologies are needed for concept design, simulation, analysis with interactive control and computation steering, mining archived data, and rendering of data. • Need bi-directional engineering development processes linked with business processes. 	<ul style="list-style-type: none"> • Complex designs done via computer simulations. • All parties, including end users, participate in the process. • Safer products, higher quality, lower costs.
8. Conduct Research	<ul style="list-style-type: none"> • Research problems have become more complex and interdisciplinary. • Researchers need to find innovative ways to collaborate. 	<ul style="list-style-type: none"> • Research is conducted in virtual laboratories, interacting, accessing instrumentation, and sharing data and other resources, all regardless of physical location. • All journals are available on-line.
9. Deal with the Environment	<ul style="list-style-type: none"> • To accelerate and extend climate modeling research to improve forecasting. • Increasing computing capability by orders of magnitude. • Develop improved numerical methods and algorithms, tools for data storage, management, analysis and visualization, software development and testing, and advanced networks for distributed computing. 	<ul style="list-style-type: none"> • Reliable climate models. • Ecosystem models accurately predict responses to changes in conditions. • Fully integrated models facilitate decision making by scientists and policy makers.
10. Conduct Government	<ul style="list-style-type: none"> • Develop significant improvements in data access: high performance file systems and tools. • Develop reliable, secure networks and software to deliver and protect critical data. 	<ul style="list-style-type: none"> • Government services and information are available to all regardless of location, computer literacy, etc. • One-stop shopping for locating information. • Automated business processes accelerate responsiveness. • Enhanced responsiveness to natural disasters.

2. Setting Federal Research Priorities: Findings and Recommendations

Government-funded information technology research has produced enormous innovation. The results have been made readily available to industry. Testbed activities involving academia, government research facilities, and industry have served as powerful engines for technology transfer into the private sector, for the benefit of industry, government, and the Nation.

Today we enjoy the economic, strategic, and societal benefits of well-placed investments in long-term, wide-ranging information technology research begun during the Eisenhower and Kennedy Administrations. These modest investments have yielded massive economic benefits to the Nation. The empirical evidence is unequivocal: today's information technology sector vastly outstrips the current growth of all other sectors of the economy. The Federal Reserve reports that during the past five years production in computers, semiconductors, and communications equipment quadrupled at a time when total industrial production grew by 28 percent. These three industries account for one-third of the total real growth in production since 1992.

No other sector contributes nearly as much to the growth of our economy. The businesses spawned by these technologies employ millions of Americans in manufacturing and information processing jobs that pay wages well above the national average.

History suggests that, to be successful, Federal research investment must be sustained and flexible. Federal policies must support, encourage, and help coordinate long-range technological development. Federal R&D programs must be well designed and must not subsidize activities best left to the private sector. Only in this way can the Federal investment spur those critical areas of technology that either industry neglects or the Government overlooks in the normal course of business because they cut across Federal agency missions.

2.1 Findings

Finding: Total Federal information technology R&D investment is inadequate.

The non-inflationary growth in high-technology businesses fuels U.S. world leadership and creates whole new industries with hundreds of thousands of high-paying jobs. Furthermore, information technology's importance goes far beyond its economic benefits: it is fundamental to the solutions of many nationally important problems. But the Federal R&D investment has been compromised by a shift toward applied R&D. The Committee further finds that the amount of Federal research investment in information technology has not kept pace with information technology's growing economic, strategic, and societal importance to the Nation.

The number of proposals in computer science and information technology competitions provides evidence of the shortage of research funding. For example, the recent competition in Knowledge

and Distributed Intelligence sponsored by NSF drew over 1100 letters of intent and over 850 full proposals even though it had been announced that at most seventy-five projects could be funded. The response to the ASCI level 2 centers competition was similar. The reason for the extraordinary number of responses is that funding for information technology is extremely tight. Researchers are forced to participate in nearly every competition for which they might qualify. The end result is that researchers are spending increasing percentages of their time in proposal writing to the detriment of research itself and yet many good ideas are still not being funded.

This trend is not just bad for information technology researchers, it is bad for the Nation. Our ability to produce reliable software, build an information network on which the Nation can run, produce the high-end computing systems needed for advanced science, engineering, and defense tasks, and understand the socioeconomic effects of these revolutionary technologies is threatened by inadequate investments in research and development. By neglecting research, we do more than deplete the stock of fundamental knowledge—we endanger the long-term effectiveness of the entire R&D system and threaten U.S. leadership in the emerging 21st-century information economy.

Finding: Federal information technology R&D is excessively focused on near-term problems.

The NSF defines basic research as the study of the “fundamental aspects of phenomena and of observable facts without specific application toward processes or products.” In contrast, “applied research” is aimed at determining the means to meet specific needs, and “development” is defined as the systematic use of knowledge to produce useful materials, devices, or methods. These definitions are widely used in Federal policy and budget accounting, but they accurately describe only part of the nature of research and development (R&D). R&D is a complex non-linear interaction between concepts and theories, data and experiments, and new products and processes. Basic research is a critically important part of this interwoven system.

During the past decade both industry and Government have altered the balance between basic research and the later stages of technology development and commercialization. At the same time, major corporations have cut back on basic research expenditures, shifting staff from centralized laboratories to operating divisions where applied work is closely tied to commercial products and processes. In both the public and private sectors, the interacting reasons are 1) downward budget pressures, 2) increased focus on mission, and 3) the inefficiency of transitioning long-term research to near-term product.

Although total Federal spending for R&D has remained steady, there has been a marked shift toward support for applied R&D. One example is that DARPA, which funded much of the innovative research in the 1980's, revised its priorities in the 1990's so that all information technology funding was judged in terms of its benefit to the warfighter. In the process of making this change it decreased the time horizon for potential technology transfer. In the early 1990's, DARPA split the Information Systems Office off from its Information Technology Office to address specific military systems. Nevertheless, total funding for basic research in the DARPA Information Technology

Office is less than \$20 million out of a total office budget of more than \$200 million, an inadequate investment in our judgement. DARPA's new leadership plans to reverse these changes, but recent history shows that Government research managers correctly favor their agencies' missions when budgetary pressures grow and they have to choose between long-term research and short-term mission needs.

Most R&D investment restructuring in the early 1990's was essential for industry and the Government to maintain a competitive footing in the global marketplace and to maintain readiness against our present and future adversaries. However, this restructuring came at a high price: a serious decline in basic research activities. Research in computer science is a good example. In 1995, by the Federal Government's own measure,¹ more than two out of every three Federal dollars spent on research in computer science was for applied work.

The Committee finds that the Federal agenda for information technology R&D has moved too far in the direction of near-term applications development, at the expense of long-term, high risk, fundamental investigations in the underlying issues confronting the field. Currently, too high a percentage of research funding comes from mission-oriented agencies whose main goal is not so much to advance the knowledge-base in information technology as it is to solve the immediate problems confronting those agencies. Meanwhile, information technology has begun moving much faster than other fields. In order to intercept tomorrow's challenges we must set long-range, stretch goals rather than focusing on incremental advances. Within information technology, the scalable information infrastructure area is moving extremely fast, making it even more important to stretch for long-term research objectives.

Economic growth and defense leadership, if based on evolutionary improvements to yesterday's research results, are not sustainable—the rest of the world is not standing still in seeking economic advantage from new information technologies. By the Committee's own calculation, basic research spending may be as low as five percent of the total Federal information technology R&D budget. As a result, promising long-term research is being passed over in order to meet the goals of near-term technology development.

It is time to swing the pendulum back in the other direction and to strike a proper balance. We need more basic research—the kind of groundbreaking, high-risk/high-return research that will provide the ideas and methods for new disciplinary paradigms a decade or more in the future. We must make wise investments that will bear fruit over the next forty years.

Finding: The Federal information technology R&D funding profile is incomplete.

The current funding portfolio is not properly balanced. It is deficient in the support of multiple-institution, long-duration projects. Funding for projects of longer duration and larger scope is critical to the Federal research program. Projects of larger scope allow for multiple-investigator, inter-

¹ National Science Board, *Science and Engineering Indicators — 1996*. Washington, DC: U.S. Government Printing Office, 1996.

disciplinary collaboration, intramural research in academia and Federal research institutes, and joint industry-Government-academia experiments or proofs of concept. Projects of longer duration allow exploration of research problems with multiple-year horizons, which may lead to unexpected and significant discoveries.

It is important that Federal investments include a range of complementary funding modes, including classical single principal investigator (PI) research, multiple-PI experimental research, and multiple-institution/multiple-year efforts. Such diversity in funding approaches and tactics is important. It provides complementary modes for research, ensuring a broad perspective in addressing problems, thus increasing opportunities for discoveries.

2.2 Recommendations for Research

The current boom in information technology is built on basic research in computer science carried out more than a decade ago. There is an urgent need to replenish the knowledge base. Our legacy to future generations should be based on an intelligent, well-planned, highly disciplined investment in information technology that is commensurate with its role in the Information Age.

Recommendation: Create a strategic initiative in long-term information technology R&D.

The Advisory Committee recommends that the President create a strategic initiative to support long-term research in fundamental issues in computing, information, and communications. The initiative should endeavor to increase the total funding base to over two billion dollars per year after five years. The Federal funding agencies should use the resulting budget increases to encourage research that is visionary and high risk. The goal should be to recapture in the universities and research labs much of the excitement that existed at top-rated departments in the past.

The Administration's proposed Federal budget for FY2000 demonstrates a commitment to sustained growth in IT research through its Information Technology Initiative (IT²). This commitment is an important first step in what must be a continuing effort on the part of the Federal government to increase research dollars and to create a new management system designed to foster innovative research. But the effort cannot stop here. Further increases and continued oversight are needed to remedy the shortfall in long-term research investments that have accrued. In order to maintain U.S. leadership, 21st Century budgets should be guided by the principle that Federal support for information technology must be increased to bring it in line with its prominence in the economy and its importance to solving critical societal problems.

Recommendation: Increase the investment for research in software, scalable information infrastructure, high-end computing and socioeconomic issues.

Four areas of the overall research agenda particularly need attention, and must be a major part of the strategic initiative:

- *Software.* The science and methods for efficiently creating and maintaining high-quality software of all kinds, for ensuring the reliability of the complex software systems that now provide the infrastructure for much of our economy, for improving the interaction between people and computer-based systems and devices, and for managing and using information.
- *Scalable Information Infrastructure.* Techniques for ensuring that the national information infrastructure—communications systems, the Internet, large data repositories, and other emerging systems—is reliable and secure, and can grow gracefully to accommodate the massive numbers of new users and applications expected over the coming two decades.
- *High-end Computing.* Continued invention and innovation in the development of fast, powerful computing systems and the accompanying communication systems needed to implement high-end applications ranging from aircraft design to weather and climate modeling.
- *Socioeconomic issues.* Research on understanding the effects of information technology on our society, its economy, and the workforce should be funded. Furthermore, research should be focused on strategies for ameliorating information technology’s potentially harmful effects and amplifying the benefits.

The Committee recommends increasing investments in research in these specific areas, to meet critical national economic and defense needs and maintain U.S. leadership, as detailed in chapter 3.

Recommendation: Fund projects of larger scope and duration.

We stand at the dawn of a new century, a century where leadership in information technology may well be economically and militarily decisive. To meet this 21st century challenge, we believe we must diversify the modes of research supported by the Federal government. Computer science and information technology are collaborative research fields. Hence, special emphasis should be placed on involving and supporting researchers at many institutions in large-scale research projects that can explore technologies farther into the future with teams of researchers that may be interdisciplinary and multi-institutional.

Chapter 5 of this report discusses ways to diversify research support to accomplish this goal:

- Expand support for research carried out by teams of two to five researchers, possibly at different institutions, to address a single research project.
- Fund large centers for “Expeditions to the 21st Century,” which would involve large teams of researchers in explorations of future information technologies and their societal effects, and
- Establish “Enabling Technology Centers,” which would conduct research on the application of information technology to particular problems of national importance.

These approaches would allow the Government to adapt its funding mechanisms appropriately to critical research goals.

3. Technical Research Priorities

3.1 Software Research

Software is the new physical infrastructure of the information age. It is fundamental to economic success, scientific and technical research, and national security. Software is increasingly important for commerce, communication, information access, and the Nation's physical infrastructure. The continuing emergence of cheaper and faster microprocessors has allowed more and more functions to be performed by software.

While we may think of the phone system, the Internet, or even a camera or a car as devices, these devices cannot function without software. We rely on software to work without fail, to be modifiable as requirements change, and consistently to provide more functions with better performance.

However, because the demand for software has grown at such an explosive rate, demand now far exceeds supply. Furthermore, the Nation depends on software that is often fragile, unreliable, and extremely difficult and labor-intensive to develop, test, and evolve. Our ability to construct the software to manage information and our ability to analyze and predict the performance of the enormously complex software systems that lie at the core of our economy are painfully inadequate. We are neither training enough professionals to supply the needed software² nor adequately improving the efficiency and quality of our construction methods. The recent Government emphasis on short-term applied R&D has hurt software research, reducing the generation of new ideas and approaches that would address the demand.

3.1.1 Findings

Finding: Demand for software far exceeds the Nation's ability to produce it.

The explosive growth of information technology has fueled an unprecedented demand for new software. Software is needed to support millions of legacy systems as well as new products, to provide fundamental services on the Internet, and to solve important national problems. The resources needed to develop this software have not kept pace with the demand, producing what might be called a "software gap."

An interaction of factors has caused this software gap: accelerated demand for software, increased complexity of systems, labor-intensity of development, variable quality in the labor pool, labor shortages, and lack of adequate science and technology to support robust development. Since hardware improvements create demand for more software, the demand will continue to accelerate. Today's systems and applications software are substantially more complex than in the past. Software systems are now among the most complex human-engineered structures.

² Bureau of Labor Statistics. "Employment Outlook 1996-2006." *Monthly Labor Review*, November 1997.

This situation threatens to inhibit the progress of the current boom in information technology. It may also threaten the health and welfare of the Nation by reducing the rate at which solutions to software-intensive problems, like aviation safety and crisis management, can be solved.

