



Networking Health: Prescriptions for the Internet

Committee on Enhancing the Internet for Health Applications: Technical Requirements and Implementation Strategies, Computer Science and Telecommunications Board, National Research Council
ISBN: 0-309-51560-2, 388 pages, 6 x 9, (2000)

This PDF is available from the National Academies Press at:
<http://www.nap.edu/catalog/9750.html>

Visit the [National Academies Press](#) online, the authoritative source for all books from the [National Academy of Sciences](#), the [National Academy of Engineering](#), the [Institute of Medicine](#), and the [National Research Council](#):

- Download hundreds of free books in PDF
- Read thousands of books online for free
- Explore our innovative research tools – try the “[Research Dashboard](#)” now!
- [Sign up](#) to be notified when new books are published
- Purchase printed books and selected PDF files

Thank you for downloading this PDF. If you have comments, questions or just want more information about the books published by the National Academies Press, you may contact our customer service department toll-free at 888-624-8373, [visit us online](#), or send an email to feedback@nap.edu.

This book plus thousands more are available at <http://www.nap.edu>.

Copyright © National Academy of Sciences. All rights reserved.
Unless otherwise indicated, all materials in this PDF File are copyrighted by the National Academy of Sciences. Distribution, posting, or copying is strictly prohibited without written permission of the National Academies Press. [Request reprint permission for this book](#).

NETWORKING HEALTH

PRESCRIPTIONS FOR THE INTERNET

Committee on Enhancing the Internet for Health Applications:
Technical Requirements and Implementation Strategies

Computer Science and Telecommunications Board
Commission on Physical Sciences, Mathematics, and Applications
National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C.

NATIONAL ACADEMY PRESS • 2101 Constitution Avenue, NW • Washington, DC 20418

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

Support for this project was provided by the National Library of Medicine under Task Order No. 42, Sponsor Award No. N01-OD-4-2139. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsor.

Library of Congress Cataloging-in-Publication Data

Networking health : prescriptions for the Internet / Committee on Enhancing the Internet for Health Applications: Technical Requirements and Implementation Strategies, Computer Science and Telecommunications Board, Commission on Physical Sciences, Mathematics, and Applications, National Research Council.
p. cm.

Includes bibliographical references and index.

ISBN 0-309-06843-6 (casebound)

1. Internet (Computer network) in medicine. I. National Research Council (U.S.).

Committee on Enhancing the Internet for Health Applications: Technical Requirements and Implementation Strategies.

R859.7.I58N48 2000

362.1'0285'4678—dc21

00-008698

Additional copies of this report are available from:

National Academy Press
2101 Constitution Avenue, NW
Box 285
Washington, DC 20055
(800) 624-6242
(202) 334-3313 (in the Washington metropolitan area)
<http://www.nap.edu>

Copyright 2000 by the National Academy of Sciences. All rights reserved.

Printed in the United States of America

THE NATIONAL ACADEMIES

National Academy of Sciences
National Academy of Engineering
Institute of Medicine
National Research Council

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce M. Alberts is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. William A. Wulf is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce M. Alberts and Dr. William A. Wulf are chairman and vice chairman, respectively, of the National Research Council.

**COMMITTEE ON ENHANCING THE INTERNET FOR HEALTH
APPLICATIONS: TECHNICAL REQUIREMENTS
AND IMPLEMENTATION STRATEGIES**

EDWARD H. SHORTLIFFE, Columbia University, *Chair*
RUSS BIAGIO ALTMAN, Stanford University
PATRICIA FLATLEY BRENNAN, University of Wisconsin at Madison
BRUCE DAVIE, Cisco Systems, Inc.
WILLIAM M. DETMER, University of Virginia
VALERIE FLORANCE, Association of American Medical Colleges
ANDREW FRIEDE, Cerner Corp.
MARK FRISSE, Express Scripts, Inc.
JOHN GLASER, Partners Healthcare System, Inc.
JOHN HUFFMAN, Stentor, Inc.
ISAAC KOHANE, Children's Hospital, Boston
CARL E. LANDWEHR, Mitretek Systems
DANIEL R. MASYS, University of California at San Diego
JANE E. SISK, Mount Sinai School of Medicine
THORSTEN VON EICKEN, Cornell University

Staff

JERRY R. SHEEHAN, Study Director
RITA GASKINS, Project Assistant
MICKELLE RODGERS RODRIGUEZ, Senior Project Assistant

COMPUTER SCIENCE AND TELECOMMUNICATIONS BOARD

DAVID D. CLARK, Massachusetts Institute of Technology, *Chair*
JAMES CHIDDIX, Time Warner Cable
JOHN M. CIOFFI, Stanford University
W. BRUCE CROFT, University of Massachusetts at Amherst
SUSAN L. GRAHAM, University of California at Berkeley
JUDITH HEMPEL, University of California at San Francisco
JEFFREY M. JAFFE, Lucent Technologies, Inc.
ANNA KARLIN, University of Washington
BUTLER W. LAMPSON, Microsoft Corporation
EDWARD D. LAZOWSKA, University of Washington
DAVID LIDDLE, U.S. Venture Partners
TOM M. MITCHELL, WhizBang! Labs, Inc.
DONALD NORMAN, UNext.com
RAYMOND OZZIE, Groove Networks
DAVID A. PATTERSON, University of California at Berkeley
CHARLES SIMONYI, Microsoft Corporation
BURTON SMITH, Tera Computer Company
TERRY SMITH, University of California at Santa Barbara
LEE SPROULL, New York University

Staff

MARJORY S. BLUMENTHAL, Director
HERBERT S. LIN, Senior Scientist
JERRY R. SHEEHAN, Senior Program Officer
ALAN S. INOUE, Program Officer
JON EISENBERG, Program Officer
GAIL PRITCHARD, Program Officer
JANET D. BRISCOE, Office Manager
DANIEL LLATA, Senior Project Assistant
SUZANNE OSSA, Senior Project Assistant
MICKELLE RODGERS RODRIGUEZ, Senior Project Assistant
DAVID DRAKE, Project Assistant
MARGARET MARSH, Project Assistant
BRANDYE WILLIAMS, Office Assistant

**COMMISSION ON PHYSICAL SCIENCES,
MATHEMATICS, AND APPLICATIONS**

PETER M. BANKS, Veridian ERIM International, Inc., *Co-chair*
W. CARL LINEBERGER, University of Colorado, *Co-chair*
WILLIAM F. BALLHAUS, JR., Lockheed Martin Corporation
SHIRLEY CHIANG, University of California at Davis
MARSHALL H. COHEN, California Institute of Technology
RONALD G. DOUGLAS, Texas A&M University
SAMUEL H. FULLER, Analog Devices, Inc.
JERRY P. GOLLUB, Haverford College
MICHAEL F. GOODCHILD, University of California at Santa Barbara
MARTHA P. HAYNES, Cornell University
WESLEY T. HUNTRESS, JR., Carnegie Institution
CAROL M. JANTZEN, Westinghouse Savannah River Company
PAUL G. KAMINSKI, Technovation, Inc.
KENNETH H. KELLER, University of Minnesota
JOHN R. KREICK, Sanders, a Lockheed Martin Company (retired)
MARSHA I. LESTER, University of Pennsylvania
DUSA M. McDUFF, State University of New York at Stony Brook
JANET NORWOOD, Former Commissioner, U.S. Bureau of Labor
Statistics
M. ELISABETH PATÉ-CORNELL, Stanford University
NICHOLAS P. SAMIOS, Brookhaven National Laboratory
ROBERT J. SPINRAD, Xerox PARC (retired)

MYRON F. UMAN, Acting Executive Director

Preface

Considerable attention centers on the Internet and health care. The popular press tends to focus on the growing use of the Internet to support consumer health via health-related Web sites that provide information on specific diseases, contain guidance on healthy lifestyles, host chat and support groups, and sell a range of health-related products. Federal programs tend to focus on using networking of various kinds to support telemedicine, especially for rural and underserved areas. These applications represent just a small sampling of the ways the Internet can be used to support health and health care. Many other applications exist in public health, biomedical research, health care finance and administration, and the maintenance of electronic health records.

Each of these applications demands different capabilities of its underlying networks, whether high-bandwidth connections, rapid delivery of data, tight security, reliability, or widespread access. As a result, the existing Internet cannot support them all. Nevertheless, a number of programs, such as the federal government's Next Generation Internet and the private-sector Internet 2 initiatives, are under way to enhance the capabilities of the Internet and to develop technologies that support high-performance networking. These programs could help the Internet meet the needs of the health and health care communities. But what capabilities must the Internet provide in order to support health and health care? How do they differ from those that might be developed anyway to support other applications of the Internet in sectors as diverse as commerce, entertainment, and defense?

THE COMMITTEE AND ITS CHARGE

To obtain preliminary answers to these questions, the National Library of Medicine (NLM) asked the Computer Science and Telecommunications Board (CSTB) of the National Research Council (NRC) to conduct a study on the technical capabilities the Internet needs to support health applications and on ways of ensuring that these capabilities are implemented. The objectives were to (1) identify a range of health applications that could reasonably gain widespread, routine use over the Internet in the foreseeable future, (2) examine the technical capabilities these applications would demand, (3) define the characteristics of the Next Generation Internet and an associated infrastructure that would meet these requirements, and (4) recommend an appropriate strategy for achieving this infrastructure in light of other activities under way to enhance current Internet capabilities. The project was intended to address questions such as the following:

- What characteristics of the current Internet limit its utility for various types of routine health and biomedical research uses? For different types of telemedicine? For plausible future health-related uses that will require high bandwidth? What would be the likely future levels of various types of health traffic on the network if these limitations were overcome?
- What quality of service and security characteristics or tools will the Next Generation Internet require to be suitable for various types of health-related applications, given technical requirements and estimated traffic levels? Are the requirements of different health applications compatible with one another and with the requirements for other projected uses of the Internet? To what extent are they supported by technology available now or already in development?
- What specific strategies are likely to ensure that the United States attains a communications infrastructure to support the full potential range of routine health and biomedical research uses?

To conduct the study, CSTB assembled an expert committee consisting of 15 members drawn from the networking and health communities. The committee met five times between September 1998 and June 1999 to solicit testimony from outside experts, deliberate on its findings and recommendations, and draft its final report. It met again in September 1999 to discuss its plans for modifying the draft report in response to comments from many outside reviewers. The committee also conducted a series of site visits to gather information firsthand on the ways the Internet could be—and is being—used to support health and health care. Members of the committee visited with researchers and health practitioners at

Stanford University, NASA Ames Research Center, the University of California at San Francisco, Kaiser-Permanente of Northern California, East Carolina University, the University of North Carolina at Chapel Hill, the University of Washington, Regence BlueShield, and the Washington State Department of Health. These visits provided an opportunity to directly observe Internet-based systems that had been developed for health care, biomedicine, and other health-related activities. It also provided an opportunity to learn more about the kinds of applications that cannot yet be implemented in an Internet-based system.

The committee used the information gathered during site visits and presentations by other briefers to synthesize a comprehensive view of the ways the Internet could transform health-related activities. It attempted to identify the technical and nontechnical challenges that need to be overcome in order to expand the use of the Internet in health and biomedicine and to devise a set of recommendations that will help make the Internet a more useful communications medium within the health sector. The site visits are summarized in Appendix A. Specific material from the visits has been incorporated into Chapter 2 of this report.

ACKNOWLEDGMENTS

This report benefited from the combined talents of many people, including those who were directly associated with the project and many who were not. First, thanks are due to members of the committee itself, all of whom maintained a high level of enthusiasm, energy, and dedication over the course of the project. Committee members found time for project meetings, site visits, and drafting portions of the text despite their many other responsibilities and commitments. Approximately half of the committee members changed jobs and/or affiliations as the project unfolded—a living testament to the dynamic nature of the Internet in health applications—yet they remained committed to this project throughout.

Many other people volunteered their time and expertise to help the committee to better understand the ways in which the Internet might be used to support health objectives and the technical capabilities that health applications demand of the Internet. Special thanks are due to those who hosted, coordinated, and participated in its site visits and to those who met with the committee at its three open meetings (see Appendix D for a list of participants). The information gathered during these interactions proved invaluable to the committee's deliberations and forms the backbone of this report. Other people also provided useful information and advice to the committee. Mark Ellisman at the University of California at San Diego and Martin Hadida-Hassan at the San Diego Supercomputer Center helped to clarify the discussion of telemicroscopy in Chapter 2 of

this report; Stewart Streimer, John Parmigiani, and Sandy Haydock from the Health Care Financing Administration provided updates on the agency's security policy, pilot programs for electronic submission of data, and information technology strategy, respectively. Grant Miller and Yolanda Comedy, from the National Coordination Office for Computing, Information, and Communications, supplied information and funding data on the federal government's NCI initiative. Kenneth Birman at Cornell University contributed valuable insight into security and reliability concerns associated with use of the Internet in health applications. Donald Simborg described early efforts to network computers in health care organizations.

The committee stands in awe of the remarkable, patient work of the CSTB and NRC staff in supporting its deliberations over the course of this study. Staff members kept the committee on track and helped its members to put their ideas and analyses into coherent prose. The committee is further indebted to the reviewers of an early draft of this report, whose thoughtful comments and criticisms challenged committee members to strengthen and refine their arguments—and to articulate them more clearly. This final report is considerably improved thanks to their input. It is a much more readable document thanks to Laura Ost, a free-lance editor who assisted the NRC's internal staff in editing the manuscript, and to James Igoe, from the National Research Council Library, who tracked down numerous references and helped to complete the citations in the reference lists.

Finally, thanks are due to Donald A.B. Lindberg, director of the NLM, Michael Ackerman, assistant director for High Performance Computing at the NLM, and Betsy Humphreys, associate director for Library Operations at the NLM. Their dedication to improving health, health care, and biomedical research through Internet technologies and their financial support made this project possible. The committee hopes that its findings and recommendations will assist them in leading the health and the networking communities to achieve their vision of a "healthier" Internet.