The Federal Government must take steps to address the situation by investing in software research. Note that “software research” means more than just research on new strategies to improve software development productivity. Software research also includes research into new tools to manage information and support commerce, development of new software to make computers easier to use, creation of software for embedded devices and new applications, and research to make software systems interoperate more effectively. It also means more than just investing in refining standard approaches to software development. It means funding well-defined but radical approaches that are high risk but could have a major impact. It will require conceptual breakthroughs to invent new, “outside-the-box” approaches. The Federal Government should also help increase the pool of information technology professionals capable of developing good software.

Finding: The Nation depends on fragile software.

The Nation needs robust systems, but the software our systems depend on is often fragile. Software fragility is its tendency not to work properly—or at all. Fragility is manifested as unreliability, lack of security, performance lapses, errors, and difficulty in upgrading.

Examples can be found everywhere, from our huge information systems for air-traffic control to the personal computers on our desks, from the Pentagon to the Internal Revenue Service (IRS). One aspect of our software’s lack of security is evidenced by the fact that teenage hackers have easily penetrated our telephone and military systems; in reality the software-security problems we face go far beyond calculated external attacks. The Federal Aviation Administration (FAA) and IRS systems have proven difficult to upgrade. Large telecommunications networks have crashed, and banks have been robbed electronically. Even after large, expensive testing efforts, commercial software is shipped riddled with errors (“bugs”). Software producers rely on their users to discover the remaining errors in actual use, making it even more likely that our systems will fail.

Not only are our products inadequate—so are our processes. The Standish Group³ reports that 73 percent of software projects are late, substantially over budget, canceled, or outright failures. The well publicized, potentially catastrophic “Year 2000 problem” is an example of known unreliability and illustrates the difficulty and cost of upgrading to meet new circumstances. Companies spend enormous amounts to back-up their on-line data in order to compensate for the risk that their own systems will corrupt it.

³ The Standish Group International, Inc., is a market research and advisory firm specializing in mission-critical software and electronic commerce. They are located in Dennis, Massachusetts.

The Nation cannot afford to let the current situation continue. We must commit to develop the science, technologies, and methods needed to build robust software systems—ones that are reliable, fault-tolerant, secure, evolvable, scalable, maintainable, and cost-effective.

Finding: Technologies to build reliable and secure software are inadequate.

During the past 40 years, computing hardware has seen an increase in performance of at least eight orders of magnitude. Our ability to develop software has not kept pace with the opportunities those hardware advances provide. Our ability to construct the systems we require and our ability to anticipate the performance of the enormously complex software systems that lie at their core are inadequate.

Large software systems are beyond our capability to describe precisely. Consequently, there is little automation of their construction, little re-use of previously developed components, virtually no ability to perform accurate engineering analyses, and no way to know the extent to which a large software system has been tested.

Having meaningful and standardized behavioral specifications would make it feasible to determine the properties of a software system and enable more thorough and less costly testing. Unfortunately such specifications are rarely used. Even less frequently is there a correspondence between a specification and the software itself. Often software behavior and flaws are observable only when the program is run, and even then may be invisible except under certain unusual conditions. Programs written in such circumstances frustrate attempts to create robust systems and are inherently fragile.

Software development relies on individual genius and creativity, and, as with all design-based disciplines, will continue to do so. But it has become clear that the processes of developing, testing, and maintaining software must change. We need scientifically sound approaches to software development that will enable meaningful and practical testing for consistency of specifications and implementations. This requires long-term research in languages, theories, simulation, analysis, and testing that could lead to standardized multilevel mechanisms similar to those which have created the success in computer-aided design for digital hardware.

The construction and availability of libraries of certifiably robust, specified, modeled, and tested software components would greatly aid the development of new software. Libraries of software components for developing business applications are beginning to be constructed commercially, but only in limited circumstances.

Finding: The diversity and sophistication of software systems are growing rapidly.

The enormous advances in computing capability, the integration of computing and communications, and the striking reductions in the costs of those technologies have increased the opportuni-

ties for new and different software-based applications. The algorithms and the information structures needed for these new application domains require increased research in order for the applications to flourish. For example, real-time control is essential for transportation systems, for observational input to weather and climate systems, and for telemedicine. Scene recognition systems play an important role in national security. Ubiquitous computing and collaborative decision-making have very broad and important military and civilian uses. Whether the purpose of software is systems-oriented (e.g., an operating system, a compiler, or a protocol for wireless communication) or applications-oriented (e.g., a billing system, a speech interface, or a simulation), the software technology must continue to advance.

Finding: More and more often, common activities of ordinary people are based on software.

Twenty years ago, most computer use was job-related. Today, people use computing services all the time—for banking and financial matters, for entertainment, to arrange travel, to correspond with friends and family, to reach other people by phone, and so forth. Not only must the services they use be reliable and available, but they also must be sufficiently easy to use, both physically and intellectually. As software processes replace manual activities, important characteristics of the older processes must be preserved (three examples: historical records must be preserved, errors and omissions must be tolerated, and it must be possible to get assistance and to pause a partially completed activity and resume later).

Computers provide communication and information access that previously required specialized devices, such as telephones, music players, televisions and VCRs. Making these devices more intelligent promises to make user interfaces much richer and more pleasant to use. These new devices will allow us to use all our senses when working with computer applications much as we do in the real world. Since computers are much more complicated than consumer devices, we should not introduce complexity as we integrate these new modes of interaction into computer user interfaces. For example, playing music with a computer or new computer-based client device should be simpler and better than playing music with a CD player.

The needs of inexperienced users are not well understood. New modes of interaction must be developed, along with a better sense of what constitutes ease-of-use. More research is needed in human-computer interfaces and in human-centered information management. Using computers—particularly computers embedded in devices such as teller machines, “smart” phones, electronic kiosk systems, and the like—must be natural and instinctive. Additionally, the integration of separately developed services must appear seamless to the user. Research in middleware (the standard interfaces and standard services that enable interoperability, including those supporting distributed computing), in scripting of complex applications, and in overall systems integration is required.

Finding: The Nation is underinvesting in fundamental software research.

Over the past decade the U.S. has underinvested in research to create fundamentally new software technologies.⁴ Most of today’s software technology is based on research performed 15 or more years ago. If we fail to invest in far-reaching, high-risk research now, on what ideas will the commercial advances of the year 2015 be based?

In recent large Federal information technology initiatives, the Committee found systematic underinvestment in the software side of the research agenda, leaving a growing gap between the theoretical potential and actual performance of hardware. The Federal High Performance Computing Program (HPCC) plan, for example, stated, “advances in software will be critical to the success of high performance computers with massively parallel architectures.”⁵ HPCC researchers, however, were encouraged to focus on the immediate use of their results on Grand Challenge problems or in semi-production parallel computer systems. This successful, short-term strategy emphasized software to solve particular scientific or systems problems, but it failed to develop the fundamental software technology needed to solve wider classes of problems such as interoperability across different computing architectures.

Additionally, underinvestment in long-term software research has interrupted the pipeline that produces software research professionals. The Committee strongly believes that increases in research funding are needed to adequately develop the next generation of promising young researchers.

3.1.2 Recommendations

Software is increasingly important to the fabric of our society. By failing to improve the quality of the software we develop, and by failing to provide adequate tools to manage information, we put the Nation at risk. Long-term research is needed to strengthen our software enterprise. That research is not being adequately supported. This research should be closely coordinated with research undertaken in response to the President’s Commission on Critical Infrastructure Protection designed to protect the Nation’s infrastructure from intentional attacks.

The Committee recommends that a variety of additional investments be made to enable fundamental improvements in the Nation’s software quality and its development processes. In particular, major improvements must be made to methods for software development, verification and validation, maintenance, user interfaces to computing systems and electronically represented information, software for high-end computing, and software to support emerging ubiquitous and collaborative computing.

⁴ Council on Competitiveness. *Endless Frontier, Limitless Resources*. Washington, DC. April 1995: 112-113.

⁵ Executive Office of the President, Office of Science and Technology Policy. *The Federal High Performance Computing Program*. Washington, DC. September, 1989.

Technology for building software is only part of the solution, since we also need to know what software to build. Current computer systems are too difficult to use. New theories and techniques to simplify the way people interact with information appliances are urgently needed. This will involve new techniques for seeing, hearing, and understanding the environment, techniques for representing and organizing this information within computer systems, and new techniques for analyzing, summarizing, and explaining information.

Chapter 1 described how information technology will transform science, medicine, commerce, government, and education. These transformations will require new algorithms, new tools, and new ways of using computers. Research ideas and prototypes will fuel the transformations. The huge demand for commercialization of current prototypes is draining the pipeline of talented university faculty and researchers—essentially eating the seed corn. The Committee recommends expanding and refocusing university research funding to address longer-term research issues.

Major Recommendation: Make fundamental software research an absolute priority.

The Committee recommends that the Government make fundamental research in software both for computer systems engineering and for applications one of the Nation's highest R&D priorities. We recommend that a focus on software research be a mandatory element of all the Expeditions proposed in Chapter 5. In addition, significant added research funding within current modes of funding should be focused on fundamental software research.

Recommendation: Fund more fundamental research in software development methods and component technologies.

The Committee recommends that research in software methods be aggressively pursued, especially in the area of automated support for software development and maintenance. Such research should:

- Explore and create component-based software design and production techniques, and the scientific and technological foundations needed for a software component industry.
- Explore and create theories, languages and tools that support automated analysis, simulation, and testing of components and their aggregation into systems.
- Create a national library of certified domain-specific software components that can be reused by others.
- Explore and create techniques for using measurably reliable and secure components and their aggregation into predictably reliable and secure systems.
- Explore and create protocols, languages, and data structures to promote interoperability of applications running concurrently across wide-area networks.

Recommendation: Support fundamental research in human-computer interfaces and interaction.

The purpose of many computer applications is to inform people. Moreover, most computer applications require information and guidance from us. Yet the interaction between people and computers is still primarily limited to very rudimentary actions and outputs, such as keyboard strokes, simple gestures using pointing devices such as a mouse or touch-sensitive screen, and textual and simple graphical output on displays. People are capable of far richer interaction. More to the point, humans are perceptual creatures, using all their varied sensory systems in concert to form a harmonious interpretation of the environment. The sensory-motor systems of the human brain are tightly integrated, yet in our computer programs, computer inputs (from human motor systems) and computer output (which go to human sensory systems) are treated as independent and somewhat unrelated activities.

People and computers both deserve better. Fundamental research in human sensory-motor systems, perception, attention, pattern recognition, and decision-making has the potential to make dramatic improvements in the interaction of people and machines.

Today, simulators for commercial aircraft, military systems and games provide the richest examples of what is possible. Full-immersion environments are beginning to explore these areas. However, without sustained, fundamental research into the possible interactions, we will not make much progress beyond today's limited systems. Moreover, these investigations will have rich implications: reducing errors and misinterpretations, enhancing complex control settings, whether in industrial process control, air traffic control, or military applications, as well as helping to overcome some of the access problems faced by people with sensory, motor, and perceptual handicaps.

There are many facets to human-computer interfaces and interactions. We must provide easy access for all people, regardless of economic circumstances, physical impairment, or intellectual limitations. To that end, user interfaces must become considerably richer, depending less on textual interfaces and manual dexterity. They should take more advantage of the decreased cost and increased miniaturization of audiovisual devices, increase the use of natural language, and provide more assistance to the user through techniques ranging from help systems to inference. Progress in ease-of-use is a good example of the coupling of software advances with those of the underlying hardware devices.

Recommendation: Fund more fundamental research in information management techniques to (1) capture, organize, process, analyze and explain information and (2) make information available for its myriad uses.

We already have on-line access to vast quantities of information. Now the challenge is finding the right information quickly and easily, and using it effectively. Ideally, the information will come to us in anticipation of our needs. The Committee recommends increased funding for basic research

on capturing, organizing, processing, and using information. Information management is based on the classic computer science disciplines of new and better data structures and algorithms, but also includes theories and new approaches to digital libraries, databases, knowledge discovery, data visualization, and information-intensive applications. Software tools that augment our intelligence and increase our productivity will be key components of the Nation's prosperity in the future. The research in this area should:

- Extend both existing information technologies and those of the future so that non-textual information (pictures, audio streams, animations, etc.) is fully incorporated and integrated.
- Develop improved technologies for filtering information, for data mining, and for tracking the lineage and quality of information, while protecting privacy and intellectual property rights.
- Create digital libraries that can enhance the Nation's scientific and educational processes, and that include mechanisms for the long-term preservation of valuable information.
- Create new technologies for incorporating rich information management services into complex, distributed, demanding applications.
- Develop tools and technologies for modeling and understanding the semantics of information.
- Create tools that make it easier for people to manage, visualize, and exploit information and knowledge.

Recommendation: Make software research a substantive component of every major information technology research initiative.

In the past, major initiatives, such as HPCC, have suffered from underinvestment in software and information management research. When considering new initiatives like the ones proposed later in this Chapter, we must ensure that the software research necessary for success (not just the software development) is included.

For example, we cannot build a scalable information infrastructure without an adequate investment in the software that will provide both system-level and user-level services on that network. For this reason, the Committee recommends in Section 3.2 substantive new software investments to support scalability, distribution, and reliability.

Similarly, there is a continuing need to raise the level at which users of high-end computing create applications software and migrate their software solutions to new architectures. Thus an adequate investment in software to support those activities is called for in Section 3.3.

In short, Federal research support programs have a history of underestimating the software research and development investment needed for success. This tendency must be reversed if new information technology initiatives are to be as effective as the Nation needs.

Budget Recommendation: Increase current funding for software research as follows over fiscal years 2000-2004.

The recommended software research budget increases must account for two important factors. First, the need to increase the average funding per senior investigator by at least \$100,000 per year, these increases to start immediately. Second, based on the technical objectives in the above recommendations, the need to increase the number of researchers by approximately 1000 and to support that larger community. The latter increase must be phased in over the five-year period.

The following numbers are in addition to current (FY 1999) funding levels.

Funding increases for Software Research (\$ millions)					
	FY 2000	FY 2001	FY 2002	FY 2003	FY 2004
Software Engineering and Component Technologies, Human-Computer Interaction and Information Management	112	268	376	472	540

The Committee recommends that a significant portion of the new monies be used for thematic programs. Rather than establishing large centers with structured management, thematic programs should fund a diversity of groups that are small enough to be well-focused but are large enough to reach “critical mass.” Groups might have several cooperating senior investigators, together with staff, graduate students, and undergraduates. Funds must be awarded for long enough periods to address the truly difficult but important software problems described above.

3.2 Scalable Information Infrastructure

Over the next several decades, the transformations described at the beginning of this report will dramatically change our country. The Internet is at the heart of this transformation.

3.2.1 Findings

Finding: Our Nation’s dependence on the Internet as the basis for its information infrastructure continues to grow at a dramatic rate.

Within the next two decades, the Internet will have penetrated more deeply into our society than the telephone, radio, television, transportation, and electric power distribution networks have today. For many of us, the Internet has already become an integral part of our daily lives. Communications by electronic mail is commonplace, and many prefer it to the telephone or fax machine.

Many more people will soon depend on the Internet to communicate with friends, family, and businesses; for banking and financial transactions; to purchase goods and services; for entertainment; as well as for vital medical and governmental services. As we come to rely on using the Internet every day, and as billions of dollars of financial transactions are conducted electronically, this information infrastructure becomes even more critical to our Nation's well being. Our ability to manage and control it will become as important as our ability to manage our electricity distribution system. It is rapidly becoming an essential component of our economic and social infrastructure.

America leads the world in developing and applying information and communications technologies, and it is crucial that we continue to invest in them. At stake are the technologies that will determine our Nation's ability to sustain its economic well being, to compete successfully in the global marketplace, to maintain world leadership in basic and applied scientific research, and to provide for national security and defense.

Finding: The Internet is growing well beyond the intent of its original designers, to such an extent that we no longer understand it and cannot confidently continue to extend it.

First, the Internet is growing rapidly in size. A network that connected 2,000 computers in 1985 now connects 70 million devices and traffic is doubling in size every hundred days. More than 50 million Americans use the Internet today, and there will likely be one billion Internet users worldwide early in the next century. No one has ever built a network of such size that provides this level of access to information resources. The combination of common services, a common global namespace, and interactive access creates a new cultural environment on a global scale.

Second, the Internet is growing in capability. We are using it in more and more varied ways. Electronic mail servers are increasingly used for "anytime" communications. Satellite and wireless systems will soon provide many users with "anywhere" connectivity. World Wide Web directory and search services help users find a growing fraction of all the information known to humanity. Electronic payment services handle more of the Nation's commerce. Streaming audio and video is being shown and distributed. Building blocks for new applications are being developed, enhanced, and used: digital signatures, secure transactions, advanced modeling and simulation software, shared virtual environments for collaboration, intelligent agents for discovering and retrieving knowledge, speech recognition, and low-cost networked sensors.

Finally, extending the Internet poses challenges we may not find so easy to meet. We cannot confidently extend today's technologies to new networks and information systems that are orders of magnitude larger, more complex, and that can handle more types of information than today, while expecting to achieve the quality and reliability demonstrated by today's U.S. telephone systems. We understand how to build some, but not all, of the individual elements of the future infrastructure. We do not know how to make these elements work together in reliable, efficient, and robust ways. Even greater challenges will arise when we try to build scalable services, such as search engines that

can index exabytes of data, payment services that can process billions of transactions per day, and scalable applications, onto the existing infrastructure.

Finding: Learning how to build and use large, complex, highly-reliable and secure systems requires research.

The Internet-based information infrastructure of the future must provide high performance-reliable, robust, and secure connectivity and functionality among an increasing number of networks and computing systems. The combinatorial character of such highly capable networks of computing systems filled with disparate software defies easy comprehension. We need to analyze this large, complex system in order to understand its behavior better.

We need a better scientific and engineering base to build network links and computing system nodes that can handle radically enlarged and diversified scales of usage and traffic in the expanding information infrastructure. Many technical obstacles—including ease of management, ease of use, integration, privacy, reliability, and security—stand between the current state of the art and the global-scale, heterogeneous, ubiquitous (anytime, anywhere) infrastructure implied by having one billion Internet users.

Networking research must allow us to keep the capability of the Internet ahead of its unprecedented growth. Moore's Law, which has long and successfully guided the semiconductor industry by projecting when performance and technology requirements are needed, can suggest future Internet scalability and performance needs. But the rate of Internet growth is even faster than Moore's Law would predict, and telecommunications researchers will need to develop revolutionary rather than just evolutionary new ideas and theories in order to keep ahead of the Internet's growth curve.

To support the growing demand and dependence on the information infrastructure, research advances in scaling are needed in many different dimensions:

- Scaling to provide robust, reliable, high-speed access. This will improve the quality and richness of the content that users can enjoy and that the information infrastructure can support.
- Scaling to provide assured quality of service. These advances will improve the quality of a user's interaction with the information infrastructure through multi-model communications including speech, sight, or touch, and incorporating isochronous streaming data such as voice and video.
- Scaling to provide ubiquitous access. These advances will enable more people to have continuous access to information and services regardless of their location.
- Scaling of infrastructure services to handle users and requests reliably. These services include authentication, resource directories, search engines, and banking. These advances will improve the quality of information.

- Scaling of the security of the infrastructure to safeguard intellectual property rights, to protect against failures and attacks, and to provide privacy of access when needed. These advances will increase the trustworthiness of both the information and the infrastructure.
- Scaling to support huge servers that can service billions of requests per day against exabyte data stores. These stores need to always be available, must never lose data, and must be able to search the entire data store quickly.

This research should not stop at network creation, but encompass the operation of the complete end-to-end system. Applications research is needed to better understand how to take best advantage of these networks and their large-scale information servers. One example would be scientific and engineering applications that are most efficient when they use different types of computing systems distributed over a national network. Another research area is network-aware applications that can interact with network management software to improve both application and network performance. A third area focuses on designs for large high-traffic servers. A fourth research area is efficient management of data flow among distributed data stores and the applications that access those data. Managing the flow of information from original data through processing to publication and usage will create an information environment that is the core of our scientific, commercial, and governmental systems. We need to learn how to build and use such a system if we are to take full advantage of information technology's potential.

3.2.2 Recommendations

The following recommendations are listed in their logical order and touch upon all aspects of the scalable information infrastructure, ranging from analysis and modeling, to implementation through middleware and scaling, to deployment on testbeds. The Committee considers these recommendations to be of equally high priority in a balanced research program.

Recommendation: Fund research in the behavior of the global-scale network and its associated information infrastructure. This should include collecting and analyzing performance data as well as modeling and simulating network behavior.

It took more than 30 years for the telephone industry to develop useful mathematical models that allowed them effectively to configure, manage, and plan the future growth of telephone networks. Unfortunately, those models do not apply to the Internet, and indeed have become less useful in their own right as telephony has expanded to include fax and Internet communications. The telephone network, though not ideally suited for this task, is the principal method of accessing the Internet. However, the shortcomings of the telephone network are increasingly limiting our future deployment capability. Fundamental differences between telephony and the Internet include:

- The packet-switched Internet Protocol (IP) as opposed to the reserved end-to-end bandwidth of telephony.
- The wide range of applications that run over the Internet (from e-mail to the Web to teleconferencing) as opposed to voice-only telephony.

- The wide range of networking technologies (wireline, wireless, satellite, cable, broadcast) and vendors as opposed to the mostly twisted-pair technology of the then largely monopolistic telephone industry.
- The wide range of performance requirements of both networks and end-node devices, again as opposed to the simpler requirements of telephony.
- The considerably broader statistical distribution of characteristics such as service times compared to those found in conventional voice telephony.

We are not yet able to model this complex system, either at the level of an individual node or the behavior of Federal, commercial, or corporate subnets of the Internet. We must develop a new set of models. This will require a long-term effort in developing new data-collection techniques, mathematical and statistical modeling and simulation methods and tools, and analytical techniques. It will also require encouraging data sharing among network service providers.

Recommendation: Support research on the physics of the network, including optical technologies, wireless technologies including satellites, wired technologies including cable, and the related bandwidth issues.

Improvement in transmission technologies drives many advances in networking. Improving how fast, how far, or how clearly we can transmit a data packet has a profound impact on what the information infrastructure can do. These advances primarily require research in physics and signal processing: waveforms, signal quality, and manufacturing techniques. Particularly pressing issues surround optical, wireless, and local-loop transmission. Successful research today will help meet the Internet's demands for quality, speed, and ubiquitous access tomorrow.

Recommendation: Support research to anticipate and plan for scaling the Internet.

The most pressing problem in the Internet today is scalability. In whatever way it is measured, it is becoming larger and more diverse:

- Number of nodes: Millions of nodes and links are added each year, and it is not unreasonable to expect the Internet to have more than 100 billion nodes and links by 2025. These will include networked systems embedded in objects ranging from cars to homes to ordinary appliances.
- Number of users: Every person will be able to communicate with every other person and every node on the network virtually instantaneously.
- Geographical dispersion: the Internet is global.
- Heterogeneity of the networking environment (including wired such as optical, cable, etc; wireless, including satellite transports; symmetric and asymmetric two-way communications; and broadcast).
- Heterogeneity of its end nodes (from small embedded to large computers and data stores and other devices from a variety of vendors).

- New types of traffic such as voice, audio, video, or other streamed content.
- Availability, bandwidth, efficiency, reliability, security, and survivability.
- Networks with low to high latency characteristics, as well as the ability to manage them.
- New features such as differentiation of service.
- New uses such as information management.
- New applications such as telemedicine.

We can summarize the research challenges simply. How do we design and build systems and software for deployment on an Internet that is many thousands of times larger and more diverse than today's network? How do we test a system that cannot be built today? How do we manage such a system once it is deployed? How do we add new capabilities without degrading or interrupting the operation of the underlying system on which so much depends? How do we create a system that can evolve capabilities that we cannot imagine today? Answering these questions requires new mathematical and statistical theory dealing with new classes of data distributions, modeling, simulation, experimentation, and testbeds.

Recommendation: Support research on middleware that enables large-scale systems.

Middleware for the scalable information infrastructure is the collection of shared software and rules that will improve the software development process, help make the infrastructure operate properly and efficiently, and make large new software systems possible. It will provide reusable software that performs common functions, making software development more efficient. This will let applications software developers concentrate on their applications and not on data transmission or computing issues. It will enable the development of applications that communicate with each other. It will also enable the development of middleware that adapts to a changing environment through services that coordinate and supervise the infrastructure itself. Middleware research topics fall into the two broad categories:

Information management

Scalable and widely application information management services require that we build systems to organize and manage vast amounts of information. The needed technologies include those for:

- Managing exabytes of data and billions of requests for information each day.
- Managing text, semi-structured information, relational data, multimedia data, and active data.
- Translating between information representations.
- Discovering and monitoring resources.
- Analyzing, summarizing, integrating, and fusing information from diverse heterogeneous sources.
- Presenting information in meaningful ways.
- Tracking the lineage or provenance of information, so users can know the sources of original data and how information was derived.

- Tracking and controlling access to information, protecting privacy, and charging for information where appropriate.

Information and services survivability

Our Nation's security, commerce, education, and well-being depend increasingly on our information infrastructure. It is thus critical to ensure the survivability of that infrastructure in the face of malicious attacks or viruses, equipment or software failures, and overload. Survivability means that services will be available when needed and information will be delivered in a timely fashion.

Services must operate correctly and the information they deliver must be of high quality or be identified if otherwise. Survivability includes long term preservation of information; a document of historical importance should be preserved, even as underlying storage technologies and information representation evolve and even if it is not accessed for hundreds of years.

Survivability technologies include:

- Authentication and security mechanisms for a large, heterogeneous, and evolving infrastructure.
- Mechanisms for detecting system intrusion, and information and software corruption.
- Mechanisms for detecting, mitigating, responding to, and recovering from, or for preventing, human error in the creation and use of the infrastructure.
- Mechanisms for assuring information quality.
- Scalable information and service replication strategies.
- Mechanisms for monitoring services to ensure correct operation within given quality-of-service bounds.
- Repositories for guaranteed long-term preservation of information.

Recommendation: Support research on large-scale applications and the scalable services they require.

Very large applications have unique needs that today's technologies cannot meet. The process of developing these applications has been found to be difficult to manage, and product quality is difficult to engineer. Today's commercial-off-the-shelf technologies do not scale for use in these large applications. In addition, university-based researchers are not well equipped to work on these problems, which require substantial labor and infrastructure. Research focused on these scalability issues is needed.

Several of these applications come from Federal agencies with science and engineering missions, including the Department of Energy's Accelerated Strategic Computing Initiative (ASCI) for nuclear stockpile stewardship, NASA's Earth Observing System Data and Information System (EOSDIS — one of the most complex data handling systems ever to be built), and the FAA's air traffic control system. These applications combine communications-intensive, data-intensive, and computation-intensive requirements.

While these applications are small in number and, by today's standards, are large in scale, they will be quite common in 20 years. In the meantime, today's demanding applications will have grown in size and will continue to make scalability demands, so the research requires a long-term approach.

The technologies developed to address large-scale Federal science and engineering applications will be useful in addressing non-science Federal applications such as improving the capabilities and efficiency of large Federal entitlement programs, and in addressing non-governmental applications that are not being addressed today. We examine two such applications in some detail.