Edward H. Shortliffe, *Chair*
Committee on Enhancing the
Internet for Health Applications:
Technical Requirements and
Implementation Strategies

Acknowledgment of Reviewers

This report was reviewed by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the NRC's Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the authors and the NRC in making the published report as sound as possible and to ensure that the report meets institutional standards for objectivity, evidence, and responsiveness to the study charge. The contents of the review comments and draft manuscript remain confidential to protect the integrity of the deliberative process. The committee wishes to thank the following individuals for their participation in the review of this report:

Dixie Baker, Science Applications International Corporation,
William Boebert, Sandia National Laboratories,
James Bradley, Abaton.com, Inc.,
Scott Bradner, Harvard University,
Charles Brownstein, Cross-Industry Working Team,
Paul Clayton, Intermountain Health Care,
Don Detmer, Cambridge University,
Mary Fennell, Brown University,
Thomas Ferrin, University of California at San Francisco,
Alan Garber, Stanford University,
John Halamka, CareGroup Healthcare System,
Michael G. Kienzle, University of Iowa,
Clement McDonald, Regenstrief Institute,

Satyanarayanan Mahadev, Carnegie Mellon University,
David Pryor, Allina Health System,
Thomas Rindfleisch, Stanford University,
Jay Sanders, Global Telemedicine Group,
Elliot Stone, Massachusetts Health Data Consortium,
Peter Szolovits, Massachusetts Institute of Technology,
Elizabeth Ward, Foundation for Health Care Quality, and
Betsy Weiner, University of Cincinnati.

Although the individuals listed above provided many constructive comments and suggestions, responsibility for the final content of this report rests solely with the authoring committee and the NRC.

Contents

EXECUTIVE SUMMARY	1
1 OVERVIEW AND INTRODUCTION	27
A Systems Perspective, 29	
The Internet and Health, 33	
Drivers of Internet Applications in Health, 35	
Impediments to Broader Adoption of the Internet, 36	
Technical Considerations, 38	
Networking Alternatives, 41	
Enhancing the Internet, 45	
The Next Generation Internet Initiative, 46	
Private-Sector Efforts: Internet 2 and Abilene, 50	
Deploying Enhanced Internet Technologies, 51	
Organization of This Report, 52	
References, 53	
Notes, 55	
2 HEALTH APPLICATIONS OF THE INTERNET	57
Consumer Health, 58	
Consumer-Oriented Health Web Sites, 59	
E-mail Between Patients and Providers, 62	
Online Health Records, 64	
Patient Monitoring and Home Care, 66	
Technical Requirements for Consumer Health Applications, 69	

Clinical Care, 71	
Remote Consultation, 72	
Medical Imaging, 76	
Clinical Transactions, 80	
Technical Requirements for Clinical Care, 87	
Financial and Administrative Transactions, 88	
Technical Requirements for Financial and Administrative Applications, 93	
Public Health, 94	
Public Health Surveillance, 96	
Integrating Data Sources for Improved Decision Making, 98	
Responding to Bioterrorist Attacks, 99	
Technical Requirements for Public Health Applications, 100	
Professional Education, 102	
Graduate Education, 102	
Continuing Education, 105	
Technical Requirements for Health Professional Education, 107	
Biomedical Research, 108	
Biomedical Databases, 109	
Linked Simulations, 112	
Remote Control of Experimental Apparatus, 113	
Publication on the Internet, 116	
Collaboration Among Researchers, 118	
Clinical Research, 120	
Technical Requirements for Biomedical Research, 121	
Summary, 123	
Bandwidth, 123	
Latency, 124	
Availability, 124	
Security, 124	
Ubiquity, 124	
References, 126	
Notes, 128	
3 TECHNICAL CHALLENGES	132
Quality of Service, 133	
Increasing Bandwidth, 135	
Differentiated Services, 138	
Integrated Services, 140	
Alternative Quality of Service Options, 141	
Quality of Service Policy, 141	
Multicast, 143	

Security, 144	
Elements of Security, 145	
Firewalls, 148	
Security Protocols, 150	
Access Controls, 157	
Network Availability, 160	
Broadband Technologies for the Local Loop, 162	
Privacy-Enhancing Technologies, 167	
Anonymous E-mail, 169	
Protected Web Browsing, 170	
Anonymous Payment, 172	
Anonymous Data Released from Sensitive Databases, 172	
Conclusion, 173	
Bibliography, 174	
Notes, 176	
4 ORGANIZATIONAL CHALLENGES TO THE ADOPTION OF THE INTERNET	178
Lessons from Other Industries, 179	
Advancing the Strategic Interests of Health Care, 181	
Impediments to Adopting Internet Applications, 184	
Barriers to Change, 185	
Uncertainties Surrounding Internet Strategies, 189	
Establishing Organizational Leadership for Information Technology, 197	
Summary, 199	
References, 200	
Notes, 201	
5 ISSUES FOR PUBLIC POLICY	202
Protection of Personal Health Information, 203	
Access to Information Infrastructure, 209	
Intellectual Property Protection, 215	
Electronic Publishing, 215	
Distance Education, 217	
Regulations Affecting Electronic Delivery of Health Services, 219	
Payment Policies, 219	
Liability and Licensure, 221	
Federal Support for Health-Related Information Technology Research, 223	
Workforce Issues, 227	
Conclusion, 230	
References, 230	
Notes, 232	

6	CONCLUSIONS AND RECOMMENDATIONS	235
	Conclusions, 236	
	Recommendations, 249	
	Research, Development, and Deployment of Needed Technical Capabilities, 250	
	Demonstration and Evaluation of Health Applications of the Internet, 257	
	Addressing Educational Needs, 261	
	Addressing Policy Issues, 263	
	A Final Word, 265	
	References, 266	
	Notes, 268	
	APPENDIXES	
A	Site Visit Summaries	271
B	National Library of Medicine Awards to Demonstrate Health Applications of the Next Generation Internet	314
C	Biographies of Committee Members	334
D	Individuals Who Participated in Site Visits or Briefed the Study Committee	342
	INDEX	345

NETWORKING HEALTH

Executive Summary

The Internet has great potential to improve Americans' health by enhancing communications and improving access to information for care providers, patients, health plan administrators, public health officials, biomedical researchers, and other health professionals. Ongoing research and development (R&D) efforts, such as the federal government's Next Generation Internet (NGI) initiative and the complementary Internet 2 program of the private sector, could help to realize that potential. Such efforts promote the creation and deployment of new networking technologies to enhance the Internet's capabilities, enabling a growing range of applications in health and other sectors. But what technical capabilities do health applications demand of the Internet? How do these capabilities differ from those needed by applications in other sectors, such as banking, defense, and entertainment? What types of experiments and demonstrations should be undertaken now to learn quickly about the requirements and benefits of different health applications of the Internet? And how can the health community ensure that its needs are considered within the networking research community and in standards bodies that are defining future capabilities?

Questions of this nature prompted the National Library of Medicine (NLM) to request a study by the Computer Science and Telecommunications Board of the National Research Council that would evaluate the technical capabilities demanded by health applications of the Internet. As the health community's primary representative in the NGI initiative and a longtime supporter of R&D focusing on health applications of informa-

tion technology (IT), NLM sought advice on which capabilities should be deployed in the NGI testbed networks and, ultimately, the Internet. It recognized that the potential for health applications of the Internet had contributed to policy discussions of information infrastructure for several years but that progress in realizing that potential had been slower than in other economic sectors. This report responds to the NLM request by examining applications of the Internet in six health-related areas: consumer health, clinical care, health care financing and administration, public health, professional education, and biomedical research. It draws on a series of visits by members of the committee to organizations that are actively designing, developing, and in some cases operating networked applications. It identifies the technical capabilities that these applications demand of supporting networks and makes recommendations regarding the capabilities that need to be deployed to enable the health community to take fuller advantage of the Internet. It also identifies additional work that is needed to develop complementary and appropriate information technologies, such as tools to help consumers evaluate the quality of the information they find on the Internet and access controls to reliably limit Internet users' ability to access resources such as patient medical records.

But the report does not focus exclusively on networking technologies, since the capabilities needed in networks are intertwined with other technical, organizational, and policy considerations. As the committee learned during its site visits, an adequate communications infrastructure is not the only prerequisite for expanded Internet use within the health community. Efforts are also needed to surmount organizational and policy impediments to the adoption of the Internet and Internet-based applications. At present, health care organizations are ill prepared to deploy Internet-based applications, because they lack information upon which to base investment decisions, face an uncertain financial environment, and have difficulty attracting the talent needed to design, develop, and implement such applications. A number of public policy issues, ranging from concerns about patient privacy to the lack of payment mechanisms for some medical consultations delivered remotely, also stand in the way of greater deployment of Internet applications. All of these issues need to be addressed if health organizations are to take advantage of the capabilities offered by an enhanced Internet.

HEALTH APPLICATIONS OF THE INTERNET

The most visible examples to date of the Internet's role in health-related activities are in the consumer domain. Tens of thousands of sites on the World Wide Web (the Web) offer information on health topics, and a growing number of companies have established Web sites to provide

consumers with information on specific diseases, therapies, and healthy lifestyles. Some sites allow consumers to evaluate risks to their health, manage chronic medical conditions, purchase health-related products, pose questions to health professionals, or engage in discussions with other consumers. These systems take advantage of the Internet's broad, public reach to engage significant portions of the online population, often with information that is specially tailored to their needs. An estimated 30 million users searched for health information on the Internet in 1999 alone, and in 1998 consumers and students—as opposed to practitioners and researchers—accounted for roughly 30 percent of the use of the NLM's MEDLINE system, which contains references to millions of journal articles (Lindberg, 1998).

Although health-related Web sites garner considerable media attention, they represent only a small sampling of the ways in which the Internet can be used in health, itself a large sector embracing health care, public health, health education, and biomedical research. Because the Internet, in theory, can link all the participants in the health community, it can be used to improve consumer access to health information and health care, to enhance clinical decision making and improve health outcomes by making better information available to clinicians on demand, and to reengineer the processes of care to make them more efficient. The Internet can also be used to improve the education of medical professionals, enhance public health surveillance, and facilitate biomedical research. In each of these domains, specific applications can be envisioned in which the Internet is used to transfer text, graphics, or video files (and even voice); control remote medical or experimental equipment; search for needed information; and support collaboration, in real time, among members of the health community (Table ES.1). For example, the Internet could do the following:

- Enable consumers to access their health records, enter data or information on symptoms, and receive computer-generated suggestions for improving health and reducing risk;
- Allow emergency room physicians to identify an unconscious patient and download the patient's medical record from a hospital across town;
- Deliver care instructions to a traveling businessperson who begins to feel chest pains while in a hotel room;
- Enable homebound patients to consult with care providers over real-time video connections from home, using medical devices capable of transmitting information over the Internet;
- Support teams of specialists from across the country who wish to

TABLE ES.1 Primary Technical Challenges and Limiting Technical Factors in Selected Health Applications of the Internet

Application Domain	Class of Application		
	Real-Time Video Transmission	Static File Transfer	Remote Control
Consumer health	Remote medical consultations to the home, office, or wherever the patient is located.	Accessing personal health records online. Downloading educational videos. Sending periodic reports on health conditions to a care provider.	Remote control of patient monitoring equipment.
Clinical care	Remote medical consultations between clinician and patient or between two clinicians.	Transfer of medical records and images (e.g., X rays, MRI, CT scans).	Remote and virtual surgery (a long-term possibility being examined by the defense and space communities).
Administrative and financial transactions	Videoconferencing with real-time sharing of documents.	Payment of services, enrollment of patients, quality reviews, etc. Large medical records and images may be transmitted in support of some claims.	N/A
Public health	Videoconferencing among public health officials during emergency situations, such as chemical or biological attacks by terrorists.	Incident reporting. Collection of information from local public health departments and laboratories. Surveillance for emerging diseases or epidemics. Transfer of epidemiology maps or other image files for monitoring the spread of a disease.	N/A

Information Search and Retrieval	Real-Time Collaboration	Primary Technical Challenges
Online searching for health information or self-assessment guides. Looking for a doctor or hospital.	Collaboration with care providers. Participation in chat groups and support groups.	Protection of sensitive patient information from breaches of confidentiality and from corruption. Ubiquity of access so that all health care consumers can be reached at the location at which care is needed. Tools and policies for validating the quality of online information.
Practice guidelines. Searches of professional medical literature.	Consultation among care providers, such as for surgical planning, which may involve manipulation of digital images.	Access to sustained bandwidth and low latency for remote consultations and collaboration. Security of clinical records. Network reliability. Ubiquity of access for care providers.
Consumer access to information about health plans, participating practitioners, eligibility for procedures, covered drugs in formulary.	N/A	Security to ensure confidentiality and integrity of records. Network reliability sufficient to support regular use for business transactions. Standards for data exchange and definitions of data elements.
Access to published literature and research results as well as epidemiological data. Delivery of alerts and other information to practitioners or other health workers.	Videoconferencing among public health officials during emergency situations, such as chemical or biological attacks by terrorists.	Security to ensure confidentiality and integrity of laboratory reports and other public health information that may contain personal identifying information. Network reliability. Security from information warfare or attacks on the network's physical infrastructure.

continued

TABLE ES.1 Continued

	Class of Application		
	Real-Time Video Transmission	Static File Transfer	Remote Control
Professional education	Distance education: either real-time transmission of lectures or on-demand streaming video with integrated graphics. Real-time consultations with experts about difficult cases.	Accessing electronic medical records from remote clinics. Downloading sets of reference images or prerecorded videos of lectures.	Simulations of surgical procedures. Virtual environments for exploration of three-dimensional environments.
Biomedical research	Visual feedback from remote instrumentation. Online conferences. Collaboration among distant researchers.	Transferring large data sets between computers for high-speed computation and comparisons. Reviewing results of remote experiments. Searching archives of three-dimensional medical images.	Controlling experimental equipment, such as electron microscopes.
Limiting Technical Factors	Availability of sustained, predictable, high-bandwidth connections to many locations, including rural health clinics and patients' homes (to support remote consultations).	Authentication of source and recipient of information. Security of personally identifiable information in transit across the network and in storage at either end of the network. Availability of sustained high-bandwidth connections for transfer of large, time-critical files.	Network latency and bandwidth. Ability to obtain guaranteed bandwidth for predictable periods of time.

Information Search and Retrieval	Real-Time Collaboration	Primary Technical Challenges
Accessing reference materials and course materials.	Virtual classrooms. Distributed collaborative projects. Distributed discussions.	Sufficient bandwidth to accommodate large numbers of transactions from a single educational institution or to support access to remote scientific and clinical simulations. Ubiquity of access for students in remote clinical rotations and to support educational applications in the home.
Searching remote databases and professional literature.	Collaboration among researchers. Peer review. Interactive virtual conferences.	Sufficient bandwidth to support rapid transfers of large sets of data for distributed simulations. Low latency to accommodate remote control of equipment.
Tools for locating information of interest and for determining the quality of retrieved information. Means of allowing anonymous searches.	Sustained access to high-bandwidth, low-latency networks for collaborations involving real-time video or manipulation of images. Multicast protocols to make more efficient use of networking resources.	

plan particularly challenging surgical procedures by manipulating shared three-dimensional images and simulating different operative approaches;

- Allow a health plan to provide instantaneous approval for a referral to a specialist and to schedule an appointment electronically;
- Enable public health officials to detect potential contamination of the public water supply by analyzing data on nonprescription sales of antidiarrheal remedies in local pharmacies;
- Help medical students and practitioners access, from the examining room, clinical information regarding symptoms they have never before encountered; and
- Permit biomedical researchers at a local university to create three-dimensional images of a biological structure using an electron microscope a thousand miles away.

A number of these applications have been demonstrated in localized settings, such as individual hospitals or health care delivery systems. For reasons of technology, organizational capabilities, and public policy, many of them have yet to be deployed more broadly across the Internet or on private networks that rely on dedicated communications links. As a result, little is known about their costs and benefits—whether they would improve health or research capabilities, how much they would cost to implement, or whether they would reduce health costs if deployed on a larger scale. That kind of knowledge will require continued exploration and evaluation, as well as an understanding of how the economics of the large but decentralized health sector can influence the development of the Internet, driving decisions about which capabilities will be deployed, and when.

This report addresses a broad spectrum of health applications in an attempt to demonstrate the diversity of needs and the degree of commonality in the technical capabilities they require. It is intended to guide a faster realization of the Internet's potential for health, a potential that has eluded the health sector for too long. The report recognizes that the applications themselves—and the technical capabilities they demand—are moving targets with uncertain trajectories. While today's demonstration programs hint at the kinds of capabilities that will be needed in the future, the evolutionary path of health applications of the Internet is unclear. Will, for instance, remote medical consultations become viable between any patient and any care provider connected to the Internet, or will this capability remain more localized in its reach and limited to patients and providers in the same health plan? The answer depends on technical, economic, social, and policy considerations that are difficult to predict, and different answers could drive the need for significantly different technical capabilities, as well as a different scale and scope of

deployment. The report attempts to recognize these uncertainties and to derive conclusions that are reflective and cognizant of them.

TECHNICAL CONSIDERATIONS

The technical capabilities needed to support health-related use of the Internet vary considerably from one application to another. The relative importance of bandwidth, latency, availability, security, and ubiquity in six different classes of health application is shown in Table ES.2 (see Box ES.1 for a definition of the technical terms used in this report). For the most part, these considerations are common to Internet applications in other sectors, and that broader base increases the likelihood of affordable solutions. But in communicating with Internet researchers and technology developers, the health community (i.e., all those active in health-related activities, such as provision of care, public health, professional education, and biomedical research) can call attention to its need for particular attributes, and it can point out the characteristics of the health sector that differentiate its needs from those of sectors such as entertainment, defense, or finance.

For example, security is a primary concern in virtually all health applications of the Internet because the extreme sensitivity of personal health information demands high levels of confidentiality. Furthermore, the paramountcy of safety—individuals' health and lives are at stake, after all—requires that information not be corrupted before, during, or after transmission across the network from one party to another. Although security is also important in many other Internet applications, including electronic commerce (e-commerce, itself a player in the evol-

TABLE ES.2 Technical Demands of Health-Related Applications of the Internet

Application Area	Bandwidth	Latency	Availability	Security	Ubiquity
Consumer health	++	+	++	++++	++++
Clinical care	++++	+++	++++	++++	++
Financial and administrative transactions	+	+	+++	++++	++
Public health	+	+	+++	+++	++
Professional education	+++	++	++	+	+++
Biomedical research	++++	+++	++	++	++

NOTE: Plus signs (+) denote the relative importance of the technical feature within the designated application area. A single plus sign denotes minimal importance; four plus signs signify great importance.

BOX ES.1 Glossary of Technical Terms

A range of technical capabilities must be considered in determining the suitability of different networking technologies for particular applications. Among the more important are five that are emphasized throughout this report:

1. *Bandwidth* is the data-carrying capacity of a network, usually expressed as the number of bits per second that can be transmitted across a particular link or the network as a whole.

2. *Latency* is the time required for an individual packet of data to be transmitted between communicating entities on a network. A related concept is *response time*, which refers to the time required for an entire message or file to be transferred across the Internet and acknowledged.

3. *Availability* is the likelihood that the network is available for service and functioning properly. Availability can be compromised by the failure of individual components or network links, by hostile attacks that overload the system, or other causes discussed in Chapter 3.

4. *Security*, as used in this report, refers to the capability of a network to ensure the confidentiality and integrity of information transmitted across it. An important part of ensuring confidentiality is authenticating the identity of participants in a network-based transaction.

5. *Ubiquity* is the degree of access to a network. The telephone system is highly ubiquitous because access can be achieved by almost anyone in the United States from almost any location. Access to private networks is, by design, less ubiquitous because it is constrained to a limited number of people and/or a limited number of geographic locations.

Related to the first two of these terms is *quality of service* (QOS), which refers to the capability of a network to provide a range of guarantees about its performance, measured in terms of sustained bandwidth, latency, and/or packet loss rates. The current Internet contains no provisions for QOS, offering only best-effort delivery of packets of data, although several protocols have been developed for implementing QOS.

ing health environment), health applications pose special challenges, the solutions to which may lie in the computers attached to the network rather than in the network itself. For example, the exchange of electronic medical records, payment data, or prescription information demands that the identities of both the sender and recipient of the data be validated (authenticated) with high levels of assurance. Mechanisms for authenticating individuals that are more secure than passwords are not in widespread use across the Internet. This situation has not, however, impeded consumer-oriented e-commerce applications, because online vendors

have robust means of authenticating themselves to their customers' Web browsers (using electronic certificates provided by a handful of certificate authorities, as described in Chapter 3). Moreover, the vendors do not necessarily require strong authentication of users who present a valid credit card number: credit card companies and vendors who accept credit cards expect to incur some costs from fraud, and consumer losses are generally capped at nominal levels. By contrast, health has a low tolerance for losses and other kinds of mistakes: before an electronic prescription can be filled or a copy of an electronic medical record sent, the identity of the requester must be verified as rigorously as the identity of the supplier. The constantly shifting relationships among health organizations further complicate security considerations. Other aspects of security also present challenges in health applications, as outlined in Chapters 2 and 3 of this report.¹

Network availability is also important in health applications of the Internet. High levels of availability are needed in mission-critical applications in many industries, and similar needs obtain in health: if insurance companies and managed care organizations are to rely on the Internet for claims processing, referrals to specialists, or checks on eligibility for particular services, they must be sure the network will be running when needed and that data will not be corrupted. But the health sector's need for high levels of network availability to and from a large number of possible locations can also be greater than in other sectors, because health, well-being, and even life may be at stake. If care providers are to use the Internet to access electronic patient records when treating patients in the emergency room, they must know that the network and the applications are operational 24 hours a day, 7 days a week. Accordingly, health applications add to the call for the Internet to be made resistant to malicious attacks and resilient in the face of failures of hardware, software, or human operators.

Many of the applications that can be envisioned in the health domain demand high levels of bandwidth or timely delivery of data, often for an extended period of time.² Consider the case of remote medical consultations, which could make expert care more equitably available across the country, regardless of the location of the patient. Video consultations demand high-bandwidth connections (roughly 384 kilobits per second) in both directions between two communicating sites for the duration of the session—as long as 30 minutes in some cases—even if one of the sites is a small medical practice or a patient's home. Although the backbone networks that make up the Internet have sufficient capacity to accommodate such needs, they cannot currently guarantee that adequate bandwidth and latency will be available whenever needed, because other traffic with unknowable bandwidth needs will also be traversing the network. These

types of applications therefore demand mechanisms for ensuring quality of service (QOS) across the Internet, whether by allowing users to subscribe to higher-end services (referred to as differentiated services, or diff-serv) or allowing them to reserve capacity on an as-needed basis (referred to as integrated services, or int-serv).

The Internet Engineering Task Force (IETF) has codified standards for both diff-serv and int-serv, but neither has yet been deployed across the Internet. Moreover, it is not clear that Internet service providers (ISPs) will deploy them in the near future—or in ways that support the health industry. For example, the highly decentralized nature of the health industry implies that health organizations will obtain their Internet service from many different ISPs rather than from a single provider. To provide QOS between the many different sets of communicating parties that are possible, QOS mechanisms would need to be deployed through the entire Internet, not just across a single ISP's network. However, mechanisms do not yet exist for supporting QOS (either diff-serv or int-serv) between ISPs, precluding the possibility of end-to-end QOS guarantees any time soon. Furthermore, the protocols for supporting int-serv will not necessarily scale sufficiently to allow their use across the Internet. The decentralized structure of the health industry makes it hard for health organizations to come up with viable business models whereby they can pay ISPs to deploy the kinds of QOS they need. Almost any solution to this problem will require the participation of the insurance companies and other third-party payers who finance health care in the United States.

Ubiquity of access is particularly important in health applications of the Internet because people in need of health care and related services can be almost anywhere. Indeed, the most significant advance in health care brought about by the Internet may prove to be better access for care providers, consumers, and administrators operating in relatively isolated environments. Although many near-term applications that extend to individual consumers do not require high-bandwidth connections, future applications—whether remote medical consultations or the downloading of educational videos—could easily drive a need for ubiquitous, broadband access technologies. Of particular interest could be broadband technologies for residential access that provide sufficient bandwidth both upstream (from the end user to the Internet) and downstream (from the Internet to the end user). Most existing residential broadband technologies—such as cable modems and digital subscriber line service using the telephone network—allocate much more bandwidth downstream than upstream, consistent with a view of the Internet as a mechanism for distributing content from a centralized source (such as an entertainment company) rather than facilitating collaboration and interaction among multiple participants.

ORGANIZATIONAL BARRIERS TO THE ADOPTION OF INTERNET APPLICATIONS

A handful of pioneering health organizations are developing and deploying innovative applications of the Internet, but such capabilities are diffusing slowly throughout the sector—in traditional care provider organizations, in particular. A number of factors have impeded the broader deployment of Internet-based systems within the health sector, including the structure of the sector itself. Despite some consolidation over the past decade, the sector is very diverse and decentralized and marked by local solutions to problems—it has been characterized as a “trillion-dollar cottage industry.” As a result, effecting wide-scale change can be difficult, as is achieving a unified voice on issues of technology and its application.

Further slowing adoption is a paucity of reliable information on the costs and benefits of Internet-based applications in operational settings. How much will Internet-based systems cost to deploy, operate, and maintain? How will they improve care and/or reduce costs? How well can Internet-based systems be integrated with legacy databases in large health care organizations? Health care professionals tend to be cautious in adopting unproven technologies because of the overwhelming need to ensure patient safety and positive health outcomes. Organizations that pay for health care (including traditional insurance companies and the Health Care Financing Administration, or HCFA, which processes Medicare and Medicaid payments) also want evidence of cost savings or medical effectiveness if they are to pay for services based on the new technology. Although Internet applications have been demonstrated to improve efficiency in some applications run across enterprises, Internet technology is still fairly new and untested in health care applications, making evaluations and comparisons difficult and prompting caution in the pursuit of Internet strategies.

Evaluating the costs and benefits of health applications of the Internet is made more difficult by the uncertainties surrounding the effects of Internet-based communications on relationships among the numerous entities involved in health care. For example, little is known about the ways in which the Internet will alter the traditional relationships among patients, primary care physicians (PCPs), medical specialists, and hospitals. Will the Internet change the way consumers seek care, enabling them to learn enough about their health to bypass PCPs and go directly to specialists, or will they be confused by all the information and need to consult their PCP more often? Will Internet-based care improve management of the chronically ill, and if so, how will it affect the cost structure of health care organizations? Restructuring may also be needed within indi-

vidual organizations. As has happened with other applications of information technology, Internet applications have been shown to alter work patterns within organizations in unanticipated ways. What types of skills will information systems staffs need to develop and implement Internet-based systems for health care? What types of skills will administrative staff and health professionals need to work with Internet-based health care systems? What type of training do intended system users need? Answers to these questions will come only after additional experimentation and evaluation.

Internet applications also tend to demand new (or modified) organizational policies and procedures. For example, when should electronic mail (e-mail) be used between patients and care providers? What types of liability does an organization assume for the quality of the information that is relayed in the online discussion groups it hosts? How can patient privacy be protected in electronic transactions, and what balance between security and access is acceptable to consumers who want online access to their health records? Progress is being made on a number of these issues (e.g., the American Medical Informatics Association has developed a set of guidelines for clinical uses of e-mail, and the Department of Health and Human Services has promulgated draft regulations governing the privacy and security of electronic health information), but new issues continue to arise and managers of health organizations are struggling to keep up, at times slowing the broader deployment of new applications.

PUBLIC POLICY ISSUES

Public policy influences the ways in which health organizations can use the Internet to achieve their goals. For example, state-based practices for licensing health care professionals and resolving malpractice suits hamper efforts to provide remote medical consultations across state lines. Uncertainties over evolving federal regulations for the privacy and security of electronic health information continue to deter organizations from implementing systems for sharing health records or administrative and financial information across the Internet. Other issues, such as the protection of intellectual property contained in materials developed for educational purposes, affect a broad base of constituents, including some in the health community. So does the issue of unequal access to the information infrastructure, the so-called digital divide. Reports show that people in different geographic regions and socioeconomic classes and with different levels of educational attainment have considerably different degrees of access to the Internet (NTIA, 1999). Such differences are alarming in a number of contexts—the delivery of government services and educational opportunities among them—but take on added significance in a health

care setting, where limited access to the information infrastructure could exacerbate the existing differences in access to quality health care. Furthermore, many of the near-term remedies proposed for enhancing Internet access for the general public—such as the wiring of schools, libraries, and community centers—do not necessarily translate well to health care, because consumers may be reluctant to conduct transactions in public settings and may need access outside normal business hours. Policy issues such as these have to be resolved if health applications of the Internet are to become more pervasive and more effective. Their resolution will require the health community to become more actively engaged in the policy-making process to ensure that health-related interests are addressed.

EVOLVING THE INTERNET TO MEET HEALTH NEEDS

Before Internet use can become widespread throughout the health community, action is needed in four areas: (1) research on, and the development and deployment of, technologies suitable for health applications of the Internet, (2) continued demonstration and evaluation of health applications of the Internet, (3) educational needs of health care organizations and their workers, and (4) resolution of policy issues that impede the use of the Internet in health applications. The committee's recommendations are intended to guide efforts in each of these areas, and progress is needed in all of them. As a group, the recommendations recognize that what differentiates health from other sectors is the juxtaposition of exacting technical requirements with a vast geographic expanse and a highly decentralized industrial structure and economic base. This combination exacerbates the difficulties arising from the failure to articulate technology needs or to devise means to ensure that needed services are provided. By advocating the needs of a broad constituency, NLM can provide timely leadership in enhancing the Internet for health.