Application: A National Digital Library

A National Digital Library would integrate all electronic knowledge sources, including books, journals, music, films, and informal "documents." The Nation's scalable information infrastructure would manage this shared information base. It might begin with a digital library for the country's scientific literature, which while largely produced at taxpayer expense is not widely available even to scientists, engineers, and students. Such a library would have a strong positive impact on scientific and engineering research and development and on education by providing users with more up-to-date information, reducing effort and cost, and preserving this body of knowledge for future generations. The university-led Federally-funded Digital Libraries projects, begun in 1994, have developed a base upon which to develop this library. Some of the challenges a National Digital Library presents are:

- Mechanisms for protecting and managing intellectual property.
- Strategies for managing the data that describe the library's holdings and mechanisms for translating document formats.
- Technologies for multi-lingual interoperability, allowing users to access information in different languages.
- Interfaces that let users visualize and understand the wealth of available information.
- Digital journals with on-line collaborative authoring, submission, reviews, publishing, and reader annotation.
- Technologies for organizing large bodies of diverse knowledge (for example, clustering, summarization, and classification).
- Technologies for efficiently digitizing paper documents.

Application: The Next Generation World Wide Web

Today's Web and its browsers can be viewed as a primitive information infrastructure. It has evolved rapidly from a simple design that did not anticipate either today's size or the wide class of applications for which it is now used. Designing and developing a Web for the year 2010 is a research challenge that requires addressing capability, flexibility, power, robustness, and scalability needs and applications far different from today. While a Next Generation Web shares many research issues with information management and digital libraries described above, it has its own

unique needs. These include research and development in translating among the protocols and policies of different countries and providing a graceful transition path.

Recommendation: Fund a balanced set of testbeds and research infrastructure that serve the needs of networking research as well as research in enabling information technologies and advanced applications.

The Federal government has funded two testbeds as part of its Next Generation Internet (NGI) initiative. One connects more than 100 sites at end-to-end speeds 100 times faster than the Internet of 1997, and the other connects about 20 sites at end-to-end speeds 1,000 times faster than 1997's Internet. The former is used primarily for developing “revolutionary” applications and the latter for research in networking technologies. While interconnected, these testbeds have different characteristics.

The 100x testbed connects several Federal research networks, including NSF's vBNS, DoE's ESnet, NASA's NREN, and DoD's DREN. Each of these networks is a service provided by a U.S. company—MCI WorldCom, Inc., Sprint Communications Company, and AT&T, respectively. In addition to being used by researchers in almost 100 universities and major Federal research facilities to develop applications, the testbed provides a national-scale environment for experimenting with new capabilities such as multicast protocols and protocols enhanced to include differentiation of service, privacy and security. It also serves as a testbed for the interoperability, robustness, and scalability of networks provided by different vendors, and the management of such interconnected networks. The testbed uses networking technologies more advanced than those in the commercial Internet. These technologies have been made stable and robust largely through collaborative efforts by the Federal agencies and the providers. Users appropriately expect high availability and high performance from every NGI site nationwide, because their research depends on network availability and performance.

The 1,000x testbed is used to explore networking technologies for a future generation (beyond the next generation) Internet. It must stress the limits of the new technologies being developed. It is smaller than the 100x testbed, is more fragile, and may experience outages.

Information infrastructure facilities are increasingly serving a similar role for scientists and engineers that large-scale shared physical instruments have for decades. As is the case for national telescopes, microscopes, and particle accelerators, agencies like NSF and DoE provide peer-reviewed access to leading edge IT facilities to enable a wide range of disciplines to do research. It is clear that such facilities are only useful to maintaining our Nation's research leadership if they are continually improved to keep them at the state of the art. However, this need to continually improve specialized IT facilities, such as high-end computers, high performance networks, and large discipline databases, must be managed by the Federal agencies in a manner that does not directly conflict with funding the basic research that is enabled by the facilities.

We recommend a periodic review of testbed and research infrastructure activities. This will help ensure that funding and usage is consistent with the research missions for which they are designed. The participants in this review should include representatives from the Federal Large Scale Networking Working Group, the Federal agencies that provide these resources, user communities, and the appropriate private sector communities.

Recommended funding levels

Scalable Information Infrastructure is a relatively new and especially fast moving field, making it difficult to estimate its research and testbed needs. The large number of research topics explicitly identified above, including many topics never before or at best minimally addressed in the past, and the growing demand for new Internet capabilities, suggests that research in scalable technologies can be highly effective. The Federal government's responsiveness to the rapid evolution of this field supports this position. The Government has modified its plans (for example, by increasing the size of its testbeds to meet demand), accelerated its research schedule, and consistently addressed issues that the private sector will not address (such as interoperability across providers).

The considerations in proposing the scalable information infrastructure budget included the growth rate of the Internet and acknowledged cost savings such as in-kind contributions from the private sector. The Committee recommends that funding for scalable information infrastructure research and testbeds grow from the current approximately \$200 million per year to \$500 million per year in a linear fashion between FY 2000 and FY 2004. The emphasis should be on networking research. The research should be a mix of small (now on average less than \$100,000 per year, growing to \$200,000 per year), medium (approximately \$1 million per year growing to twice that amount), and large (more than \$4 million per year) research projects, plus networking research as part of two expedition centers (each approximately \$5 million per year). The relative ratio of the number of projects should be approximately 1 large: 4 medium: 20 small projects. Costs for applications development testbeds should be shared with the applications communities whenever possible.

The following numbers are in addition to current (FY 1999) funding levels.

Funding increases for Scalable Information Infrastructure Research (\$ millions)					
	FY 2000	FY 2001	FY 2002	FY 2003	FY 2004
Scalable Information Structure	60	120	180	240	300

3.3 High-End Computing

Since its creation under the High Performance Computing Act of 1991, the Federal High Performance Computing and Communications (HPCC) Program has contributed greatly to U.S. technological leadership. The Program has created and disseminated technologies to speed the pace of innovation, enhance national security, promote education, and help us understand the global environment. Furthermore, advances in high-performance technology have become an important component of mid-range computing and communications systems. Through the dedicated effort of many scientists and engineers, HPCC has succeeded in laying a foundation for future growth.

Historically, high-performance computing was a specialized market segment within the computer industry, serving the focused needs of science, engineering, and national defense. High-performance supercomputers were based on unique architectures and technologies, optimized to the requirements of these compute-intensive applications and were expensive relative to mainstream computers. However, over the last few years, the use of high-performance computing has undergone a major transition from purely compute-intensive applications to include data- and communications-intensive ones. As the scope broadened, the description “high-performance computing” was broadened to “high-end computing,” and now encompasses an increasing number of applications beyond science, engineering and defense.

With the advent of powerful microprocessors used in workstations and servers, a large portion of the high-end computing market shifted from supercomputers built with specialized processors to ones built with microprocessors. Due in part to Federally funded research efforts more than ten years ago, these systems, based on commodity parts, are now available at a fraction of the cost of earlier supercomputer systems. But for various technical and economic reasons, highly parallel systems (with hundreds or thousands of processors)—with the potential of ever-increasing performance into the teraops range (10^{12} operations per second) and beyond—have not yet been used widely for critical high-end applications.

Research in high-end computing has produced far-reaching benefits to society. Many of the present uses of high-end computing are well-documented in previous reports.⁶ They include:

⁶ Brooks, Jr., Frederick P., and Sutherland, Ivan E., co-chairs. *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*. Prepared by Committee to Study High Performance Computing and Communications: Status of a Major Initiative. Washington, DC: National Academy Press, 1995.

Lax, Peter D., [chairman]. *Report of the Panel on Large Scale Computing in Science and Engineering*. Under the sponsorship of Department of Defense (DoD) and National Aeronautics and Space Administration (NASA). Washington, DC, 1982.

Department of Defense, “National Security High End Computing Integrated Process Team for National Security High Performance Computing” briefing, 1997.

- Designing new cancer fighting and anti-viral drugs.
- Understanding the causes and sources of air, water and ground pollution, and devising solutions to these problems.
- Forecasting local weather and predicting long-range climate changes.
- Designing safer, fuel-efficient vehicles.
- Ensuring the safety and effectiveness of the Nation's nuclear stockpile.
- Designing new aircraft, such as the Boeing 777.
- Making scientific discoveries in a broad range of disciplines from elementary particles to cosmology.
- Simulating the effects of attacks on weapons systems, such as the New Attack Submarine.

High-end computing has three distinct and highly inter-related markets, each with its own dynamics and issues:

- The broader data-intensive and communications-intensive market, with applications such as data mining, web serving, digital libraries, and transaction processing. This market is doing very well, and growing very rapidly.
- The scalable, microprocessor-based compute-intensive market. This market is doing quite well at the mid-range of the market, with symmetric multiprocessor systems and parallel systems of moderate-size. However, the high end of this market, which highly parallel systems are expected to fill, remains in the research and experimental stage primarily because we lack adequate software technology, application-development tools, and, ultimately, well-developed applications.
- The vector-architecture-based sector. This market has hundreds of installed systems but is shrinking, primarily because most new applications are being developed for the more modern and economical microprocessor-based systems. Thus, the majority of the market for vector systems consists of older applications that have not been converted and tuned to high-end parallel systems referred to at the end of the previous paragraph. Currently, it is not known whether all vector-architecture applications, some of which are critically important to selected Federal missions, can be efficiently executed on high-end commodity processor-based parallel systems.

3.3.1 Findings

Finding: High-end computing is essential to science and engineering research.

The successes of the HPCC program have led to widespread use of computational simulation and modeling as a means to understand natural phenomena and to explore and optimize engineering designs. Over the next few decades, computation will become even more essential to the Nation's fundamental and applied science and engineering endeavors, both civilian and military. Currently, the extreme high-end of computing is done under the auspices of focused mission agencies such as the Department of Energy (DoE), while widespread academic research is primarily sponsored by

the National Science Foundation (NSF). Both for the sake of fundamental scientific research and to enable applications to benefit from the research, the research community needs access to systems at the leading edge of capability. The power of these systems should be comparable to that of systems-like the DoE Accelerated Strategic Computing Initiative (ASCI) systems available to the mission agencies. Furthermore, if the scientific community is to continue to benefit from high-end computing, it is important to ensure that there are enough computer and computational scientists to collaborate in making leading-edge computation an effective research tool.

Finding: High-end computing is an enabling element of the United States national security program.

Government customers account for at least half the market for supercomputer systems, with national security programs historically driving the market. These programs are very important to the Nation. Such Government organizations face the challenge of converting their applications to new high-end architectures without disrupting their missions. This transition has proved to be technically difficult and is exacerbated by the lack of mature software tools. The DoD HPC Modernization Program, DoE's ASCI, and the National Security Agency are examples of national-security users of high-end computing.

Through the DoD HPC Modernization Program, the Office of the Secretary of Defense is investing more than a billion dollars over five years in high-end computing to provide the United States military with a technological advantage. This program will provide advanced hardware, computing tools, and training to DoD developers using the latest technology to support the warfighter. The incorporation of high-end computing will allow the United States to maintain its technological supremacy in weapons systems design into the foreseeable future, while decreasing the total life-cycle costs of fielding new defense systems.

The DoE ASCI program is critical to the Stockpile Stewardship Program. ASCI is designed to assist in ensuring the safety, reliability, and performance of the United States nuclear stockpile in the absence of underground testing. In addition to addressing National security needs, the computer technology and products developed by the ASCI program will be applied to a broad spectrum of national needs. The computational techniques developed by the ASCI program have great potential, if widely adopted by the civilian researchers, to bring economic and scientific benefits to the United States in environmental studies, biology, drug design, automobile and aircraft development, consumer product safety, and information management and access.

For decades, the National Security Agency (NSA) has influenced high-end computing and has been an early and sophisticated user of the most advanced commercial computer, storage, and networking systems. In its mission to protect U.S. information systems and produce foreign intelligence information, the NSA has stimulated both industry and academia with some of the most challenging problems in the Nation. NSA research is directed to the discovery and application of methods seeking order-of-magnitude improvement for deriving intelligence by use of mathematical

and signal processing methods. Activities range from invention and prototyping of new concepts, to improving the ability to use leading-edge commercial products.

Finding: New applications of high-end computing are ripe for exploration.

In addition to advancing the long-term research agenda in traditional high-performance computing, the Committee believes that the Nation needs to develop new uses of high-end computing to promote a better understanding of our world and to improve services to all citizens. This expansion of high-end applications will attract computer, software, and application vendors to support a healthy high-end computing marketplace. Today the high-end computing market is a shrinking fraction of the information technology industry. Among the areas in which high-end computing could provide enormous benefits are:

- “Intelligent” systems: using sophisticated data mining algorithms against very large data bases to make expert decisions in a variety of application areas, from health care to market research.
- Design: using computer assisted design and engineering technologies in the many, more mundane, but far more broadly applied sectors such as fashion and architecture rather than just in extremely important but rarefied fields such as aircraft design and nuclear weapons stewardship.
- Transportation: managing route optimization and fuel efficiency for aircraft more effectively, at a savings of billions of dollars per year.
- Crisis management: simulating crisis scenarios, particularly weather and groundwater flow, critical to the management of emerging crises as well as to effective training and preparedness.
- Infrastructure support: managing the large-scale networks that are increasingly essential to the Nation’s economy and government.

Finding: United States suppliers of high-end systems suffer from difficult market pressures.

Despite the economic and societal benefits of solving problems like those described above, the market for high-end computing systems is mixed. The low- and mid-range parts of the high-end market—workstations and servers—are doing well financially, but the highest end, sometimes referred to as “supercomputing” to reflect its origins, has not been growing. Its sales are less than \$2 billion per year. This part of the market (for instance the fastest 100 computers installed worldwide—the Top100), historically dominated by U.S. vendors, has gone through a period of intense consolidation in the past several years, leaving only SGI (with its Cray subsidiary) and IBM as significant U.S. participants. No well-established U.S. company remains with its primary focus on technical, high performance computing. Recently established companies Tera and SRC have yet to show a profit.

At the same time, foreign competitors with substantial government subsidies have enlarged their efforts to capture the very-high-end market, which creates a potential reliance on foreign vendors for future specialized high-end systems. For instance, as of November 1998, whereas only 10 percent of the more broadly based Top 500 fastest computers installed worldwide are of non-U.S. manufacture, the foreign market share of the higher-end Top 100 rises to 25 percent. The Committee believes that this trend will continue unless the Nation invests in research and development programs in cooperation with industry and academia to increase U.S. leadership.