Research, Development, and Deployment of Technical Capabilities

Broadening the utility of the Internet for health applications demands that the Internet possesses the necessary technical capabilities. This can be done by deploying the technologies that will soon be available for improving security, availability, QOS, and ubiquity across the Internet and by continuing to research and develop improved capabilities over the long term. Efforts must also be made to ensure that the health community's needs are relayed to the networking research community and that advances are made in complementary technologies that will enable health organizations to take advantage of the networking infrastructure. These

efforts need to reflect the many uncertainties surrounding health applications of the Internet and their technical needs.

Recommendation 1.1. The health community should ensure that technical capabilities suitable for health and biomedical applications are incorporated into the testbed networks being deployed under the Next Generation Internet initiative and eventually into the Internet.

As a first step toward enhancing the Internet to support health applications, the health community should push to have the capabilities described below deployed in the testbed networks being constructed under the federal government's NGI initiative. Without these capabilities, future health applications could be thwarted or delayed and the opportunities the Internet offers could be lost. While the entire nation has a stake in ensuring that these capabilities are deployed, it is the health community itself that is in the best position to identify the capabilities it needs, to communicate them to the network research and development community, and to help shape the business case that will impel their deployment. The networks being deployed under NGI will support a range of experimental health applications, such as remote medical consultations, collaboration among practitioners and researchers, and access to online repositories of information (see Appendix B for a listing and brief description of ongoing NLM projects). The testing of technical capabilities in these testbed networks will provide an opportunity for evaluations and refinements that will be incorporated into the demonstration projects, enabling the health community to better assess the capabilities its applications demand. Those that prove effective should be deployed in the public Internet as they become more stable. These technologies are described in more detail in Chapters 3 and 6 of the report.

- *Quality of service.* QOS protocols should be deployed across the NGI testbed networks so that users are guaranteed access to needed capabilities (e.g., bandwidth and latency). A number of academic medical centers will have access to the NGI via their universities and have received funding for projects to demonstrate a variety of applications that demand high bandwidth—from remote medical consultations to real-time transmission of high-resolution radiological or biological images. The deployment of differentiated services would allow users to experiment with premium services that could eventually be offered across the Internet. The deployment of integrated services would allow further experimentation with protocols for reserving capacity as needed for particular events and would allow further evaluation of the scalability of existing proto-

cols. By experimenting with these protocols, users may be able to better understand the specific capabilities required and devise business models that will support the deployment and effective use of the new protocols across the public Internet.

- *Security.* Both Secure Socket Layer (SSL) encryption and IPsec (IPSec) should be deployed in the NCI testbed networks to allow the continued evaluation of different modes of securing transactions across the Internet. Although SSL is already in widespread use across the Internet, the broader deployment of IPsec would provide a complementary means of protecting information exchanges among organizations, and it might prove effective for financial and administrative exchanges among affiliated organizations. Before either of these protocols (especially SSL) can be used successfully in health applications, a public key infrastructure must be established, along with the technical mechanisms needed to support stronger authentication of all parties involved in transactions across the Internet. Such mechanisms are generally lacking across the Internet, although there are enclaves where they are used in the private sector and within the federal government. More research is needed to develop means of authenticating large numbers of users, many of whom need to communicate securely despite having no established relationships.

Recommendation 1.2. To ensure that the Internet evolves in ways supportive of health needs over the long term, the health community should work with the networking community to develop improved network technologies that are of particular importance to health applications of the Internet.

Continued research will be needed to make the Internet even more capable of supporting health—and other—applications in the long term. The technologies of most interest to the health community include the following:

- *More readily scalable techniques to guarantee bandwidth on demand.* Existing protocols for providing QoS on demand across the Internet, such as the integrated services model, may not scale sufficiently to allow widespread use. To enable applications such as remote consultation, new protocols will be needed.

- *Stronger forms of authentication.* Continued effort will be needed to find ways of identifying participants in Internet transactions, including participants who have not previously communicated with each other. The new techniques will need to scale to cover all Internet users and be simple to administer. Work on smart cards, token-based authentication,

and biometric authentication devices should be pursued. This kind of R&D may fall under two headings—high-confidence systems and NGI research—within the federal government’s portfolio of information technology research programs.

- *Symmetric or dynamically reconfigurable broadband technologies for the last mile.* Users of residential-grade access technologies (e.g., cable modems and digital subscriber lines) will need either access to a more balanced allocation of bandwidth into and out of their homes or the capability to reconfigure the allocation as needed to support applications such as remote medical consultations, which could extend to many small health clinics, places of employment, and patients’ homes.

- *Hardened quality-of-service guarantees.* Mechanisms will be needed to ensure that critical applications in health (and other sectors) do not lose QOS guarantees except in extreme circumstances, such as a major network outage. One area of interest is techniques for rapid reconvergence after link failures to ensure that new paths across the Internet are found quickly in the event that a particular link fails. Many parties and sectors want QOS guarantees, and many Internet users are responding to offers by large ISPs who make such guarantees within their own large networks. The challenge for health (which serves a dispersed national population) is finding a way to improve QOS and availability of service across multiple ISPs.

- *Disaster operations.* Techniques are needed for delivering mission-critical, health-related traffic even in a major natural or man-made disaster.

Recommendation 1.3. The National Library of Medicine should forge stronger links between the health and networking research communities to ensure that the needs of the health community are better understood and addressed in network research, development, and deployment.

The diverse and decentralized nature of the health sector impedes the development of a unified voice through which it can express its needs to those involved in networking research, development of Internet standards (e.g., by the IETF), and deployment of Internet services. The NLM, by virtue of its leadership in health informatics, could play a more pronounced role in this area, actively forging links between the health and networking communities, which have historically had limited interaction. This could be done in several ways, perhaps by providing special funding to recipients of NLM grants and contracts that would support their participation in conferences and meetings of the networking community or by funding projects that explicitly involve researchers from the health and networking communities. Additional activities would undoubtedly

be needed to help the health community find ways to more effectively identify and communicate its needs to the networking community. An ad hoc task force could be set up to explore additional ways to accomplish this goal. The NLM itself could work more closely with the networking community, leveraging its long-standing attention to information technology development and speaking for the health sector as a whole. It could also serve as a focal point for contact with ISPs, the business entities responsible for deploying the capabilities that would benefit health. It could advance the perspective that health is a leading example of a peer-to-peer application (as opposed to the more asymmetric application associated with many other kinds of content distribution) that requires advanced networking services from the Internet, helping create a more unified voice for the decentralized health sector.

Recommendation 1.4. The National Institutes of Health and its component agencies should fund information technology research that will develop the complementary technologies that are needed if the health community is to take advantage of the improved networking technologies that can be expected in the future.

Health applications of the Internet pose a number of challenges for information technology research on topics other than networking. The National Institutes of Health and its constituent centers and agencies should pursue research in those areas that are of particular importance to the health community, such as (1) validation of information retrieved from the Internet, (2) tools for protecting the anonymity of Internet users, (3) access controls governing the ability of many different types of users to access different resources on the network, (4) controls on the secondary distribution of information, (5) improved capabilities for auditing the logs of accesses to databases and information, (6) QOS policies that are suitable for health and health care applications, and (7) applications that are alert to QOS offerings and that use them appropriately. Other technical needs will undoubtedly emerge as new applications are developed and gain acceptance within the health community. The constantly changing context of the Internet implies that the set of applications will evolve and that the need for research will remain.

Demonstration and Evaluation of Health Applications

Continued experimentation and evaluation will be key to the development of a better understanding of the types of health applications that may become popular on the Internet and of the technical capabilities they

demand. Through demonstrations of applications such as remote consultation, remote control of experimental equipment, and online access to electronic medical records, members of the health community will gain an opportunity to examine the relative costs and benefits of these applications, the business models needed to support them, and the organizational policies needed to govern their use. A number of public and private organizations have supported programs to allow these types of demonstrations. Such efforts need to continue as new Internet technologies become available and new applications are envisioned. Demonstrations will serve as venues for continued identification of technical needs that the networking community can address and other problems and issues for the health community to resolve. The process will be increasingly important to the health community if it is to establish a dialog with the Internet community about evolving needs and technical requirements and if it is to leverage that dialog to grow capabilities from the confines of a demonstration to widespread deployment. To provide information that will inform this dialog, a number of parallel efforts will be needed, as recommended below:

Recommendation 2.1. The Department of Health and Human Services should fund pilot projects and larger demonstration programs to develop and demonstrate interoperable, scalable Internet applications for linking multiple health organizations.

Pilot projects are needed to explore the full range of health uses of the public Internet, particularly projects that link multiple distinct organizations in an operational context. They could include projects to allow the patients of one organization to obtain remote consultations with specialists at other organizations or to allow the transmission of financial and administrative information among organizations that provide, pay for, and manage health care. Few health care organizations have a strong incentive to implement such systems on their own, given the significant uncertainties surrounding the effectiveness of different Internet-based systems in health care, the fragmented and proprietary nature of the industry, and the scale at which such systems would need to be built. Federal funding could play an important role in stimulating such work, especially if it focused on applications that link multiple organizations.

Recommendation 2.2. Federal agencies such as the Department of Veterans Affairs, the Department of Defense, the Health Care Financing Administration, the National Institutes of Health, and the Indian Health Service should serve as role models and

testbeds for the health industry by deploying Internet-based applications for their own purposes.

Federal agencies that operate large-scale health care programs should, whenever possible, attempt to be leading-edge users of Internet technologies. By doing so, they could not only demonstrate the feasibility of deploying different health applications but also provide a testbed for developing needed standards and supporting technologies. The Department of Defense already has a sizeable program under way for delivering health care at a distance (i.e., telemedicine), the Department of Veterans Affairs has a network of hospitals that share patient information as needed, and the HCFA processes Medicare and Medicaid claims. Each of these programs, as well as those of the Indian Health Service, could serve as a testbed for Internet applications while helping to fulfill important government missions. Additional support might come from other ongoing efforts to reengineer federal activities.

Recommendation 2.3. Health organizations in industry and academia should continue to work with the Department of Health and Human Services to evaluate various health applications of the Internet in order to improve understanding of their effects, the business models that might support them, and impediments to their expansion.

Work is needed to evaluate the effectiveness of different forms of Internet-based health care and to compare their effectiveness against applications run across different network infrastructures. Health care organizations have little evidence or data on which to base their decisions about Internet strategies. Because such evaluations would benefit a wide range of health-related organizations, not just those directly involved in the studies, active federal support would be justified.

Recommendation 2.4. Public and private health organizations should experiment with networks based on Internet protocols and should incorporate the Internet into their future plans for new networked applications and into their overall strategic planning.

By using networks that incorporate Internet protocols—whether the Internet protocol suite per se or those associated with the Web—health organizations could gain a better understanding of the capabilities and trade-offs inherent in the use of the Internet for health applications without exposing themselves to the associated risks and uncertainties. Using

these protocols locally would also prepare health organizations to take better advantage of the Internet—and the continued advances in its abilities—once technical tools are in place to make it safer and reliable enough for health applications.

Addressing Educational Needs

Wider deployment of Internet-based applications in health care will require that organizations in the health sector adopt, adapt, and extend Internet technologies to fit their missions and develop the internal capabilities to do so. The Internet promises to radically transform the provision of health care and the education of health professionals, and organizations that fail to take steps now may find themselves ill prepared when improved Internet technologies become available. To make better use of the Internet, health care organizations will also have to learn how to evaluate the benefits of Internet technologies and develop effective policies for guiding their use, just as they had to learn how to use earlier and more localized forms of information technology, an effort in which the health care system is still lagging. Efforts are recommended in three areas:

Recommendation 3.1. Professional associations with expertise in health issues and information technology should work with health care organizations to develop and promulgate guidelines for safe, effective use of the Internet in clinical settings.

Part of the challenge of Internet use in health care is the development of suitable policies, practices, and procedures to guide its use. For example, how should providers handle e-mail from patients to ensure timely responses, maintenance of patient confidentiality, and the incorporation of necessary information into the medical record? How can care providers be sure of the identity of a patient to whom they are sending e-mail? What is the role of a health care organization in monitoring discussion groups that operate under its initiative or that of affiliated care providers? Health care organizations have little experience upon which to base such policies, but they can learn from each other's experiences. Professional associations have a significant role to play in helping define industrywide guidelines for safe, effective use of the Internet. The American Medical Informatics Association has developed guidelines for clinical uses of e-mail (Kane and Sands, 1998). Similar guidelines on other topics would support industry efforts to develop Internet-based systems.

Recommendation 3.2. Government, industry, and academia should work together and with professional associations with

experience in health and information technology to educate the broader health and health care communities about the ways the Internet can benefit them.

One obstacle to the greater use of the Internet in health care is that health workers at all levels (care providers, administrators, and information systems staff) do not fully appreciate the ways in which the Internet can improve the provision and administration of health care. The growing amount of publicity for e-commerce and even consumer-health-information Web sites does not translate into the kinds of institutional and procedural changes that would make the most of Internet capabilities in health care. Educational outreach programs would create a more receptive audience for new technologies. Academic health centers and professional associations have unique capabilities to educate members of the health community.

***Recommendation 3.3.* The Department of Health and Human Services should commission a study of the health information technology workforce to determine whether the supply of such workers balances the demand for them, to identify the kinds of training and education that workers at different levels will need, and to develop recommendations for ensuring an adequate supply of people with training at the intersection of information technology and health.**

The process of developing, deploying, and evaluating health applications of the Internet demands workers with a solid understanding of the Internet, of other information technologies, and of the processes involved in health care. The policy community has already expressed concern about a perceived shortage of skilled information technology workers.³ Anecdotal evidence indicates that similar concerns may apply to the field of health informatics, and in June 1999 DHHS announced its Biomedical Information Science and Technology Initiative, which would boost the pipeline of people educated as computational biologists. However, there is little documentation with which to evaluate these concerns or to project the types of IT skills that workers at different levels within a health organization will need. Additional study would be required to determine the extent of the problem and the best way of solving it.

Resolving the Policy Issues

Public policy issues that impede Internet-based activities in health, health care, and biomedical research need to be addressed. These include

issues specific to the provision of health care services over the Internet, such as payment for services, professional licensure, and liability, as well as issues of patient/consumer privacy, intellectual property protection, and equitable access that extend far beyond the health domain. Such issues could stand in the way of use of the Internet in health care and in the education of health professionals. Accordingly, although the committee was not constituted with the range of expertise needed to make recommendations for solving these problems, the report offers the following recommendation for advancing the debate on these policy issues:

Recommendation 4.1. The Department of Health and Human Services should more aggressively address the broad set of policy issues that influence the development, deployment, and adoption of Internet-based applications in the health sector.