Finding: Innovations are required in high-end systems and application-development software, algorithms, programming methods, component technologies, and computer architecture.

These innovations need to be aimed at breakthroughs in sustained performance on real applications, programmability, ease-of-use, and scalability. Inadequate funds for software and algorithm research and development, as well as excessive focus on short-term research, have created an unduly limited software base usable by only the most intrepid users.

Current Federal programs that address innovative high-end architectures and technologies are too small, having shrunk dramatically in the last decade. Petaflops-scale computing is required for many of the applications listed above. To realize petaflops computing, we need advances in architecture, system software, algorithms, programming methods, and hardware technology.

Finding: The high-end computing capability available to the civilian science and engineering community is falling dangerously behind the state of the art.

Information infrastructure facilities increasingly serve a role for scientists and engineers similar to the role that large-scale shared physical instruments (for example, national telescopes, microscopes, and particle accelerators) have served for decades. Agencies like the NSF and DoE provide peer-reviewed access to leading-edge IT facilities to enable a wide range of disciplinary research.

For the past decade, science and engineering researchers working on civilian applications have typically had access to computing capability that lagged the national security community by a factor of two to five. This situation resulted from smaller budgets for the civilian-access machines and the timing of the national security requirements that required the first available systems. Today, the civilian science and engineering community finds its computing capability a factor of ten to twenty behind the largest installed systems, those used in the ASCI program. Furthermore, the decision to reduce the number of NSF supercomputer centers from four to two without increasing the funding for high-end systems in the remaining centers has substantially reduced the overall capacity available to academic researchers.

It is clear that these facilities are useful for maintaining our Nation's research leadership only if they are continually improved to keep them at the state of the art. However, Federal agencies need

to manage this need for continual IT facility improvement (for example, of high-end computers, high-performance networks, and large disciplinary databases) in a manner that does not directly conflict with funding the basic research enabled by the facilities.

3.3.2 Recommendations

In high-end computing as in other areas of information technology, we need more fundamental research—the kind of groundbreaking, high-risk research that will provide the ideas and methods for new disciplinary paradigms a decade or more in the future. Our greatest needs are improving systems software and algorithm-level software support at the high end, exploring innovative architectures and devices, and making it possible for the academic research community and the Federal Government to conduct essential research and development on computers of the highest possible performance. Research on software, architecture, and hardware can be effectively driven by an initiative to reach sustained performance in the petaflops and petaops range by the year 2010.

To have the greatest effect, the NSTC Subcommittee on Computing, Information, and Communications should coordinate all research on and access to high-end computing systems across all the Federal agencies. The program-planning processes should be coordinated, and the programs that result from the following recommendations should be closely coupled and mutually supporting. For example, architecture and software efforts are mutually dependent, and the recommended petaops milestone can drive both the hardware and software efforts.

Recommendation: Fund research into innovative computing technologies and architectures.

There is evidence that current scalable parallel architectures may not be well suited for all applications, especially where the computations' memory address references are highly irregular or where huge quantities of data must be transferred from memory to support the calculation. To address these limitations, the Committee recommends expanding research and development funding for new computer architectures. In particular, we need substantive research on the design of memory hierarchies that reduce or hide access latencies while they deliver the memory bandwidths required by current and future applications. Progress in this area would improve the performance of almost all programs and would make high-end software development less difficult.

The Committee recommends funding research into promising new technologies such as optical, quantum, biological, and neuromorphic computing. Ultimately, silicon chip technology will run up against the laws of physics. We do not know exactly when this will happen, but as devices approach the size of molecules, scientists will encounter a very different set of problems fabricating faster computing components. Potential areas for investment are improved simulation and analysis tools for investigating innovative architectures and full-system experiments to explore and validate architectural design alternatives for real-world application workloads. The Government must begin to invest now in these new technologies to ensure that computing capability continues to improve well into the 21st century. Architectures for high-end computing must be developed in concert and

iteratively with hardware and software technology investigations. The high-end computing research funding strategy should consider the coupling between these domains. For many of the larger projects, industry participation will be critical as only industry can provide important practical insight. As innovative research succeeds, industry will bring the new architectures and technologies into the open marketplace.

Recommendation: Fund R&D on software to improve the performance of high-end computing.

The Committee recommends substantial investment in software to improve the performance and efficiency of high-end computers. The software investments fall into three categories: system software, algorithm development, and software to manage integrated systems in a balanced fashion. Investments in system software—languages, compilers, runtime libraries, operating systems, file systems, I/O drivers, debuggers, programming interfaces, performance tuning tools, and so on will lead to improved efficiency and performance and will make high-end systems usable by a much larger community. Algorithm development is necessary to ensure the efficient use of leading-edge machines. Work on system software and algorithms will permit the effective exploitation of parallelism and will contribute to the efficient use of memory hierarchies.

For maximum performance, we require an integrated computer environment consisting of software to perform resource allocation, parallel input and output, ultra-high-speed and high-capacity intelligent storage, scientific data management and visualization. We need to do research on software that will improve overall system efficiency by eliminating potential bottlenecks between the high-end computer and its peripheral equipment. For example, petaflops computers will generate exabytes (10^{18} bytes) of data that the system must move rapidly to high-speed storage and then make readily available to users. Scalability of these tools is critical since they must keep pace with the rapid development of processors used in high-end computers. These efforts will not only increase usability, but also facilitate transition of applications, including those required for national security, to new computer architectures.

Cost effective and efficient use of emerging terascale systems (10^{12} operations per second) requires major advances in software research, development, and deployment. Without effective building blocks for software, algorithms, and libraries, high-end hardware is an expensive, underutilized technology. And adequate software will emerge only from an aggressive, stable, sustained, and long-term software research program. The software problem will become even more critical with the possible realization of a petaop capability (10^{15} operations per second) by 2010, perhaps earlier. The time to explore petascale software aggressively is now, not after we build hardware prototypes.

Even more alarmingly, limited funding for high-performance software research is causing an exceptional group of researchers to leave this research area in favor of other areas where funding is more readily available. If the Nation is to increase, or even retain, the pool of intellectual talent exploring terascale and petascale software, the Government must fund long-term software research.

Recommendation: Drive high-end computing research by trying to attain a sustained petaops/petaflops on real applications by 2010 through a balance of software and hardware strategies.

The Committee reviewed the CIC/HECC (High-End Computing and Computation) working group's proposed petaflop activity. The preliminary studies that group sponsored have determined that we will need substantial technological advances to achieve the desired performance levels by 2010. The goal of building systems capable of achieving sustained performance on important applications at the petaops or petaflops level can be an effective way to drive research on the hardware technologies, new architectures, software systems, and algorithms that will be needed to make such systems possible. We, therefore, recommend that funding for this activity should be increased to ensure that the necessary technological advances—in algorithms, systems software, application-level software components, and hardware—are achieved.

The goal of sustained petaops and petaflops performance must be understood as a technology driver, not as a goal unto itself. It is conceivable but unlikely that sustained, real-world performance in this range could be reached by faster hardware technology and increased hardware parallelism alone, but these sustained performance levels are almost certain to require a tightly coupled, well-balanced effort on both software and hardware. It would be counterproductive to produce a computer system capable of petaops performance without the software and algorithms needed to make it useful for the full range of high-end applications. An appropriate program will, for example, accommodate a technical comparison and evaluation of current architectures with promising new ones such as those based on processor-in-memory components. It should also address the software challenges of programming computers with tens or hundreds of thousands of processors and deep memory hierarchies.

Adequate funding would permit a petaflops program to accelerate the rate of progress towards petaflops computing by directly investing in and developing key hardware and software technologies that may prove crucial to long term computing objectives, even if they are not immediately suitable for short-term, medium-scale computing systems.

Recommendation: Fund the acquisition of the most powerful high-end computing systems to support science and engineering research.

High-end computing is becoming critical to an increasing breadth of science and engineering research. Recently the Nobel Prize in Chemistry was awarded to two pioneering computational researchers whose efforts led to the development of widely used methods in computational chemistry. Research like this helps us unlock the behavior of new materials, advance our understanding of biological structures, and understand and enhance the environment in which we live.

If the United States is to continue as the world leader in basic research, its scientists and engineers must have access to the most powerful computers. Therefore, the Committee recommends that the

Federal government continue to provide these computing systems to the research community through major, shared-facility centers. To increase long-term, fundamental research across all scientific and engineering disciplines, the first priority should be to increase the computing capacity of the centers that can best serve the entire research community. These are the NSF Partnerships in Advanced Computational Infrastructure (PACI) centers and the specialized centers that support specific disciplines like NSF's National Center for Atmospheric Research.

The first priority for high-end acquisitions should be to bring the performance level of academic computing close to that of mission agencies; DoE's ASCI program is the current benchmark for computational performance. The second priority should be to increase the capacity of focused-mission centers. Mission agencies like DoD, DoE, and NASA should seek additional funds to acquire high-end systems. We encourage the NSF PACI centers and mission-specific agencies to cooperate in providing computing capability to scientists and engineers independent of sponsoring organization.

Efforts should be made to expand the use of high-end computing in information technology research. For example, the simulation of proposed petaflops and petaops architectures will be essential to developing software for those systems, yet such simulators cannot even run simple computational kernels without using computers of the highest possible performance. The Committee, therefore, recommends funding additional high-end capacity to support the Research Expeditions to the 21st Century described in Section 5.3. These high-end systems will enable truly innovative research and thereby assist in developing a cadre of researchers who can exploit high-end computing in the 21st century. These high-end systems will also create an environment for new applications and will expand the critical, domestic high-end computing market.

Recommendation: Expand the NSTC CIC High End Computing and Computation (HECC) Working Group's coordination process to include all major elements of the government's investment in high-end computing.

The Committee recommends that all of the Federal Government's major research, development, and acquisition investment budgets and plans in high-end computing be provided to the NSTC CIC Subcommittee's HECC Working Group. Currently, some large investments in high-end computing such as the DoE's ASCI program, the Department of Defense's High Performance Computing Modernization Program (HPCMP), and some of NSA's investments are excluded. Thus, these programs and their contributions to high-end computing are generally excluded from the Federal Government's attempt to coordinate high-end computing efforts across agencies. Exclusion of any large investment makes it very difficult to coordinate and plan research and development of high-end computing.

Budget Recommendation: Increase current funding for high-end computing research and acquisitions as follows over the fiscal years 2000-2004.

The budget recommended for the high-end effort was developed using a three-step process. First, four subpanels were formed—innovative technologies and architecture, high-end software, sustained petaops/petaflops, and high-end acquisition. Second, based on the findings in this chapter of the report, each subpanel independently determined the technical objectives underlying the recommendations. Third, each subpanel determined a budget required to accomplish these objectives. The budget below totals the research budgets of the first three subpanels, keeping the acquisition budget separate.

The following numbers are in addition to current (FY 1999) funding levels.

Funding increases for High-End Computing Research (\$ millions)					
	FY 2000	FY 2001	FY 2002	FY 2003	FY 2004
High-End Software, Architectures, and Petaops Computing	180	205	240	270	300
High-End Acquisitions	90	100	110	120	130
Total	270	305	350	390	430

This table covers only those efforts described in the first four recommendations above. It excludes efforts currently excluded from the HECC crosscut and coordination process addressed in the fifth recommendation.

4. Socioeconomic Research and Policy Priorities

Our National well-being depends on understanding the potential social and economic benefits of on-going advances in information technology. We must also understand the problems that are arising from the increasing pace of information technology-based transformations. To realize the promise of the new technologies, we must invest in research to identify, understand, anticipate, and address these problems. We need more data, and we need to understand social, economic, and policy issues in much greater depth. The research that is required to develop this knowledge should be broad-based, long-term, and large-scale, in its scope and in its promise for continuing the beneficial transformation of our society and economy.

We must ensure that all Americans are well-equipped to live in this changing world. To that end, we must improve education beginning at the K-12 level and extending through post-secondary education. To improve the skills of our information technology workforce and to increase the information technology literacy of our overall workforce, we must also concentrate on education and retraining for all of the information technology skills that are important to our society. Furthermore, we must concentrate on programs that will enable the United States to use its diverse human resources to the fullest capacity.

The promise of the information age must include all citizens, regardless of ethnicity, gender, physical ability, income level, or geographic location. Access to information technology must be accompanied by the assurance of a high degree of personal privacy. Programs to provide broader access should be based on sound understanding of the social and technical conditions and emerging developments to ensure that we reach out to people wherever they are through the most efficient, effective means.

Information technology can significantly improve the flow of information to all people and institutions in the Nation, and can thus be a powerful tool for democratization, equity, and access. However, to achieve the potential, we need concrete objectives and comprehensive metrics. We must conduct careful research to measure the effects of the transformations against the objectives, and we must develop appropriate policies based on knowledge gained from the research.

4.1 Findings

Finding: The use of information technology—in particular the growing popularity of the Internet and the emergence of global commerce—has introduced a series of important and complex policy issues.

In addition to the many technical challenges discussed in the preceding chapters of this report, there are many barriers to realizing the benefits of information technology that are social, political, or legal in nature. For example, consider the host of problems associated with implementing our common vision of transforming the practice of health care, laid out in Chapter 1 of this report.

Several challenging technical problems remain to be solved through IT research, before telemedicine can become a common practice. However, the solution of problems arising from several socioeconomic and policy issues will, most likely, present even more formidable barriers to the deployment of telemedicine. A few of those issues include the privacy and security of patient records, use of records by insurance companies, current requirements for state-by-state licensing of doctors, and the development of mechanisms for creating trust between patient and doctor in an on-line environment. Similarly, widespread use of electronic commerce will require mechanisms for insuring even greater security and trust between parties in an on-line environment. These are some of the important questions that can be addressed only by further research on the economic and social effects of computing and communications.⁷

Further, the use of information technology, particularly in a widely accessible networked environment, has led to new policy issues. Among the areas in which concerns have arisen are personal privacy, intellectual property rights, digital signatures, and the control over the availability of particular kinds of content (for example, pornography or inflammatory rhetoric.)

An example of an important set of complex policy issues of particular importance are those related to privacy. The management and analysis of information is critical to many applications. Information about users and their usage of applications will assist service providers in business, government and the medical profession to provide informed, cost-efficient, and highly-personalized service. However, this will be possible only if individuals are guaranteed that their information is truly protected. If there is no trust in the system, its benefits will be severely limited. Guaranteeing the privacy of information is a prerequisite for many of the exciting applications that are possible to build. Our citizens must work together with government and business to ensure that only the most appropriate regulations are put into place. There is a very fine line between helping citizens to feel safe and creating an undesirable and onerous bureaucracy.

On an international scale, the European Union Privacy Directive heightens the urgency of information privacy policies. Under European law, firms doing business in the E.U. are prohibited from sending personal information to countries, such as the U.S., that do not adhere to the privacy principles embodied in the European directive. This is likely to have a dramatic impact on U.S. electronic commerce.

Finding: Policy decisions and IT investments are being made on the basis of incomplete research and data concerning the effects of IT on our society.

Integration of IT into our personal and professional activities has not come about by accident: decisions have been made that have brought about the many transformations we are currently

⁷ National Research Council, Computer Science and Telecommunications Board. *Fostering Research on the Economic and Social Impacts of Information Technology. Report of a Workshop.* Washington, DC: National Academy Press, 1998.

experiencing. As the pace of IT research and integration accelerates, more decisions will be made which have ever-expanding consequences and effects on our daily lives. Consequences will range enormously in their scope and scale, from determining the format through which our children are first introduced to IT in early primary school, to determining policies which have long-lasting effects on our national transportation, communications infrastructures and economic practices. In too many cases, decisions about IT implementation have not been properly informed by adequate research. In fact, it is often the case that the implementation of information technologies has a considerably different set of consequences than were originally intended or anticipated.

Much more social science research on the impact of IT on our society is needed to inform ongoing debates and policy decisions on IT-related issues. Such research can also help IT researchers develop technical solutions for some difficult policy problems (e.g., development of new metadata tagging standards and micropayment technologies for managing intellectual property). Moreover, insights derived from social science research may be able to contribute to the better designs of information systems. The design of groupware, for example, should be driven by research on how groups of people share information and make decisions.

Equity and Access

Finding: All of our citizens must have access to information technology.

In his remarks at the MIT commencement ceremony on June 5, 1998, the President reiterated his vision of limitless possibilities for all Americans brought about by the very technical advances that we are recommending in this document. One of the challenges before the Nation is to extend those possibilities to *all* Americans. In an increasingly competitive global economy, our Nation cannot afford to squander our human resources by providing opportunities only to those Americans who are favored by geographic or economic circumstance. Access to and use of information technology, particularly in educational settings (K-12 through university) is a prerequisite to building the skills that will allow our citizens to function productively in the 21st century information society. It is a critical stepping stone for instilling interest and developing the skills of the budding information technology researchers who are essential to sustaining our national research capabilities.

Information technology can help eliminate social and economic inequities and mitigate problems arising from physical disabilities. Information technology tools and applications can provide opportunities that transcend barriers of race, gender, disability, age, income, and location. The enabling quality of the technology, in addition to the cultural values cultivated through its best-known application, the Internet, carry a democratizing potential that already has transformed our social interactions and economic opportunities, both at home and abroad. Public libraries have always played a critical role in ensuring equity of access to information resources and the role of public libraries is even more vital today in guaranteeing free and equitable access for all to Internet resources and communications access. It is important that we study the impact, costs and potential funding sources to support this continuing essential role for ensuring equity of access by public libraries across the country.

Yet we see a gap in access to the Internet and the accompanying participation in the information technology revolution among racial groups. For instance, whites are more likely than African-Americans to have Internet access within or outside of their households, to have access to a PC at work, and to have used the Web.⁸ We expect there are similar gaps with other minority groups, such as Hispanics and Native Americans. Recent research released by the National Telecommunications and Information Administration (NTIA) suggests that the racial gap in Internet use is increasing.

The consequences to our Nation of a persistent racial divide on the Internet will be severe. If a significant segment of our society fails to have equal access to the Internet, U.S. firms will have difficulty hiring personnel with the necessary technological skills, and competitiveness will suffer. Employment opportunities and income differences among whites and minority groups will be exacerbated, with further adverse consequences for the Nation. The Internet may provide equal opportunity and help level the playing field, but only for those with access.

Finding: Full participation in information technology research requires access to high-bandwidth connectivity.

Information technology research and education rely increasingly on high-speed network connectivity for access to information, collaboration between geographically separated colleagues, and access to advanced technology resources such as astronomical telescopes and high-power electron microscopes. Scientists, educators, and students at rural, small, and ethnic universities are often unable to participate in new research and educational opportunities because they are not connected to the vBNS⁹ or other fast Internet backbones. The experimental program to stimulate competitive research (EPSCoR) exists to promote the full participation of institutions in eighteen states and Puerto Rico that have been underrepresented in research partnerships. However, small or rural universities exist in other states as well. Full participation, particularly in research partnerships such as the proposed Expedition centers, is not possible without high-bandwidth connectivity.

Workforce

Finding: The supply of information technology workers does not meet the current demand.

By virtual unanimity, chief executive officers of a cross-section of America's leading corporations have identified the need to strengthen the technological workforce as the single greatest challenge to U.S. competitiveness over the next decade.¹⁰ It is crucial that we produce a continuous supply of well-trained, high-quality professionals in engineering and computer and information science, not

⁸ Hoffman, D. L., and Novak, T. P. "Bridging the Racial Divide on the Internet." *Science*, April 17, 1998: 390-391.

⁹ vBNS is NSF's very high performance Backbone Network Services.

¹⁰ Council on Competitiveness. *Competing Through Innovation: A Report of the National Innovation Summit*. Washington, DC: March, 1998.

merely skilled users, but researchers, creators, and designers of advanced technology. In this fast-moving field, those people must continue to update their knowledge and skills.

Today we fall far short of meeting these needs, and projections for the future are not encouraging. While the information technology sector and demand for skilled personnel are growing rapidly, the pipeline for computer engineering and computer science graduates is not filling fast enough. Beginning with skills that require a BS-level of training, qualified information technology workers are in extremely scarce supply. The largest fraction of open positions exists in those requiring graduates with advanced degrees. In its Interim Report to the President, the Committee recommended increasing the number of H-1B visas as a short-term measure, but this solution is untenable for the long term.

There is evidence that large pools of potential information technology personnel in the U.S. workforce exist and employers are recruiting new workers from outside information technology fields. However, there is also the problem of what economists call “appropriability”, a type of “market failure.” Many information technology employers who are finding hiring more difficult are unable to provide the required education or training due to resource limitations. Others are unwilling to do so because of competition for their newly trained employees. Meanwhile, many potential information technology workers may be motivated to get the education and training needed for information technology careers but unable or unwilling to pay the (often substantial) costs involved—they are often otherwise employed and precluded from fulltime study. It should be recognized that a complete solution to the workforce problem would involve recruiting from underrepresented groups as well as nationwide excellence in K-12 education. These issues are discussed more fully in the following findings.

Finding: A diverse workforce literate in information technology is critical for ensuring that our Nation is prepared to meet the challenges and opportunities of the Information Age.

It is critical to tap all of this country’s talent for the information technology workforce. There is, without any doubt, a vast untapped talent pool in the United States among women and minorities, currently under-represented in engineering and information sciences. African Americans, Latinos, and Native Americans constitute a fourth of the total U.S. workforce, 30 percent of the college-age population, and a third of the birth rate. Yet members of these minorities collectively comprise only 6.7 percent of the U.S. computer and information science labor force (all degree levels), 5.9 percent of the engineering workforce, 1.7 percent of the U.S. computer science faculty, and 4.9 percent of the engineering faculty.¹¹ Women participate in both computer science and engineering at low rates for all degrees and subsequently in academia and industry.

¹¹ Campbell Jr., George, Ronni Denes, and Catherine Morrison, eds. *Access Denied: Race, Ethnicity and the Scientific Enterprise*. New York: Oxford University Press. Forthcoming, August 1999.

These facts, trends, and issues represent a significant opportunity to build the information technology workforce we need. However, we can achieve the goal of a stronger technological workforce, remain economically competitive, and prosper only by developing and employing the talents of the currently under-utilized majority of the population.

Finding: Both K-12 and post-secondary education are inadequate to meet the challenges of the information age.

If our citizens are to reap the benefits of the information technology revolution, personally and professionally, and if the Nation is to have the information technology-literate, well-educated, and highly skilled workforce it needs, every student must learn some aspects of information technology and many must become highly educated experts. That is not happening everywhere now. Many children do not have the opportunity to acquire the information technology skills they will need in the future. Many adults have no access to continuing education or the retraining that would enable them to hold good information technology jobs. Furthermore, the educational pipeline is not preparing enough people for the information technology workforce.

4.2 Recommendations

Our recommendations are organized into two categories: recommendations involving education, training, and workforce policies and those calling for additional research. To begin to solve some of the problems described above, new research and substantive policy changes will be needed. Some of the research we recommend could well be supported within the agencies that traditionally fund science and engineering, but other appropriate sponsorship for this research should also be considered.

Overarching Recommendation: Expand Federal initiatives and government-university-industry partnerships to increase information technology literacy, education, and access.

In implementing the research agenda recommended in this report, the Government must increase its efforts to provide broader opportunities for developing information technology literacy, access, and research capabilities. The Government must also expand opportunities for information technology literacy and access as it addresses issues related to workforce skills and development. Specific recommendations are provided below.

Recommendation: Expand Federal research into policy issues arising from information technology.

As the information revolution introduces changes in society, the United States must continue to protect its citizens' information privacy, the rights to intellectual property, the ability to provide verifiable signatures, and the ability to exercise control over appropriate access to information. New policies and procedures are certain to be needed. As an example, we consider information privacy—an issue that is increasingly visible.

It is appropriate that industry take the lead in establishing and implementing privacy procedures. Government regulation might be required if industry fails to act properly, or in very sensitive areas like health care and protection of minors.

Policies for the protection of individuals' information privacy on the Internet should be evaluated and established. Industry education programs should be developed that carry the message that protecting consumer privacy on-line will actually be good for business. The government and industry should jointly fund this effort.

The Committee recommends that research in privacy issues be aggressively funded to help address the following questions:

1. What privacy policies would both protect individual privacy in electronic environments and permit the conduct of business on the Internet? For example, common United States practice is to permit users to “opt-out” of information sharing, but our notification procedures are often lacking. The E.U.’s directive requires full disclosure (notice) and informed consent (“opt-in”). Other areas of study include the conditions under which individuals should be able to change information recorded in databases and how noncompliance with privacy policies would be redressed.
2. What mechanisms can best address the protection of individual privacy on the Internet? For example, should the U.S. try to develop regulatory mechanisms for enforcing privacy policies, or are self-regulatory efforts more effective? Should an industry trade group take the lead? Should a nonprofit organization be developed? How can the U.S. best foster dialogue on privacy?
3. How is the E.U.’s Privacy Directive likely to affect the development of electronic commerce, both locally and globally?

Recommendation: Fund information technology research on socioeconomic issues.

Since an understanding of the socioeconomic issues that arise because of technological change is intrinsically a multi-disciplinary investigation, there should be research on social, economic, and policy issues. The studies could be tied to particular aspects of information technology, such as the Internet, or they could be more broadly based. As appropriate, research on socioeconomic issues should be carried out in Expedition or Enabling Technology centers, including some of those with a primarily technical focus. These centers should define and legitimize research areas, provide the scale of research activity necessary for experimental research in socioeconomic topics, allow for research teams of sufficient size for multidisciplinary research, and allow sufficient time for the maturation of a research area.

Possible areas of research for such centers include (but are not limited to):

- Transformations of social institutions.
- Governance.
- Electronic commerce research.
- Social and economic simulation and modeling.
- Sustainable use of large information infrastructures.
- Electronic groups and communities research.
- Barriers to information technology diffusion.
- Human interaction and communication laboratories.
- Innovative scientific development of a national information infrastructure to support research on social and economic issues.
- Digital government.

Recommendation: Expand the participation of underrepresented minorities and women in computer and information technology careers.

Systematic research efforts are needed to identify specific barriers experienced by underrepresented groups in their pursuit of careers in information technology. The research program should include investigations of pre-college preparation issues; barriers to undergraduate enrollment, financial aid and admissions policies; retention in information technology degree programs; graduate education issues, career access; and career advancement obstacles. Experimental programs should be developed to test the effectiveness of different approaches toward reducing those barriers. In particular, it will be important to determine if (and how) the tools of IT can themselves be used to break down these barriers. A significant proportion of this research could be carried out at the Enabling Technology Centers, and at an Expedition Center focused on socioeconomic and workforce issues.

Based on the data obtained from this research, and from other effective programs already in place, the IT community should increase its efforts to recruit women and minorities into IT careers. We note that several private-sector programs have been successful in producing minority and women graduates, particularly in engineering. They have developed a high-quality research and support infrastructure that provides such services as alternative assessments to identify potential recruits among educationally disadvantaged high school students, essential academic enrichment programs for those students, scholarships, mentoring, internships and professional development. The National Science Foundation has also developed several programs aimed at improving education in under-served communities, and at increasing participation of minorities in science and engineering. Those successful programs should be extended to encompass information technology and related fields. Government support for exemplary programs operated through public and private foundations should be increased substantially.

Recommendation: Create programs to remove the barriers to high bandwidth connectivity posed by geographic location, size, and ethnic history of research, educational institutions, and communities.

Programs to facilitate high bandwidth connectivity, such as the NSF vBNS connectivity grant program, should be continued and expanded to address problems posed by geographical location, size or ethnic history. It is essential that such limiting factors are removed so that all institutions with efforts related to the network testbeds are included, both to provide the necessary research and to contribute needed students to the information technology workforce. Special supplements to connection funding, such as the present EPSCoR supplements in the NSF Connections program are one approach to this problem. Other approaches such as research to develop and apply new technology to facilitate access and mitigate policy/regulatory concerns should be addressed, as well. In addition, we should support community efforts to provide connectivity through local institutions, such as libraries, schools, places of worship, and community centers.