Ensuring that the Internet evolves in ways that meet the needs of the health care community and enabling the health sector to better take advantage of these capabilities will require the continuous coordination of many independent activities and stakeholders in the public and private sectors. The concerns and needs of the health community must be reflected in efforts to resolve national policy issues such as intellectual property protection, privacy, and access to information infrastructure, and specific efforts are needed to ensure that policy issues of concern only to the health community, such as licensure of care providers, payment policies, federal funding for health informatics research, and the supply of health information technology workers, are addressed. While many of these issues are being addressed by various elements of the federal government—including agencies within DHHS—other issues have seen little input from the health community. The constituent agencies of DHHS vary in the importance they attach to these policy issues and in their approaches to resolving them. Strong, stable leadership is essential to keep these policy-related activities focused and sustained.

DHHS should assert itself more aggressively in this arena. Private-sector organizations also have significant leadership roles to play, but their effectiveness in bringing about industrywide change can be limited because the private sector is so highly decentralized. DHHS is not the only federal agency with responsibilities in health (the Department of Veterans Affairs, the Department of Defense, the Indian Health Service, and the National Aeronautics and Space Administration all have health-related programs), but the breadth of its programs and its mission argues for it to play the lead role within government for coordinating Internet-related activities, especially as they relate to the health community.

The establishment of a data council within DHHS and the realignment of the National Committee on Vital and Health Statistics into an advisory committee on health data, statistics, and national health information policy are positive steps that should be built upon. They have enabled DHHS to make significant strides in policy areas such as the development of regulations for protecting electronic health information. There are other roles for DHHS to play in this effort: (1) providing strategic leadership for Internet-related efforts within the department and its constituent agencies (this would include the use of the Internet in support of department and agency missions) and coordinating them with those of other federal agencies, (2) convening public and private bodies to identify, examine, and propose mechanisms for addressing issues related to the Internet and health care, (3) exploring cross-cutting issues that affect many health agencies and developing programs for addressing them (e.g., implementing a public key infrastructure that would support a range of federal health activities), (4) encouraging federal health agencies to share information and perspectives on their many responsibilities and interests, including the provision of care, payment for care, monitoring of care, health-related research, and public health, (5) advancing national debate about key information technology issues that affect health care, including the technical, organizational, and policy issues identified in this report, and (6) creating the organizational structures needed to ensure that issues at the nexus of health and information technology are identified and addressed promptly and efficiently. Although these activities will not by themselves resolve the issues, they will set in motion processes that can lead to a resolution.

LOOKING FORWARD

These recommendations are intended to help the nation move forward on technical, organizational, and policy fronts so that it can reap the benefits of the Internet for health applications. Additional work will be needed to identify other networking technologies of interest to the health community and to ensure that related information technology needs are met. This report prescribes the actions needed now to develop a truly *healthy* Internet in the future.

REFERENCES

- Kane, Beverley, and Daniel Z. Sands. 1998. "Guidelines for the Clinical Use of Electronic Mail with Patients," Report for the AMIA Internet Working Group, Task Force on Guidelines for the Use of Clinic-Patient Electronic Mail, *Journal of the American Medical Informatics Association* 5(1). Available online at <<http://www.amia.org/pubs/pospaper/positio2.htm>>.

- Lindberg, Donald A.B. 1998. Fiscal Year 1999 President's Budget Request for the National Library of Medicine. National Library of Medicine, Bethesda, Md., March 18. Available online at <<http://www.nlm.nih.gov/pubs/staffpubs/od/budget99.html>>.
- National Telecommunications and Information Administration (NTIA). 1999. *Falling Through the Net: Defining the Digital Divide*. U.S. Department of Commerce, Washington, D.C.

NOTES

1. As an example, a large number of individuals may have legitimate needs to review a patient's medical records, making the determination of access rules extremely complicated. In an emergency room situation, information may have to be accessed by a care provider with whom the patient has had no prior relationship, perhaps even at a hospital the patient has never visited.
2. Timeliness is not critical in many health care functions, such as when information is transmitted for review at a later time. But in some cases, such as acute trauma and remote consultation, timeliness can be important.
3. CSTB has a project under way to examine issues related to the information technology workforce. Information is available online at <www.cstb.org>.

1

Overview and Introduction

The Internet is rapidly and radically transforming many aspects of society, reshaping industries from aircraft manufacturing to retailing by enabling the widespread sharing of information and creating new relationships between buyers and sellers of goods and services. Businesses now sell goods and services over the Internet, often dealing directly with customers rather than working through traditional distribution channels and intermediaries, tailoring products to match more closely the preferences of individual customers. Governments disseminate public information on World Wide Web sites, and consumers use the Internet to find information, communicate with friends and family, plan trips, shop, and pursue hobbies. Both the scope of applications and the number of Internet users will undoubtedly continue to grow as technologies improve and innovators continue to experiment with new online applications.

Health-related activities stand to benefit enormously from the Internet. As a highly information-intensive set of functions characterized by complex interactions among a large number of stakeholders—primary care physicians, specialists, nurses, patients, health plan administrators, public health officials, medical librarians, researchers, and others—health-related activities can take advantage of the nearly ubiquitous reach of the Internet and its capability to support communication between users who may not have interacted with each other before. Already the Internet is beginning to influence the health sector by forging new relationships among stakeholders and improving access to health information. Its application in the delivery of health care, maintenance of public health,

payment for health care services, education of health professionals, and conduct of health sciences research could improve the quality of care and access to it as well as reduce its cost.

Despite its promise, the Internet's future in supporting health and health care is far from assured. A number of technical, organizational, and policy barriers stand in the way of its adoption by health organizations and consumers. Furthermore, although much can be done with the Internet in its present form, some health applications demand greater technical capabilities than the Internet can now provide, especially in the areas of security, reliability, and timely transmission of information. As a result, some health applications cannot be implemented across the Internet and used in operational settings without potentially threatening the privacy and optimal care of patients.

Health applications have helped motivate a number of efforts to improve the nation's information infrastructure.¹ Ongoing research and development (R&D) efforts, such as those being pursued under the federal government's Next Generation Internet (NGI) initiative and the private sector's Internet 2 initiative, also hope to foster technologies that could enhance the Internet's ability to meet the needs of the health sector. These efforts will also provide testbeds for improved evaluations of the benefits of different health applications of the Internet and their technical and nontechnical requirements. But these testbeds—and ultimately the Internet itself—will not adequately support health applications unless a better understanding is developed of the technical capabilities that these applications demand.

This report explores the use of the Internet in health-related applications and attempts to delineate the technical capabilities that such applications demand. Taking a broad view of health applications, it considers uses of the Internet in consumer health, clinical care, public health, medical education, health care financing and administration, and biomedical research.² It does not, however, attempt to predict which applications are most likely to catch on or to estimate levels of use; rather, it attempts to illustrate the types of applications that are possible and to assess the technical capabilities required for their safe, effective deployment in an operational setting.

The report also addresses organizational and policy issues that stand in the way of broader adoption of Internet technologies for health applications.³ It became increasingly apparent during the course of the study that health applications of the Internet involve systems that combine network infrastructure with other computing technologies (both hardware and software) and with end users who operate in multiple organizational contexts and are influenced by the policy environment. The close coupling among these levels makes it impossible to focus on any one level to the

exclusion of the others. Trade-offs are often made between the capabilities embedded in different levels of the system,⁴ and networking can make issues associated with other levels more important. Security, for example, takes on wholly new dimensions in a networked environment in which information can be readily transferred among entities and stored in computers that are attached to a public network. Yet, many of the mechanisms for addressing security concerns will be implemented not in the network itself but in the devices or computers attached to the network. An individual's access to health information in such an environment, and the circumstances under which such access is allowed, will be determined by a confluence of organizational and national policies for protecting health information.

The strong interrelationships between the network, other technology, and organizational and national policy introduce great uncertainties into the evolutionary path of the Internet with respect to health applications. For example, although many would agree that the Internet will enhance the role of the consumer in health care, the future of specific applications, such as remote medical consultations or online access to patients' medical records, is more difficult to discern because of the range of technical, organizational, and policy issues to be resolved (as detailed in later chapters of this report). Further research and experimentation are needed to understand these issues more fully and develop workable solutions. Consistent with the charge to the committee, this report does not attempt to resolve these policy issues, but by highlighting their significance in enabling effective and safe applications of the Internet for health care it may hasten their resolution. In the end, the report recommends ways of helping the Internet better serve a range of health interests. It identifies both long-term needs that will require R&D and steps that must quickly be taken to help people and organizations adopt and adapt to the next generation of Internet technologies. This chapter provides a broad overview of past and present uses of the Internet in health care; technical terms and considerations; and current R&D efforts that may advance the applications of the Internet and so improve health care.

A SYSTEMS PERSPECTIVE

An example may help to demonstrate both the potential value of the Internet in health care and the close linkages between networking technology, other information technology, and nontechnical issues. Consider the following hypothetical scenario:

Alice and Bob are recovering from a particularly virulent flu that kept them both out of work for the past week. They awaken

one snowy February night to hear their 6-year-old daughter, Charlotte, coughing, wheezing, and crying. She seems warm and will not be comforted. Alice and Bob are worried, but they have recently joined a plan that offers them the option of an in-home consultation. Because packing up their daughter and driving to the emergency room of the nearest hospital would take at least half an hour, they telephone the on-call pediatrician. After hearing the symptoms, the pediatrician decides to ask for basic measurements and have a quick look at Charlotte right away to decide whether she needs to be brought to the emergency room.

Alice turns on their Internet access device (a set-top box) and their television, while Bob sets up the home health assessment pack, including a digital thermometer, heart rate monitor, stethoscope, and video camera. Alice uses the keyboard to navigate to the health plan's Web site and inserts a smart card into the box that authenticates them to the health plan server. While they wait a few moments, their access device exchanges digital certificates authenticating both the server and their device and establishes an encrypted session with the server. Because videoconferencing will be used, the device also reserves a suitable level of bandwidth from Bob and Alice's Internet service provider to carry the quality of video needed for the consultation (a few hundred kilobits per second).

Once connected to the health plan Web site, a menu of options appears, and the couple make a video call to the pediatrician. A live image of the pediatrician appears in a video window. Alice transmits an authorization code to the pediatrician enabling her to access Charlotte's medical record from the online repository in which Alice and Bob maintain all their family medical records. The pediatrician asks them to take Charlotte's temperature and pulse and to position the microphone so that she can hear the child's breathing. Alice first uses the thermometer and heart rate monitor, which transmit results to the set-top box over wireless links. Guided by the pediatrician, Alice then places the stethoscope around various landmarks on Charlotte's chest and back to listen to the child's respirations. The pediatrician can see an image of Charlotte beamed to the set-top box from Bob's video camera. Alice and Bob can see a split-screen image on their television showing the pediatrician on one side and the image from their video camera on the other.

The pediatrician determines that Charlotte's condition does not require her to come in to the emergency room. From her remote observations, she concludes that the most likely diagnosis

is acute asthma. Charlotte has had two previous episodes of asthma during the past year, and in both cases she responded well to inhalants. The pediatrician asks the parents to administer a dose of the inhalant. Because it is possible to determine within 10 minutes whether the inhalant will work, the pediatrician opts to keep the video call running. Bob makes Charlotte comfortable, seating her within range of the video camera. During the ensuing 10 minutes, the pediatrician engages the parents in a brief review of the events leading up to the evening, exploring such things as exposure to dust and toxins as well as stress events in the family. Recalling that Charlotte's school has some major renovations under way, Alice asks the pediatrician about a possible connection between dust from the renovation and Charlotte's asthma flare-up.

The pediatrician guides Alice to the American Lung Association's Web site, and together they review the information about asthma in children. A checklist of environmental risk factors appears simultaneously on the screen, and the pediatrician and Alice review these together. Next they listen to an audio clip of various breath sounds, with the pediatrician coaching Alice on how to identify the distinctive sound of wheezing.

The pediatrician notes that Charlotte's breathing is easing, and the little girl is no longer crying. The pediatrician asks to speak to Charlotte and asks a few questions about how she feels. Charlotte points to her chest and says it feels tight. Noting that she is able to pronounce common words and that the audible wheezing has stopped, the pediatrician judges the situation to be under control and advises the family that Charlotte should be helped back to sleep.

The on-call pediatrician also recommends that an appointment be made for Charlotte to be seen by her own pediatrician the following afternoon. Bob navigates to the health plan's scheduling program and sets up the appointment. The site provides a map to the clinic that can be printed. The next day, as soon as she arrives at the clinic, Charlotte is welcomed and escorted into the examination room. While her doctor is finishing up another appointment, the nurse takes Charlotte's vital signs and adds the information to her electronic medical record, which is accessed from the computer in the examination room. Shortly thereafter, the doctor enters the room, reviews Charlotte's vital signs, examines her, and provides a diagnosis. Once the diagnosis and a prescription for a new inhaler are entered into the electronic record, a claim for payment is automatically filed with Charlotte's

health plan and an electronic prescription is sent to the pharmacy near her house. The medication will be waiting when Bob and Charlotte stop by on their way home.

This scenario identifies a number of benefits that Internet-based communications could bring to health care and related activities. It allows the patient (and her family) to avoid a potentially hazardous auto trip on a cold and snowy night and it eliminates waiting in an emergency room, during which time Charlotte could have been exposed to other infectious patients. In addition, the remote consultation allows rapid examination of the patient and preliminary evaluation (or triaging) of needs using several data sources (e.g., sound, vision, and instrumented sensors). Had Charlotte's condition been more serious, her parents could have been directed to take her directly to an emergency room; had her condition been less serious, the system could have enabled the family to avoid an office visit altogether. Although telephone-based services can produce similar benefits, they do not enable the clinician to examine the patient visually or with medical devices. Similarly, they are not as effective at allowing care providers to teach patients and their families to distinguish among various symptoms and at providing expert educational materials for understanding a particular condition. The electronic system also supports paperless billing, which could speed payment for services and reduce error and loss as information proceeds through the system of reviews and approvals. The system also allows easy, but protected, access to the patient's medical record to give the care provider more complete information when making a diagnosis and plan of treatment. The record can be updated easily in real time as new information is collected and can be made available to any care provider who needs it.