High-bandwidth connectivity in every state is needed. Similar to historic large-scale national infrastructure programs such as the Rural Electrification Act or the interstate highway program, the construction of an “interstate information superhighway” system must be set as a high-priority, long-term National goal. We do not propose, however, that the government subsidize the deployment of a national high-bandwidth infrastructure at this time. Independent of any related policy considerations, the expense would be prohibitive using current technologies. Rather, the first and essential step toward attaining that long-term goal is to mount the type of research effort proposed in this report. By doing the research necessary to decrease connectivity costs, nationwide high-bandwidth connectivity will become possible.

Recommendation: Accelerate and expand education in information technology at all levels—K-12, higher education, and lifelong learning.

To prepare our citizens for the Information Age, education and training are needed for computer literacy, for entry into jobs and careers that require information technology skills, and for continuing to update knowledge and skills in the face of rapid technological change.

As part of preparing students to live in an information technology world, they must learn the basic use of computers and the Web to acquire information, communicate with others, and perform other daily tasks. In addition, computer science and information science material must be incorporated into the K-12 curricula as a natural augmentation of education in math and science. In order to make education in information technology available to all K-12 and college students, effective mechanisms must be developed to train highly qualified computer and information science educators, and to encourage them to teach in under-served communities.

Lifelong-learning is an extremely important part of the information technology profession—information technology professionals must continually update their knowledge and skills to keep up

with this fast-changing field. Advances in personal computers and Internet access can enable learners to acquire high-quality information technology courses and training anytime and anywhere. Information technology education and training should be made readily available to motivated learners.

Collaborations among schools, government, and industry will be an important factor in information technology education and training. An example of a program that facilitates collaboration between government and industry to enhance information technology education and training is the National Science Foundation Advanced Technological Education Program (ATE) begun in 1994.¹² Other programs and incentives for such collaborative efforts should be developed and supported.

Recommendation: Strengthen the use of information technology in education.

Information technology has great promise as a vehicle for improving both the quality and the availability of education at all levels. Although many ideas have been advanced and tried, considerably more research and experimentation are needed to realize that promise. At the same time, it is essential that the use of those aspects of information technology that are known to be effective for education become more widely available and used.

The President's Committee of Advisors on Science and Technology (PCAST) recently released the "Report to the President on the Use of Technology to Strengthen K-12 Education in the United States." They recommend that the Federal government play a leadership role in supporting an aggressive research agenda to evaluate the efficacy in both the short and long run of the educational approaches that are being fostered by the new technologies. This Committee strongly endorses that recommendation.

Interest in "electronic learning" is growing rapidly among universities, community colleges, and private education providers. Yet information technology educators and trainers often don't use it. The use of information technology would be particularly beneficial for teaching people about information technology. In particular, we believe these "anytime, anywhere" approaches offer the potential for high-quality education and training to be accessible by large pools of would-be students (including high school students) who could not otherwise participate. Such approaches could be powerful responses to needs for both increased numbers undertaking initial education in information technology fields, and for information technology professionals to constantly update their skills while continuing their employment.

¹²Advanced Technical Education (ATE) grants are specifically designed to improve science, mathematics, engineering, and technology education for community college and high school students preparing for careers as technicians in such fields as biotechnology, electronic, and communications.

Recommended Funding Levels

The following numbers are in addition to current (FY 1999) funding levels.

Funding increases for Socioeconomic Research (\$ millions)					
	FY 2000	FY 2001	FY 2002	FY 2003	FY 2004
Socioeconomic Research	30	40	70	90	100

The committee recommends ramping up to \$100 million over five years. In the first year, we recommend spending an additional \$30 million, as follows:

- 1) Add \$10 million to fund more socioeconomic research. This includes workforce data collection and multidisciplinary socioeconomic research and one to two Enabling Technology Centers at \$3-5 million each.
- 2) Add \$10 million to increasing research on educational technologies.
- 3) Add \$10 million to fund high achieving and pilot programs in engineering or science that target students from economically disadvantaged populations.

The large increase in the third year represents our recommendation to fund an Expedition Center.

5. Creating an Effective Management Structure for Federal Information Technology R&D

Given the importance of information technology to the Nation’s future, it is critical that the Federal research investment in this area be well managed. An effective management structure for information technology research can ensure that the Federal research portfolio is appropriately balanced, can encourage innovative long-term research, can help the Administration articulate research priorities, and can help assure Congress that its goals are being met. The management structure should ensure a diversity of funding mechanisms, including single-investigator grants, multiple-investigator projects within and across disciplines, and integrated research centers. Such diversity in approach and funding modes will increase opportunities for unexpected discoveries, encourage a broader attack on problems, and ensure fewer missed opportunities.

Unlike other disciplines that are usually incremental and evolutionary, computer science and engineering is a transformational and revolutionary technology. Since tomorrow’s applications will be ones that we cannot even envision today, it is critical that we as a Nation invest in core computer science and engineering to enable those revolutionary advances. The Committee feels strongly that information technology research must not be driven entirely according to the extrapolated demands of today’s commercial, scientific, and/or national security applications.

An effective management structure should accomplish the following:

- Ensure important national goals for information technology are addressed.
- Provide a focused and informed policy setting.
- Allow for strategic management and planning to ensure adequate investments.
- Identify the entire Federal information technology R&D endeavor, so that an assessment of the portfolio can be made.
- Guarantee that Federal research funds are managed by using open, competitive selection processes, thereby ensuring that the best ideas and the best investigators receive support.
- Provide mechanisms to connect complementary efforts across research disciplines and funding agencies.
- Review and assess Federal efforts.
- Provide a stable, but flexible overall program that sustains a healthy base of computing, computation, and communications science technology.

5.1 Agency Roles

Currently, funding for information technology R&D comes from several different agencies, with no single agency having information technology as its primary responsibility. This system has worked surprisingly well in the past and the Committee views the diversity of research funding as a crucial strength.

However, the Federal funding portfolio has not kept pace with the importance of information technology to the Nation, especially with respect to support for fundamental research in computer science and engineering. The Committee recommends that the proposed budget increase be used primarily to fund fundamental information technology R&D, with a significant portion of the funds directed to the National Science Foundation (NSF). The Committee believes that NSF should use this increased funding to reinvent and significantly enhance its traditional role as the champion of core computer science and engineering research.

In parallel, other agencies conducting information technology research should renew their commitment to fundamental research and engineering, with specific emphasis on making software a substantive component of every major information technology initiative. For example, the Defense Advanced Research Projects Agency's (DARPA) flexible and creative management style and its technical expertise and history of successful investments in information technology R&D has and should continue to make significant contributions. Participation by a variety of agencies, which tend to have different styles and emphases for supporting research or procuring technology, is important to establish and maintain an effective U.S. information technology R&D program. This diversity allows many views to compete and fosters pursuit of intellectual questions via a range of complementary modes.

Such information technology research will be motivated by long-term mission goals, but should be evaluated and funded with fundamental information technology challenges in mind—developments that will generalize to other domains despite their genesis in a specific mission agency. The Committee believes, for example, that biomedically motivated basic research in information technology should be supported by the National Institutes of Health (NIH) and viewed both as important information technology research and as fundamental biomedical research. Similar arguments can be made for DARPA, DoE, NASA, and the other research agencies that will be continue to be involved with mission-motivated basic research in information technology.

Recommendation: Strongly encourage NSF to assume a leadership role in basic information technology research. Provide NSF the necessary resources to play this role.

In its leadership role, NSF should foster interagency collaborations to ensure adequate funding levels in basic research. While NSF can only control its own budget, it should work with other agencies with similar research needs to assess major requirements and needs being addressed, identify shortcomings of the fundamental research portfolio, and initiate efforts to address deficiencies.

This leadership role for basic information technology research may require innovations internal to the NSF, as may be necessary to ensure that NSF is able to define, support, and coordinate a broad range of modes of research support. Example modes of support are centers of diverse sizes and multiple-investigator projects with longer terms. NSF may also have to make changes to ensure

that it fosters high-risk, high-payoff long-term research and invests sufficiently in the core of computer and information science and engineering. We also believe that there needs to be significantly more information technology representation on the National Science Board than exists now.

To successfully carry out the program of research proposed in this document, NSF will need to support a portfolio of modes of research that differ from the current mix of mostly single-investigator research grants together with a small number of centers. For example, multi-investigator projects of longer duration (5-7 years of funding) are needed in order to carry out high-impact experimental research agendas, whether or not such grants are available in other parts of the Foundation. Some of these projects will include long-term basic technology research that has an engineering orientation and involves explicit technology transfer. Some projects should be supported in response to topic-specific solicitations in order to support a DARPA-like concentration of effort on compelling research areas. Because of the nature of the field, the portfolio must successfully meld science investigations with technology exploration.

Roughly 40-50% of the proposed budget increases for information technology should go to NSF, to support basic research in information technology focused in the Computing and Information Science and Engineering (CISE) Directorate, with the rest allocated to other research agencies. The majority of the NSF increase should go to the new programs and modes of funding; the remainder should go to the traditional programs within the CISE Directorate, expanded as appropriate to projects of larger size and longer duration.

5.2 Policy and Coordination of Information Technology R&D

Given the importance of information technology to Nation's economy, security, and well being, the Committee believes there should be high-level management attention and focus by the Administration for information technology R&D.

Recommendation: Designate a Senior Policy Official for Information Technology R&D.

Information technology is of critical importance to the Nation; it is the basis for a \$700B domestic industry that drives many of today's innovations and scientific discoveries and offers even greater potential for innovation in the next century. The Committee believes that information technology R&D requires a high-level policy voice, similar to that afforded to other priority research areas in the White House. Positioning an individual at this level would help ensure the Administration receives timely advice on information technology issues and foster a strategic approach in determining R&D investments. This individual would be responsible for leading the White House effort to establish Federal policies to support, encourage, and help coordinate long-range information technology development to maintain U.S. leadership in this vital part of our economy.

Recommendation: Establish a senior-level policy and coordination committee to provide strategic planning and management.

The principal goal of coordination is to ensure that the Federal information technology research portfolio is sufficiently aggressive and properly balanced to meet the pressing needs of the Nation. Since many Federal agencies fund information technology R&D, a high-level venue for cross-agency coordination is desirable. The coordinating committee should consist of agency directors, since it is they who have policy and budget making authority. This coordination committee should advise and report directly to the President's Assistant for Science and Technology to increase the overall effectiveness and productivity of Federal R&D efforts to develop and apply information and communications technology. The coordination committee should address significant national policy matters that cut across agency boundaries. The committee should establish objectives for research programs and review them to ensure that they are meeting those objectives. In addition, it should ensure that the overall Federal program is well-balanced and has good coverage of important topics.

This coordination should include the entire Federal information technology R&D endeavor, including the efforts resulting from the proposed budget increase. A problem with the present management structure is the difficulty of getting a clear and complete picture of the total Federal investment in information technology, since there are Federal computing, information, and communications R&D programs that are not part of the formally coordinated Computing, Information, and Communications (CIC) programs. The coordination in information technology should encompass all major investments in R&D, as well as investments in people, supplies, and the like. For instance, high-end computing acquisitions are a necessary part of a coordinated program.

Recommendation: Extend the HPCC program coordination model to major Federal information technology R&D activities.

To make this coordination successful, the entire enterprise should have strong staff support in the White House and in the agencies themselves. The HPCC program, with the National Coordination Office and working groups with agency representation, is an effective model of inter-agency collaboration. The Committee recommends extending this model to all major Federal information technology R&D activities, with the NCO facilitating interagency coordination and supporting the management structure recommended by the Committee.

5.3 Support and Implementation

To achieve the goal of reinvigorating long-term, high-risk research in information technology, Federal funding agencies will need to support many different kinds of projects, including single-investigator grants, multiple-investigator projects within and across disciplines, and integrated research centers. To encourage the visionary planning needed to produce true innovation, funding

for these activities will need to be provided for longer periods—long enough for high-risk projects to be carried through to completion. Furthermore, because effective research in computer science involves collaborations of all sorts—with other computer scientists, with researchers in other disciplines and at other institutions, with industry, and with the community—the modes of support should encourage these collaborations, without turning them into strict requirements. Research projects should have the flexibility to organize themselves in the best ways to achieve the goal of technological innovation.

Recommendation: Diversify the modes of research support to include more projects of broader scope and longer duration, placing a renewed emphasis on research carried out in teams.

During the 1970s and 1980s, DARPA computer science researchers were encouraged to imagine dramatically different futures and to carry out projects that would explore those futures. Researchers were given enough resources and time so they could concentrate on problems rather than worry about their next proposals. The results were dramatic advances in speech recognition, robotics, chip design, high-performance computing, machine vision, artificial intelligence, and virtual reality. The Committee would like to see this spirit renewed and replicated across all Federal information technology R&D programs.

Currently, most NSF fundamental research support is for single-investigator grants. Although this kind of support will remain an essential component of the research portfolio, information technology can often make the most dramatic progress when carried out in teams of moderate size. DARPA has used this model effectively for several decades, and it is rapidly becoming part of the portfolios of other funding agencies including NSF. The committee believes that team research in information technology should be at least as important as single-investigator research, including the National Science Foundation. Furthermore, inter-institutional team projects help build a coherent national research community.

Recommendation: Fund collaborations with applications to drive information technology research, but take measures to ensure that research remains a primary goal.

A principal goal of information technology research is to produce results that will someday have commercial or social applicability. Thus, it is important to understand applications to which information technology may be applied and to use this knowledge to drive research into new technological strategies. Therefore, the Committee recommends that the Federal Government continue to fund collaborations between information technology researchers and applications developers. However, the funding programs must be carefully structured to ensure that the research remains a principal focus.