Of course, considerable effort would be required to transform such a scenario into a reality on a broad scale. A number of technical advances, related to both the networking infrastructure itself and the devices attached to it, would be required. For example, communication links into and out of homes would be needed that are sufficient to support color video of adequate resolution, and there must be suitable assurance that the video service will be available without significant interruption for the duration of the call. Smart cards would need to be issued to consumers to authenticate them to a health care site and support encryption for a session. This type of health care would also depend on electronic patient records, to which patients can grant providers access as needed and which can be updated during the course of a consultation. The equipment used would have to be reliable enough to create and sustain a connection between a family and a care provider for the duration of a consultation and to provide valid measurements of vital signs. Internet-compatible

medical devices would be needed to capture vital signs and transmit them to a remote physician. Nontechnical issues would need to be addressed as well. Families would have to be trained to properly use the system and the home medical equipment, such that care providers could be assured of receiving valid information remotely. Health plans would need policies on payment for remote consultations and on care providers' access to the electronic patient record. If any one of these capabilities was lacking, the system would fail.

THE INTERNET AND HEALTH

The health sector has a three-decade-long history of linking computers together to improve health care and administration. The National Library of Medicine (NLM) made its Medical Literature Analysis and Retrieval System (MEDLARS) available online to regional libraries over a time-shared network in the early 1970s. The resulting MEDLINE (for MEDLARS onLINE) system made the library's repository of biomedical references more widely available to support clinical decision making.⁵ Shortly thereafter, the first local area networks (LANs) were introduced at the University of Vermont Hospital to support clinical and administrative processes (Box 1.1).

Since these beginnings, the health care industry has gradually come to rely heavily on information technology (IT). In 1996, IT constituted 56 percent of the industry's total net capital stock—the fourth highest percentage out of 53 industries examined by the U.S. Department of Commerce (1999). Only the telephone and telegraph, radio and television, and securities and commodities brokerage industries were more IT-intensive. Nevertheless, health care expenditures on IT are relatively small in relation to the size of its labor force. The industry overall spent just \$543 per worker on IT in 1996, compared to \$12,666 for securities brokers and \$29,236 for telephone and telegraph industry workers; on this scale, health care ranked only 38 out of the 53 industries in the Commerce Department sample.⁶

Health is already a bustling area of activity on the Internet. Recent surveys indicate that more than 22 million Americans used the Internet to retrieve health-related information in 1998—a figure that was expected to grow to 33 million in 1999 (Davis and Miller, 1999). Other estimates place the number as high at 70 million (Morrison, 1999). Since it was made available to the public via the Internet in 1997, NLM's MEDLINE database, which contains more than 15 million abstracts and references from more than 3,900 medical journals, has experienced a surge in activity to 300,000 searches per day (Benton Foundation, 1999). Health is one of the more popular topics on the Internet, with estimates of the number of

BOX 1.1

Early Efforts in Networking Health

The first attempts to deploy communications networks in support of clinical records involved the use of local area networks at the University of Vermont Hospital in 1976 and at Walter Reed Army Hospital in 1977. These systems allowed users to log on to many computers from the same terminal, eliminating the need for multiple terminals at nursing stations, each connected to a different computer for a different function. Both projects used a technology pioneered by the Mitre Corporation called broadband, which at the time referred to coaxial cable similar to that used for cable television, by which multiple communication channels were carried across a single cable.¹ The system made use of frequency-division multiplexing to squeeze multiple channels onto a single cable.

Subsequent efforts at the University of California at San Francisco (UCSF) Medical Center—under the direction of Donald W. Simborg, who worked with Steve Tolchin of the Johns Hopkins University Applied Physics Laboratory—led to the development in 1979 of the first true back-end network. Four minicomputers were connected to the network to exchange transactions between the admitting office, the clinical laboratory, the pharmacy, and the radiology departments. The computers exchanged several core messages, including the synchronization of patient admission-discharge-transfer information, orders from clinical areas, and the display of results to the clinical areas. Unlike the earlier front-end networks, these networks did not require a user to be involved in the transaction; instead, the exchange of messages was handled by the computer applications themselves, using a protocol developed specifically for the system. The result was the creation of the first application-level data interchange protocol in health care.

Dr. Simborg left UCSF in 1984 to create the Simborg Systems Corporation, and a similar data interchange protocol was developed for his product. This commercial protocol was later placed in the public domain and became the core of the first version of the Health Level 7 (HL7) protocol, which today is the most widely used data interchange protocol in health care.

¹Today, the term “broadband” is used to refer to a range of technologies that offer high-bandwidth (i.e., high-data-rate) communications across telephone wires, coaxial cable, optical fiber, or wireless communications channels.

SOURCE: Donald Simborg, KnowMed Systems, Inc., personal communication dated October 31, 1999.

health-related Web sites running as high as 10,000 or more (Benton Foundation, 1999). Health-related Web sites allow consumers to search for information on specific diseases or treatments, pose questions to care providers, manage chronic diseases, participate in discussion groups, assess existing health risks, and purchase health-related products. By one estimate, the online consumer market will grow to \$1.7 billion by 2003,

fueled largely by online sales of products such as prescription and non-prescription medicines and vitamin supplements (Nash, 1999).

Beyond its popularity with consumers, the Internet is also used by health care professionals, biomedical researchers, and health care administrators. Web sites geared to health care professionals allow them to access the professional literature, consult with colleagues electronically, order medical supplies, or communicate with insurance companies.⁷ Biomedical researchers use the Internet to access online databases of journal articles and scientific information. Organizations involved in the provision of health care, whether individual hospitals, managed care plans,⁸ or integrated delivery networks (IDNs),⁹ have begun to use the Internet to reach out to consumers. Their Web sites provide information on available services and may allow consumers to change their enrollment status, select physicians, and schedule appointments electronically.

Drivers of Internet Applications in Health

The health applications available on the Internet today take advantage of the Internet's expansive reach to enable health care organizations to interact with a growing number of online consumers (Miller and Reents, 1998). Whereas just 17 percent of U.S. households had Internet access in 1997, roughly one-third did by 1998 (NTIA, 1999), and analysts predict that 90 percent of U.S. households will have Internet access by 2005 to 2010 (Rosenberg, 1999). As people become accustomed to using the Internet for routine activities, from electronic commerce (e-commerce) to homework, they are likely to use the Internet for health-related activities. Consumer experiences in other areas of Internet activity, such as e-commerce and electronic mail (e-mail), will influence the expectations they bring to online health applications (Mittman and Cain, 1999).

Care provider organizations face a number of pressures to integrate the Internet more effectively into their operations. Recent trends toward consolidation in the health care industry and the expansion of managed care have erased some of the impediments to sharing information among competing organizations. As they attempt to link individual practices, clinics, and hospitals into single entities, IDNs have a greater need to share information with affiliated institutions. As purchasers, accrediting bodies, and the general public increasingly hold managed care plans accountable for the quality of health care, plans have developed schemes for the electronic sharing of data on facilities' utilization rates and health-related outcomes. With such data, managed care plans can compile statistics on quality-of-care indicators and monitor the quality and costs of the individual care providers. As care provider networks grow and consumers become more mobile, the electronic transmission of patient information

among providers could improve care and reduce costs to the provider, the patient, and the managed care plan.¹⁰

Impediments to Broader Adoption of the Internet

Despite the flurry of Internet activity within and around health care, many potential applications have yet to be realized. Many organizations in the health sector continue to rely on private networks (e.g., leased lines) rather than the Internet for many data communications tasks, and some health-related applications have not yet been deployed across any type of communications network, public or private (Box 1.2). Few health care organizations, for example, have integrated the Internet directly into the provision of care. Remote medical consultations remain a novelty practiced by a few institutions, typically over dedicated networks, for a small subset of their patients and with support from external financial grants. Most public health offices remain unconnected to the Internet and there-

BOX 1.2 Representative Applications Conducted over the Internet and Private Networks

Functions Commonly Performed Today over the Internet

- Search for consumer health information
- Participate in chat/support groups
- Exchange electronic mail between patients and care providers (limited)
- Access biomedical databases and medical literature
- Find information about health plans, select physicians (limited)
- Purchase pharmaceuticals and other health-related products

Functions Performed Today over Private Networks

- Transfer medical records among affiliated health organizations
- Transfer claims data to insurers and other payer organizations
- Conduct remote medical consultations (limited)
- Send medical images (X rays, etc.) to remote site for interpretation (very limited)
- Broadcast medical school classes over campus networks (limited)

Functions Not Commonly Performed Today over Either the Internet or Private Networks

- Videoconferencing among public health officials
- Remote surgery or guidance of other procedures
- Public health surveillance/incident reporting
- Home-based remote medical consultations
- In-home monitoring of patients

fore are unable to accept electronic reports from testing laboratories or communicate health information over the Internet to neighboring jurisdictions. Private insurers have in general not adopted the Internet for financial and administrative transactions but instead continue to seek payment through paper-based claims or electronic data sent over direct connections via modems.

The reasons for the limited adoption of the Internet in health-related activities are manifold, but the underlying reason is a lack of demonstrated value in different applications. The Internet has been widely adopted by the public as a tool for gaining insight into issues of illness and health because it is perceived to deliver value. Many (but not all) care providers use the Web frequently for searching online databases (such as MEDLINE), also because it is perceived to deliver value. A small, but growing, number of care providers engage in e-mail discussions with their patients about health problems. Care providers do not use the Internet more broadly in the process of treating patients because the valuable, usable, affordable, and practical Internet-based solution has yet to be built. The process of determining which applications add value in health applications—and which specific capabilities and attributes provide that value—requires continued experimentation and analysis of data on the benefits and costs of the Internet relative to those of other media.

To date, little information is available with which to gauge the contributions of the Internet to the provision of health care—not to mention its potential to improve public health, biomedical research, and professional education. Emerging evidence of the benefits to health care of information systems generally bodes well for the Internet; a growing number of studies demonstrate, for example, reductions in adverse drug interactions and improved diagnoses stemming from the use of computer-based decision support tools in clinical environments.¹¹ Research has also demonstrated the positive effect of information technology applications in several other areas of health care.¹² However, the ability of the Internet (as opposed to private networks) to improve the quality of health care or expand access to it has not been demonstrated. On the contrary, there has been considerable concern about the quality of health information available on the Internet and its potential to harm consumers (Mittman and Cain, 1999; SCIPICH, 1999). In an industry already facing serious fiscal and organizational upheaval, health care organizations may remain skeptical of a range of Internet applications until there is greater evidence of their benefits, along with more information about the policies and procedures needed to avoid the potential harms.

The benefits of the Internet in health applications may prove difficult to measure because the most notable benefits may be indirect and may vary across segments of the health sector. For example, the advantages of

consumer-oriented Web sites to care provider organizations such as hospitals might include marketing, possibly advertising, and the collection of valuable data about interested consumers. The direct and indirect revenues from many of these activities, however, may be insufficient to support the development and maintenance of the applications themselves. Furthermore, the use of the Internet could stimulate changes in industry structure that are difficult to foresee at present. For example, the Internet could enable large provider organizations to extend their reach more directly into local communities, working with local care providers to provide greater continuity and consistency of care. Or, it could allow consumers to better triage their own health needs using online modules created by their health plans. These changes could improve health and disease management among local populations, but the benefits may accrue most directly to consumers. The benefits to care providers may be more difficult to measure, especially if healthy patients demand fewer health services in the long run.

Further slowing adoption of the Internet by health organizations are uncertainties about the technical capabilities needed to support health applications. Managers of many health organizations say that security concerns prevent them from using the Internet to transfer patient medical records among affiliated organizations or from allowing care providers to access such records remotely (Siwicki, 1999). At the same time, they are not certain what types of security technology are needed to adequately protect patient information in such applications. With respect to other applications, such as remote medical consultations, practitioners note that they cannot obtain sustained access to the bandwidth they might need for real-time video. It is appropriate to question whether today's Internet provides a sufficiently strong infrastructure to support applications such as critical-care monitoring and automated delivery of medication. If deployed in health care settings without proper attention to these capabilities, the Internet could have an adverse effect by eroding patient privacy and preventing care providers from accessing needed information.

TECHNICAL CONSIDERATIONS

Whether the Internet will become more widely adopted in health care will depend, in part, on the technical capabilities it can provide and how these capabilities compare with those provided by other networking alternatives available to health organizations and consumers. A number of technical factors need to be considered in such evaluations. The five primary factors considered in this report are bandwidth, latency, availability, security, and ubiquity.

1. *Bandwidth* is the rate at which information is transmitted through a network, measured in bits (or kilobits or megabits) per second. The bandwidth a network can provide is a property of the transmission medium (e.g., fiber optics, coaxial cable, telephone wire, radio waves), the network topology, and the switching or routing devices used to guide traffic through the network. The amount of bandwidth a particular application demands is determined by the amount of data to be transmitted and the time in which that transmission must be completed. Applications that must transfer large amounts of data quickly demand much greater bandwidth than do applications that transfer smaller amounts of data (such as e-mail) or transfer data more slowly (e.g., if a response is not needed quickly). From the point of view of an individual user (e.g., a consumer, a doctor, or a nurse), the demand for bandwidth is a demand not so much for a uniform increase in the bandwidth of the entire network but for access to sufficient bandwidth when needed.

2. *Latency* is the time required to transmit data across the network (i.e., the delay between a sender transmitting a message and a recipient receiving it). The minimum latency a network can provide is influenced by the speed of its switches and routers and the physical distance across which the message is sent. Data communications traverse different media at the speed of light, which places a lower limit on the time it takes for a message to travel between two points on the network. The latency an application demands can vary tremendously. Real-time, interactive applications demand low latency so that users can interact with each other easily. Many interactions, such as telephone conversations or control of remote devices, become unwieldy if round-trip latencies (i.e., across the network and back) exceed 300 milliseconds. Applications that do not demand real-time interactions between users—so-called asynchronous applications such as e-mail and store-and-forward messaging systems—have only weak demands on latency requirements. Closely related to latency is jitter, the variation in latency over time. High levels of jitter imply unpredictable degrees of latency across the network, although several techniques, including temporary buffering of information, can remove jitter at the receiver's end of the network. In some applications, the related notion of response time is more important than latency per se. Response time refers to the length of time needed to transmit a full message (such as a service request or an image) rather than an individual packet across the network and receive a response. Messages can consist of many packets, and successful transmission can depend far more on network reliability on a packet-by-packet basis than on the actual latency. The Transmission Control Protocol (TCP) used on the Internet, for example, may transmit more than one packet before it receives an acknowledgement that the first packet has been received, but if a packet is not received,

TCP has to wait for a certain period of time before transmitting that packet again. Thus, both response time and latency can be affected by the load on a network at a given point in time.

3. *Availability* refers to the continuous availability of the network, the individual links of which it is composed, and the services it offers. Availability can be measured in terms of the percentage of the time the network (or a particular link) is operational or by the average time between failures. A number of factors can render networks unavailable, including physical damage to network links or nodes, hardware or software failures (i.e., component reliability problems), operator error, software errors, and deliberate malicious attacks against the system. Steps can be taken to harden systems against these sorts of failures, and procedures can be developed for restoring some level of network services in the event of failures. Nevertheless, some applications cannot function properly if the performance of the network is degraded; and many time-critical applications cannot tolerate network failures, even if very brief.

4. *Security*, in the computer science community, generally encompasses three elements: system availability, confidentiality, and integrity. The first of these is addressed above. Confidentiality refers to the ability to prevent communications from being disclosed to unauthorized parties in violation of disclosure rules. Integrity refers to the ability to prevent malicious or accidental alteration of data. These two capabilities can be provided by a variety of technical mechanisms that support authentication of user identities, encryption of communications, and different forms of access control. The need for confidentiality and integrity varies greatly across applications. Both are important concerns in applications involving exchanges of personal health information. Integrity is of paramount importance for some applications, such as setting levels on dosimetry equipment that delivers drugs to patients and preserving the authenticity of medical images.

5. *Ubiquity* refers to the relative accessibility of a network. The ubiquity of a network is influenced by the network's geographic scope (i.e., whether it can be accessed from many places) and by rules regulating participation (i.e., whether it is open to the general public or to members only). The telephone network, for example, is highly accessible because roughly 94 percent of U.S. homes have telephone connections, and anyone is allowed to subscribe to the service. Cable television is almost as ubiquitous because it passes most homes and is also available to all who pay for service. However, cable modem service (to support data communications as opposed to television) is not yet available to all locations in the United States. Such systems stand in contrast to private networks, such as those used by financial institutions, which may have broad geographic reach but strict rules regarding membership. Applications that

serve the general public usually demand high levels of accessibility (ubiquity), whereas those that serve limited populations do not, although there may be interest in allowing access by members from multiple locations.

Another term that is important in describing network performance is “quality of service” (QOS). Network engineers use this term to refer to the ability of a network to provide a range of assured levels of performance. Performance is characterized by metrics such as the bandwidth obtained between two points in the network (which may be dramatically less than the bandwidths of the individual communications links involved, either because of other traffic on the network or because of the need to retransmit packets of information dropped during transmission); latency and jitter (defined above); and the packet loss rate, the percentage of transmitted packets that are dropped inside the network and not delivered to their intended destinations.¹³

The need for QOS stems from a design characteristic of networks whereby resources (e.g., especially backbone links) are generally shared among many users who are running applications simultaneously. Thus, even if the bandwidth of the system is measured in megabits or gigabits per second, any single user might gain access to only a small fraction of that bandwidth. How much bandwidth a single user obtains depends on the level of activity of the other users. When using the Internet, for example, a user often obtains high bandwidth and low latency over certain paths and at certain times of day, but it is not currently possible to ensure that such conditions will be available on a given path at a given time. Hence, although some users may succeed in making telephone calls of acceptable quality over the Internet, they cannot depend on this medium to make calls to any location at any time. Quality of service guarantees require mechanisms that enable applications with specific requirements to negotiate to receive appropriate treatment in the network. Those mechanisms must be able to deal with requests from huge numbers of applications running simultaneously.¹⁴

NETWORKING ALTERNATIVES

The Internet differs from other communications networks in all the dimensions outlined above. Many of these differences stem from the public character of the Internet (Box 1.3), which carries aggregated traffic from numerous parties, whereas private networks interconnect a limited number of sites using dedicated transmission links that are not shared with any other users. Two private networks are completely isolated from each other in the sense that no data can “leak” from one into the other, a user of one network cannot access resources on the other, and the level of

BOX 1.3 **Defining the Internet**

The Internet is commonly defined as a network of networks. Although it is often considered a single, uniform network, the Internet is actually composed of a large (and growing) set of independent networks that are connected to each other using a set of common standards and protocols, such as the Transmission Control Protocol (TCP) and the Internet Protocol (IP). The Internet is based on a technology called packet switching, which breaks down messages (electronic mail, data files, digital images, etc.) into smaller pieces called packets and sends them across the network in a series of hops from one routing device (or router) to another.¹ Routers determine dynamically which route to send each packet along, with the end result that the packets constituting a single message may be sent along different routes, depending on relative levels of congestion and other factors. The packets are reconstituted at the recipient's computer and assembled in the proper order.

The underlying networks that make up the Internet can assume many forms, ranging from access networks to backbone networks. Access networks connect customers (whether individuals or entire organizations with their own internal local area networks) to an Internet service provider (ISP). These access networks connect to backbone networks with wider geographic scope, providing regional, national, and international connectivity. Thousands of ISPs offer connectivity in the United States, but the major Internet backbone networks are operated by a small number of independent companies that interconnect with each other at numerous points, some of which are private (typically a connection between only two providers) and others of which are public. The public interconnection points are used not only by backbone providers but also by local providers to gain access to the backbone. Backbone networks connect to access networks and to each other at points of presence, where service providers house switching hardware and transmission equipment.

This topology has two important consequences. One is that a packet traversing the Internet will normally cross networks owned and operated by many different providers en route to its eventual destination, with implications for overall performance and security. The second is that backbone networks carry highly aggregated traffic, representing the combined traffic flows of millions of customers connected to access networks. Although individual users have some control over the capacity of access lines and the flows of information across those lines, they have virtually no control over data flows across the rest of the network—meaning they have little assurance of the speed at which their messages will transit the net. They have a similar lack of control over the other organizations and individual users that may join the network. The Internet is a public network that interconnects other networks running IP and any user with an Internet account. The public nature of the Internet therefore introduces a number of security risks not seen on private networks (see Chapter 3).

¹Packet switching differs from circuit switching used in the telephone system in that it does not establish a dedicated circuit between sender and recipient for the duration of a transmission. Rather, packets pertaining to many different messages may be routed over the same links to use communications channels most efficiently.

usage on one network has no effect on the availability of resources in the other.¹⁵ Thus, users have greater control over QOS (i.e., the bandwidth that will be available at a given time) and confront fewer security risks than do users of the Internet. For these reasons, the Defense and Energy Departments, the banking industry, and other sectors have developed communications networks that are separate from the Internet.

And yet, the distinction between public networks such as the Internet and private networks is blurred by a number of factors. For example, many private networks used for internal communications among elements of a single corporation can be connected to the Internet in a controlled way that limits the messaging traffic that can transit the interface (indeed, many of the networks attached to the Internet are private). Similarly, many organizations develop private intranets that are based on Internet technologies and protocols to facilitate data sharing within an enterprise. Other networks, called extranets, use Internet standards and technologies to support secure exchanges of information among trading partners. Many of these private networks—whether intranets or extranets—could be made accessible to valid users over the public Internet through the use of appropriate security and authentication technologies. The distinction between private and public networks is further blurred by technologies such as virtual private networks (VPNs), which provide many of the same properties of private networks while using shared network facilities (including possibly the Internet itself) to provide connectivity. In this way, VPNs (discussed in greater detail in Chapter 3) create the illusion of private, dedicated point-to-point connections, but many VPNs may be supported on the same physical network. They provide almost the same level of security as true private networks because there is no connectivity among the VPNs, but messages on one virtual circuit must contend with those on other circuits for network resources (e.g., bandwidth).

Whether a network is truly or virtually private, it differs in a very important way from the Internet. In a private network, connectivity is deliberately constrained to a limited number of sites, and these sites are known a priori. Every time a new site is added to the network, some administrative overhead is involved, whether ordering and awaiting the installation of a new circuit or provisioning a new virtual circuit. Thus, private networks are ideally suited to an environment in which the necessary connectivity is known and remains stable over time. A classic application is interconnection of the different geographic locations of a single corporation. Another application is the connection of a number of companies that have a long-standing business relationship, such as a large manufacturer and its parts suppliers. However, private networks are fundamentally unsuited to environments in which arbitrary connectivity

is required. The Internet, by contrast, enables users to connect to each other without a prior arrangement. In a business setting, this enables consumers to find suppliers readily, and vice versa.

Because of the range of technical capabilities available across different types of networks, organizations can tailor their network architectures to their specific needs. When security and QOS are important and the relationships between communicating parties are sufficiently well known in advance, private networks are likely to be chosen. Accordingly, health organizations continue to use dedicated networks to transmit sensitive patient information, share large image files, and submit claims for reimbursement. When the overriding goal is maximum connectivity without a priori knowledge of the communicating parties, as in the case of making a product or service available to a wide set of consumers, the Internet is likely to be chosen. In some cases, the most judicious choice might be a network that uses a combination of private or dedicated lines that are connected in appropriate ways to the Internet to allow broader access, but with suitable security capabilities in place. What may be most important is the use of consistent, interoperable protocols for all communications so that various networks can be connected as needed with appropriate gateways. The architectures chosen by various organizations depend on the requirements and cost-benefit analyses of technical alternatives.

The value of open (i.e., public) networks in health care is rooted in the nature of the industry, which remains highly decentralized and involves a range of individuals and organizations in providing care, paying bills, analyzing health data, conducting health services research, and monitoring public health. As recently as 1995, the United States had 1.2 million health care providers—half of whom work in private practices—and more than 3,000 private insurance payers.¹⁶ The patterns of data sharing among these organizations are complex and can change frequently. Internet use could enable IDNs, for example, not only to exchange patient records among affiliated hospitals and clinics, but also to send records to other hospitals to assist in the treatment of patients who are injured or become ill while traveling. It could enable any rural care provider to arrange remote medical consultations with any remote specialist connected to the Internet who has the expertise needed to handle a particular patient's case. It could lead to much more rapid reporting of diseases, enabling state public health officials to accept reports directly from physicians and testing laboratories throughout the state.

Furthermore, Internet connections promise to be less expensive to install and maintain than private networks. Individual lines do not need to be leased from telephone companies to connect the various partners in the network. Internet connections entail some costs, especially if high-bandwidth connections are needed, but they tend to be lower than the

costs of leased lines and can be spread among a wider range of applications and users. Finally, the Internet has been the catalyst for the integration of many applications and systems of applications. The existence of a communication infrastructure built on a common set of protocols makes it more difficult to justify, both economically and technically, the use of separate, special-purpose networks for different applications; the common infrastructure also facilitates interactions between separately developed systems. When separate information systems are integrated on the foundation provided by the Internet, the investments frequently are justified on the basis of both savings in communication costs and improved functions. With improved security and QOS, the Internet might become preferable to private networks in almost all cases.

The Internet may prove to be an ideal technology for use by willing health care organizations to simplify and standardize processes and collaborate more effectively with one another. The value of this common communications infrastructure in other sectors, from entertainment to banking to retail sales, could extend to health care if technical obstacles, organizational uncertainties, and policy barriers can be overcome. Organizations that adopt the Internet could reap significant benefits, including cost savings from support of a wide range of applications and the ability to leverage the technology investments of other communities. An analogy to the telephone system may be illuminating. There is no separate U.S. telephone system for health applications. Despite the distinct set of priorities and trade-offs associated with health applications, the health community uses "plain old telephone service" (POTS) for most of its voice communications, leveraging the telephone companies' investments in R&D as well as infrastructure deployment and accepting inconveniences such as busy signals that might not be desirable in emergencies. Features such as 911 have been added to POTS to support emergency needs (whether related to health or public safety), but even this feature leverages the existing network. Some specialized voice communications networks, such as communications between emergency rescue vehicles and hospital emergency rooms, have also been established for health needs not well served by POTS.

ENHANCING THE INTERNET

A number of efforts are under way to improve the ability of the Internet to provide QOS, security, and availability, which would enable its broader use within the health domain. These efforts include attempts by individual companies to increase the capabilities of Internet routers and deploy high-speed data services on demand, as well as attempts by the Internet Engineering Task Force (IETF) to develop new standards and

protocols for improved services (see Chapter 3). In addition, two major collaborative efforts are under way to develop and demonstrate advanced networking technologies that promise to improve the QOS, availability, and security across the Internet. The government's Next Generation Internet initiative and projects sponsored by the private-sector University Consortium for Advanced Internet Development (UCAID) are attempting to develop advanced networking technologies and applications and deploy them in testbed networks that link a limited number of sites and allow early experimentation with advanced applications. The technologies and applications to be developed under these programs, which are described below, could diffuse onto the Internet as they are demonstrated and proven.

The Next Generation Internet Initiative

Formally initiated in October 1997, the NGI initiative is a multiyear program, funded at approximately \$100 million per year, that involves a number of federal agencies: the Defense Advanced Research Projects Agency (DARPA), the National Science Foundation (NSF), the Department of Energy (DOE), the National Aeronautics and Space Administration (NASA), the National Institute of Standards and Technology (NIST), and the National Institutes of Health (NIH) (Table 1.1). The initiative is managed by individual agencies, with coordination provided by the large-scale networking working group of the White House National Science and Technology Council's Committee on Technology, Subcommittee on Computing, Information, and Communications R&D.¹⁷

The NGI initiative has three components: R&D on advanced networking technologies for improved performance and functionality; the deployment of high-speed testbed networks that emphasize end-to-end performance; and the development and demonstration of revolutionary applications that demand advanced networking and are not possible on today's Internet.¹⁸ The first component involves R&D projects in areas such as high-speed routing, security, QOS, and network management and modeling. This work will be funded primarily by DARPA but also by NSF, NASA, and NIST.

The second component will be achieved by developing and demonstrating applications of two types: (1) discipline-specific applications of interest to participating agencies, including health care, basic science, education, and environment and (2) their enabling technologies, including collaboration technologies, digital libraries, distributed computing, privacy and security, and remote operation and simulation (National Science and Technology Council, 1999). The NIH is actively involved in this effort through the NLM, which awarded 24 contracts totaling \$2.3 million

TABLE 1.1 Agency Funding for the Next Generation Internet Initiative, 1998-2000 (in millions of dollars)

Agency ^d	1998				1999 Total	2000 Total
	Research	Testbeds	Applications	Total		
DARPA	20	20	2	42	50	40
NSF	5	10	8	23	25	25
DOE	0	0	0	0	15	0
NASA	2	3	5	10	10	8
NIST	0	0	5	5	5	5
NIH/NLM	0	0	5	5	5	5
Total	27	33	25	85	110	83

NOTE: Breakouts for 1999 and 2000 funding into categories of research, testbeds, and applications are not available.

^dDARPA, Defense Advanced Research Projects Agency; NSF, National Science Foundation; DOE, Department of Energy; NASA, National Aeronautics and Space Administration; NIST, National Institute of Standards and Technology; NIH, National Institutes of Health; NLM, National Library of Medicine.