In the early 1990s, a principal focus of the Federal High Performance Computing and Communications (HPCC) Program was the use of large-scale applications efforts to drive research on parallel computing. In this vision, teams of applications developers and computer scientists

would have worked together to rebuild major applications on parallel computing platforms developing many useful computational support technologies. Unfortunately, many of the projects came under pressure to meet progress and performance deadlines for the applications, leading them to jettison all but the most immediately applicable computer science research.

To reemphasize long-term fundamental research, we must acknowledge that the time to applicability will be too long in many cases to directly affect an application project that is driving the research. For example, if a computer architecture researcher works with a scientist on computational fluid dynamics to understand how to design better computers for that application, it is unlikely that machines of the resulting architecture will run that specific application in a two to three-year time frame. In a funding program designed to foster long-term research, this situation must be completely acceptable. Indeed, it is essential that the funding programs be designed to give equal value to the information technology research as to progress on the application.

Recommendation: Fund centers for Expeditions into the 21st Century.

Projects of moderate size cannot alone address the entire spectrum of long-term research needed to sustain our national leadership. To foster research with truly dramatic impacts, the Committee recommends the creation of two types of center-sized activities: “Expeditions into the 21st Century,” which will endeavor to intersect with the long-term technological future, and “Enabling Technology Centers,” which will focus on research with application to problems of national importance.

“Expeditions into the 21st Century” will be centers, perhaps virtual, that bring together scientists, engineers, and computer scientists from academia, government, and industry to “live in the technological future.” The mission of these expeditions will be to report back to the Nation what could be accomplished by using technologies that are quantitatively and qualitatively more powerful than those available today. In essence, these centers will create “time machines” to enable the early exploration of technologies that would otherwise be beyond reach for many years. Just as the Lewis and Clark expedition opened up our Nation and led to unanticipated expansion and economic growth, the ideas pursued by information technology expeditions could lead to unexpected results and nourish the industry of the future, creating jobs and benefits for the entire Nation.

There are existing successful examples of the use of this “living in the future” approach. In the private sector one of the most famous examples is the Xerox Palo Alto Research Center (Xerox PARC) where researchers created an experimental network of computers for use by individuals. This effort pioneered many of the revolutionary technologies that led to today’s personal computers, including graphical user interfaces, pointing devices, laser printing, distributed file systems, and WYSIWYG word processing. In the university community, the MIT Media Lab has been conducting similar explorations. Finally, there is the example of ARPAnet, which evolved into today’s Internet.

The Committee recommends funding several Expeditions, each with a different focus. The focus may be on either a discipline-based theme, such as bioinformatics or multi-scale engineering, or on an infrastructure-based theme, such as distributed databases or tele-immersion. To establish a context, each Expedition should be based on assumptions not true today, for example, ubiquitous computing or a vast amount of simulation, as described in Gelernter's *Mirror Worlds*. Each Expedition need not be limited to a single such assumption, but an Expedition should invest sufficient resources to make exploration of its assumption areas, those parts of the map of the future, possible.

Each Expedition would be required to carry out several activities, including:

- *Technology testbeds.* Expeditions would establish testbeds in which the future technologies that are the focus of their activities could be effectively explored in a realistic setting. The computing environment established by Xerox PARC is a good example of such a testbed.
- *Economic and societal impact studies.* Expeditions should include research on the impact of future technology on society. Such studies should involve researchers from disciplines other than computer science.
- *Education.* Each center should have an extensive educational program to inspire students of all backgrounds to think about the long-term technological future and its implications.
- *Outreach.* Expeditions should reach out to the research community, to industry, and to the public. The purpose of these programs is to accelerate the realization of the center's vision in the form of knowledge and products. Outreach to the research community should include open meetings where the results and effects of center activities are publicly presented.

Criteria for selecting proposed expeditions for funding should include:

- Elucidation of a clear and exciting vision for the future.
- Novelty of the technological area of exploration.
- Identification of significant computer science and engineering research challenges.
- A coherent plan for carrying out and managing the proposed activities.

The Committee recommends creating up to five Expeditions, to be selected by a competition to establish both a core activity, and a standard process to allow additional researchers to participate. Competitions should be held at regular intervals e.g., every three years, to ensure a flow of new ideas. Each of these “virtual centers” should include investigators from many research institutions. Each center would receive an initial funding for five years, renewable after five years for a full term of 10 years. To encourage truly aggressive efforts, very high annual funding levels should be possible, perhaps up to \$40 million per center.

Recommendation: Establish a program of Enabling Technology Centers.

The Committee also recommends establishment of centers of excellence in computer science and engineering research applied to particular applications of information and communications technology. These Enabling Technology Centers (ETCs), located at university and/or Federal research institutes, will provide integrated environments for academia, industry, and Government to focus on the application of next-generation information technology to important national problems. There are many application domains where information and communications technologies could make a difference. These include computational science and engineering; health care; delivery of Government services/Digital Government; crisis management; environmental monitoring; life-long learning; law enforcement and public safety; arts, culture, and the humanities; intelligent transportation systems; improving the quality of life for persons with disabilities; and distributed work (e.g. telecommuting, collaboration by geographically distributed teams).

Enabling Technology Centers should be focused on applied technology and development. Researchers at the centers would conduct R&D on information technology problems arising from the center's application domain, develop new curricula for students and mid-curricula for both students mid-career professionals, participate in testbeds, and identify barriers to more widespread adoption of information technology in the application domain of concern.

In addition to research and development, these centers should perform several functions, including:

- *Education and training.* Researchers will develop educational programs and curricula the intersection between computer science and engineers and a particular applications domain.
- *Testbeds.* Researchers will participate in the experimental deployment of information systems “in the field.”
- *Research on factors inhibiting deployment of information technology in the application domain.* Researchers may discover that the use of information technology in particular application domains is being slowed by legal and regulatory barriers, lack of end-user training, the absence of compelling cost-benefit analysis, a lack of technical standards or other mechanisms for interoperability, etc.
- *Community building.* The centers could help build communities of researchers, companies, Government officials, users, and other stakeholders by convening conferences and workshops, developing research agendas, and supporting intellectual infrastructure such as electronic print archives, collaboratories, databases, and case studies of successful and unsuccessful uses of information technology.

Funding for Enabling Technology Centers should be shared between mission-oriented agencies with an interest in a particular applications domain and an agency with a broad mandate to support information technology research (e.g., the National Science Foundation).

Grants need to be of sufficient duration and size to support a critical mass of researchers interested in a particular applications domain. The Committee recommends each center be funded on the Science and Technology Center funding model. The full term of an ETC would be 10 years. Annual funding of up to \$10M per center is recommended, with up to 15 centers simultaneously in operation. Competitions should be held every three years.

5.4 Annual Review

Recommendation: Establish an annual review of research objectives and funding modes.

Both the Coordination Committee and the President's Information Technology Advisory Committee (PITAC) should be instrumental in reviewing research objectives. PITAC's role in advising the President through NSTC serves to provide high-level private sector advice on information technology. The Coordination Committee will provide high-level advice from within the government.

The Coordination Committee, with advice from the Advisory Committee, should conduct an annual review of research programs to ensure that they are achieving the goals set out for them. In particular, the Coordination Committee should ensure that the modes of support proposed for the programs— centers, multiple-investigator interdisciplinary projects, and testbeds, along with a renewed emphasis on software and long-term research—are not being compromised and that these modes of research support are meeting the goals set out for them.

This review is intended to serve as a high-level check to ensure the Federal R&D portfolio is properly balanced, comprehensive, and properly coordinated. To maximize the opportunities for full and frank exchanges among the principals, the degree of formality and process associated with the preparation for and conduct of this review should be strictly limited. Given that the purpose is to coordinate a scientific endeavor, the review might more closely resemble a scientific workshop than a traditional committee meeting.

5.5 Funding

Budget Recommendation: Increase current funding for IT R&D as follows over the fiscal years 2000-2004.

Table 5.1 summarizes the recommended funding increments for each of the four research priority areas identified in this report. Descriptions of the research topics and associated funding tables were provided in Chapters 3 and 4 of this report. The funding recommendations are in addition to those for current (FY 1999) Federal IT R&D programs and represent, rather, new or expanded activities.

Table 5.1
Recommending Funding Increases for IT R&D (\$ millions)

	FY 2000	FY 2001	FY 2002	FY 2003	FY 2004
Software	112	268	376	472	540
Scalable Information Infrastructure	60	120	180	240	300
High End Research	180	205	240	270	300
High End Acquisitions	90	100	110	120	130
Socioeconomic	30	40	70	90	100
Total	472	733	976	1192	1370

Can the Proposed Increases be Invested Effectively?

One question that arose during the public comment on the interim report was whether the IT research community could effectively invest an increase of the size we are proposing. To answer that question, we have developed a simple investment model described here. Note that this is not a proposal on how invest the funds, but rather an exercise to determine whether we can reasonably expect the community to be able to take on the challenge.

There are three principal ways that funding increases might be used to affect change in the research community:

- To add new researchers taken from the pool of graduating Ph.D.'s over the next five years.
- To upgrade the support base for existing faculty.
- To provide shared infrastructure for the community, typically supercomputers and high-bandwidth networking.

New researchers will be critical to increase the capacity for both research and teaching in IT. The current Ph.D. production in the field is approximately 1000 per year. Of these, less than half the choose academic careers. If research investments can increase that number by half or more, the Nation can add 2500 new researchers to the academic base over the next five years. The new funds can be used to add an additional 2500 researchers and

to upgrade the support base for the current IT research community. We assume a range of annual support levels from \$100K-\$250K per entry-level researcher within that community. Assuming 1000 new researchers at \$250K and 1500 new researchers at \$100K, we arrive at a total of \$400M per year at the end of five years.

Our analysis also establishes that there is unused research capacity in the existing research community. Currently the top-rated universities have an average support level of \$400K to \$500K per researcher per year, while universities not ranked in the top five have average support levels substantively lower, in the range of \$150K per researcher per year. One goal might be to bring the next twenty-five departments up to the quality and support levels of the top five and another fifty departments up to about \$250K per researcher per year. Assume that 25 organizations (with about 40 principal researchers per organization) will need to have their support base per principal researcher increased from about \$150K on average to about \$500K. We also assume that the next 50 organizations (also with about 40 principal researchers per organization) will need to have their support base increased from an average of \$100K to \$250K per principal researcher. These target levels are quite reasonable in terms of today's costs. For example, \$250K will typically cover the cost of two graduate students, a programmer, some equipment, and summer salary for the researcher. The above calculation leads us to estimate an expenditure of \$650M per year after five years.

These numbers reflect the vital importance of upgrading the current research base to improve the overall effectiveness of current funding and to insure that the universities remain competitive with the IT industry in providing computational and networking environments sufficient to be attractive to new graduate students. Without this latter investment, the needed increase in researchers may not materialize in the universities. These investments will insure that the universities can attract the needed researchers and that the existing research base can be made competitive with the state-of-the-art in industry.

Following the recommendations of the report, we have budgeted \$330M for large shared infrastructure such as supercomputers, external networks and testbeds to link the researchers and associated maintenance. This leaves us with the following totals, which are consistent with the budgets proposed in the report, for the fifth year:

Fund entry-level researchers	\$400M
Provide Supercomputers and external network infrastructure	\$330M
Upgrade 25 sites to \$500K per researcher support base	\$350M
Upgrade 50 sites to \$250K per researcher support base	\$300M
Total	\$1380M

Conclusion

In both the public and private sectors today, U. S. investments in technology R&D have slowed to a relative trickle. American businesses, in an ever-shrinking and more highly competitive world, have devoted less and less of their precious resources to long-term R&D, directing their efforts instead to reducing costs and getting new products into the pipeline today at the expense of the future. The Federal government has mirrored this trend because of dramatically increasing pressures on the research and development budgets, with only modest increases in funding levels. Funding agency managers have responded by making the natural and correct decision to favor the short-term needs of their missions over the need for long-term research in information technology. The U.S. Government's lead in research in high end computing and computation—so crucial to keeping our military edge in the competitive era of the Cold War—seems to have come down along with the Berlin Wall.

As a result, the once robust technological edge the U.S. has enjoyed over the rest of the world is actually built on an increasingly fragile technological substructure. If the trend away from long-term research continues, the flow of bold new ideas that has fueled the information economy in this decade is likely to slow to a trickle by 2010. To keep its competitive edge, the United States must rededicate itself to cutting edge high-tech research and development—or risk being overcome by nations with a clearer plan and a stronger view of the future. This is a risk the Nation cannot afford to take.

The initiatives proposed in this report would represent a major step toward restoring long-term Federal research in information technology to levels that will ensure continued prosperity and new technological solutions to national problems in the next millennium. The time to act is now and the Federal government has a unique role to play in supporting research in this critical area. The Brooks/Sutherland report stated this well:

Very few companies are able to invest for a payoff that is 10 years away. Moreover, many advances are broad in their applicability and complex enough to take several engineering iterations to get right, and so the key insights become 'public' and a single company cannot recoup the research investment. Public investment in research that creates a reservoir of new ideas and trained people is repaid many times over by jobs and taxes in the information industry, more innovation and productivity in other industries, and improvements in the daily lives of citizens. This investment is essential to maintain U.S. competitiveness.¹³

This committee strongly recommends that the Federal government embark upon the kind of leading-edge, visionary research necessary to continue the revolution that has transformed the lives of

¹³ Brooks and Sutherland, 23.

our citizens in ways not thought possible just thirty years ago. The recommendations of this committee also stress the needs to: 1) to upgrade the knowledge base and skills of our workforce, so that our citizens will be prepared to face a new century fully prepared for the technological challenges that are yet to come; and 2) to give all American the opportunity to participate in the information age, so that our citizens will be able to fulfill its promise. These steps, if taken now, will bring handsome returns to the Nation over the coming decades.

For additional copies of this report, please contact:

National Coordination Office for Computing,
Information, and Communications
4201 Wilson Boulevard, Suite 690
Arlington, VA 22230
(703) 306-4722
<http://www.ccic.gov>