SOURCE: Grant Miller, National Coordination Office for Computing, Information, and Communications, 1999, personal communication.

in October 1998 to investigate and develop health care applications of the NGI. These projects make up the first phase of a three-phase program. Several of the projects received phase II awards in late 1999 and early 2000 to allow their implementation in local testbed settings (see Appendix B for a list of all NLM project awards as of January 2000). Phase III will support scale-up to the regional or national level of successful phase II testbed projects. These projects are intended to improve the health community's understanding of the ways in which the NGI can affect health care, health education, and health research systems with respect to cost, quality, usability, efficacy, and security. Supported projects include efforts to (1) build a virtual human cadaver for educational purposes, (2) develop telemedicine technologies to support health care in rural areas,¹⁹ (3) demonstrate the feasibility of a national breast imaging archive and networking infrastructure to support telemammography, and (4) create a personal health record that can be integrated with more traditional sources of clinical information for patient use in the home, at work, or at school (see Box 1.4 for examples of these projects and Appendix A for a complete listing of NLM project awards).

Other federal agencies, including NASA and the NSF, have also funded projects that will demonstrate health-related applications of the

BOX 1.4 Examples of Projects Funded by the National Library of Medicine

- *Networked Three-Dimensional Virtual Human Anatomy.* The University of Colorado Health Sciences Center plans to build a virtual human cadaver based on the National Library of Medicine's Visible Human data set, which provides detailed information on human anatomical structures. An online virtual cadaver would be available over the Internet to a wide range of students, who could explore the virtual cadaver with a variety of tools. High-end applications will have a haptic (tactile) interface.

- *NGI-Aware, Scalable, Secure, and Adaptive Technology for Rural Telemedicine.* West Virginia University Research Corporation will develop a plan to demonstrate telemedicine applications that will use the Next Generation Internet (NGI) infrastructure. Telemedicine scenarios include nomadic clinics, public health stations, and a consulting health station in rural clinics and hospitals. These systems will be configured with a set of videoconferencing, diagnostic, and patient monitoring equipment.

- *Telemammography Using the NGI.* The University of Pennsylvania will plan and implement a test bed to demonstrate the feasibility of a national breast imaging archive and network infrastructure to support telemammography using NGI technologies. The proposed infrastructure would support traditional breast screening; provide the opportunity to maintain and apply standard image processing and computer-aided diagnosis software; permit access to breast imaging experts for primary and secondary interpretations; and provide an opportunity to study and understand epidemiological issues in breast cancer.

- *Personal Internetnetworked Notary and Guardian* (Children's Hospital, Boston). The Personal Internetnetworked Notary and Guardian (PING) project is designed to address the control of a personal record that can be integrated with more traditional sources of clinical information for patient use in the home, at work, and at school. In particular, PING is focused on (1) the reconstitution via the Internet of patient longitudinal records from both provider-based information systems and portable, personal record systems, (2) providing simple and secure authentication mechanisms, and (3) evaluation of the impact of PING on the health care process.

SOURCE: Derived from information provided on the National Library of Medicine's Web page, available online at <www.nlm.nih.gov>.

NGI. Researchers at NASA's Ames Research Center, for example, are developing a system for sharing high-resolution, three-dimensional medical images in real time for purposes of collaborative diagnosis and surgical planning.²⁰ NSF is supporting work to provide psychological services over a distance to deaf patients, to develop digital video resources for teaching and learning the life sciences (using materials that reside at

the NLM), and to allow Web-based control of a remote electron microscope for biological research, among other projects. Such efforts reflect the importance of health-related applications in motivating large-scale information infrastructure programs such as the NGI.

The third component of the NGI initiative will be carried out by constructing two types of testbed networks, one of which will link approximately 130 participating universities and federal agencies at speeds 100 times faster than those available across the Internet in 1997²¹ and the other of which will link about 10 sites at speeds 1,000 times faster than the 1997 Internet. The first testbed will be built on several existing federal networks: the NSF's very-high-performance Backbone Network Service (vBNS),²² NASA's Research and Education Network, DOD's Defense Research and Education Network, and DOE's Energy Sciences network. The vBNS, for example, operated at 622 megabits per second (Mbps) in 1998 but is expected to be upgraded to 2.4 gigabits per second (Gbps) by the year 2000. Universities connecting to the vBNS at 45 Mbps will be upgraded to 155 Mbps to help them take greater advantage of the increased backbone capacity.

The NGI initiative's other testbed will be built on DARPA's SUPERNET, a network composed of a variety of high-speed technologies and testbeds, enabling researchers to collaborate and experiment with advanced networking technologies and applications in a diverse, high-capacity, wide-area environment. It will use wave-division multiplexing technology (WDM) to allow multiple frequencies of light (and hence multiple communications channels) to share a single fiber-optic cable (see Chapter 3). DARPA demonstrated a 5-node network at 2.5 Gbps per channel in 1999 and plans to establish a 10-node network with 160 Gbps facilities in 2002. The NSF, NASA, and DOD networks will connect to this network.

The goal of these networks is to provide a cutting-edge but stable network that will support the development of revolutionary applications and serve as a testbed for new technologies and protocols. According to the 1998 NGI implementation plan, the testbed networks will be initially deployed with best-effort services using IP version 4 (IPv4) (Large Scale Networking Next Generation Implementation Team, 1998). New versions of IP (including IPv6), QOS technologies, multicast protocols (for facilitating group interactions), security protocols, and network management tools will be deployed in the networks as soon as they become stable. Feedback from application developers to network researchers, operators, and implementers will help ensure that the testbeds evolve in a manner suitable to the types of applications that are expected to be run on them.

Private-Sector Efforts: Internet 2 and Abilene

The UCAID, which was incorporated in 1998, has two related networking projects under way that promise to enhance the capabilities of the Internet. The first is the Internet 2 project, which will link more than 100 member universities and partners to an advanced academic network. Research supported by Internet 2 is attempting to enable applications that are not possible with the technology underlying today's Internet (some examples are telemedicine, digital libraries, and virtual laboratories). The program is intended to demonstrate new applications for improving research and enhancing the delivery of education and other services, including health care. It will facilitate the development, deployment, and operation of an affordable communications infrastructure capable of supporting differentiated QOS based on the applications requirements of the research and education community, and it will promote experimentation with the next generation of communications technologies.²³

Biomedical applications play a significant role in the Internet 2 initiative. The first demonstration of the network, in October 1999, consisted of an online broadcast of a gall bladder operation. The surgery team inserted light, camera lenses, and surgical tools inside the patient's body, creating internal views of the operation. Audio and video were transmitted over the network in real time, requiring network bandwidth that would support a consistent data transmission rate of 2 Mbps. Only a small audience was able to view the demonstration, but it enabled a doctor based in Washington, D.C., to assist in the surgery, which took place at Ohio State University.²⁴

A related UCAID project, Abilene, is seen as a second Internet 2 backbone. Abilene is based on a partnership with Qwest, Cisco Systems, Nortel, and Indiana University. The goals are to provide a high-availability backbone network to support the demands of the advanced research applications being developed by UCAID members; a separate network to enable the testing of advanced network capabilities (for example, QOS, multicasting, and security and authentication protocols) prior to their introduction into the application development network; and a separate network capability to conduct networking research, including the design of an alternative network capable of advancing both the Abilene network and the general state of the art.²⁵ Internet 2 member universities have committed more than \$70 million per year in new investment on their own campuses for the Internet 2 project, and corporate members have committed more than \$30 million over the life of the project.

Although programmatically distinct from the NGI initiative, the UCAID's efforts are related to federal networking activities. More than 90 Internet 2 universities have received grants under NSF's High Perfor-

mance Connections program to support links to advanced backbone networks such as Abilene and the vBNS. Internet 2 is also participating in the NGI Joint Engineering Task Force to ensure the cohesiveness and interoperability of the technologies that Internet 2 is developing. Internet 2 member institutions may receive funding in the form of competitively awarded grants from the NSF and other federal agencies participating in the federal NGI initiative. Additional cooperative relationships are being planned as part of NGI implementation.

Deploying Enhanced Internet Technologies

Although they are structured as programs with a limited number of participants, the NGI and Internet 2 initiatives are intended to serve as launching points for enhancement of the public Internet. Both programs have a stated interest in transferring new technical capabilities to the public Internet once the technologies are developed and demonstrated to be robust. Just as early DARPA support for the ARPANET and subsequent NSF support for NSFNET laid the groundwork for today's Internet by funding networking research and applications development and deploying network infrastructure,²⁶ so too, it is hoped, will the NGI and Internet 2 initiatives plant the seeds for an improved Internet that can serve the public at large. They intend to accomplish this by developing and demonstrating technologies that can later be deployed in networks maintained and operated by private companies.

Whether the public Internet will evolve into a network capable of supporting a full range of health applications will depend on many factors other than technology. Of particular importance will be economic incentives for network providers to deploy the levels of bandwidth, QOS, security, availability, and ubiquity that health applications demand. These incentives will be derived from the combined demands of many applications in different sectors, including health. The history of Internet development is one of innovation and experimentation, not planned development. The forces that drive its continuing evolution are increasingly economic, and these forces alone may not yield an infrastructure that can support the integration of critical and noncritical functions of the health community. In the end, some capabilities may prove too expensive to deploy throughout the Internet, leaving health organizations to operate with a mixture of different networking infrastructures to meet their various needs.

Only by making its needs explicit and working with organizations involved in the deployment of Internet capabilities can the health community hope to ensure that an enhanced Internet infrastructure meets health needs. This report represents the first step in that effort. By

evaluating the technical capabilities that the Internet must provide to support different health applications, the report offers the health community information that it can use to shape the networks being deployed as part of the NGI and Internet 2 initiatives and, ultimately, as part of the Internet. Clearly, ongoing evaluation and experimentation will be needed. The many uncertainties inherent in the process of developing and deploying Internet-based applications make any attempt to predict the long-term evolution of the Internet within the health community foolhardy. Sustained interaction will be needed to ensure that the emerging needs of the health community continue to be met by the evolving capabilities of the Internet.

ORGANIZATION OF THIS REPORT

The remainder of this report outlines the technical and nontechnical challenges that must be overcome if the Internet is to support a widening range of health applications. Chapter 2 examines specific applications of the Internet across this domain. The first part of the chapter focuses on applications of the Internet in the provision of health care, addressing topics such as consumer health, remote consultation, and the transfer of medical images for diagnostic purposes. The next parts of the chapter explore Internet applications in areas such as public health, health care finance and administration, and biomedical research. The chapter draws on a series of site visits by the committee that provided insight into the types of Internet applications being developed today and the networking challenges that cannot currently be ported to the Internet. The chapter reviews the technical capabilities that each application demands in terms of QOS (combining bandwidth and latency requirements), security, availability, and ubiquity. The applications examined are intended to illustrate the range of ways in which the Internet might be used rather than to identify them as likely paths.

Chapter 3 reviews the technical challenges posed by applications of the Internet in health, health care, and biomedical research. It examines ongoing efforts to enhance the capabilities of the Internet and identifies areas in which health care needs might not be addressed if they are not explicitly considered during the research process. Chapter 4 examines organizational barriers to the deployment of the Internet for health and health care. It describes ways in which the Internet can serve the strategic interests of health care organizations and identifies the range of uncertainties surrounding the Internet's use that hamper efforts to deploy it more broadly in such organizations. Chapter 5 discusses elements of public policy that stand in the way of greater use of the Internet in the health community. These barriers range from issues of payment for services and

licensure that have stymied previous attempts at telemedicine, to broad issues of intellectual property protection and privacy that have special significance in the health domain.

Finally, Chapter 6 summarizes the committee's conclusions and offers a series of recommendations for facilitating the more widespread use of Internet technologies in health care and biomedical research. The recommendations suggest ways in which technical and nontechnical barriers can be overcome to enable the design of an Internet that will more fully support the needs of the health sector.

REFERENCES

- AT&T. 1999. "AT&T and M.D. On-Line, Inc. to Promote AT&T WorldNet Internet Connectivity to Healthcare Providers," News release. February 8. Available online at <<http://www.att.com/press/item/0,1193,339,00.html>>.
- Barry, M.J., J.J. Fowler, Jr., A.G. Mulley, Jr., J.V. Henderson, Jr., and J.E. Wennberg. 1995. "Patient Reactions to a Program Designed to Facilitate Patient Participation in Treatment Decisions for Benign Prostatic Hyperplasia," *Medical Care* 33:771-782.
- Benton Foundation. 1999. *Networking for Better Care: Health Care in the Information Age*. Benton Foundation, Washington, D.C., March.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1999. *Funding a Revolution: Government Support for Computing Research*. National Academy Press, Washington, D.C.
- Davis, Robert, and Leslie Miller. 1999. "Millions Comb the Web for Medical Info," *USA Today*, July 15. Available online at <www.usatoday.com/life/health/online/lhon1012.htm>.
- Gillespie, Greg. 2000. "Online Clinical Guidelines Help Trim Costs," *Health Data Management* 8(1):39-45.
- Greenfield, S., S. Kaplan, and J. Ware, Jr. 1985. "Expanding Patient Involvement in Care," *Annals of Internal Medicine* 102:520-528.
- Gustafson, David H., R.P. Hawkins, E.W. Boberg, and E. Bricker. 1992. "CHESS: A Computer-Based System for Providing Information, Referrals, Decision Support, and Social Support to People Facing Medical and Other Health-Related Crises," pp. 161-165 in *Proceedings of the Annual Symposium on Computer Applications in Medical Care*. McGraw-Hill, Health Professional Division, New York.
- Gustafson, David H., M. Wise, F. McTavish, J.O. Taylor, W. Wolberg, and J. Stewart. 1993. "Development and Pilot Evaluation of a Computer-Based Support System for Women with Breast Cancer," *Journal of Psychosocial Oncology* 11:69-93.
- Gustafson, David H., R.P. Hawkins, E.W. Boberg, and E. Bricker. 1994. *The Impact of Computer Support on HIV Infected Individuals*. Final Report to the Agency for Health Care Policy and Research, Washington, D.C.
- Halamka, J., and M. Hughes. 1998. "A Paradigm Shift in Health Care Information Systems: Clinical Infrastructures for the 21st Century," pp. 401-405 in *Proceedings of the American Medical Informatics Association Fall Symposium*, C. Chute, ed. Hanley & Belfus, Philadelphia.
- Information Infrastructure Task Force (IITF), Committee on Applications and Technology. 1994. *Putting the Information Infrastructure to Work*. National Institute of Standards and Technology, Gaithersburg, Md.

- Institute of Medicine. 1996. *Telemedicine: A Guide to Assessing Telecommunications in Health Care*, Marilyn J. Field, ed. National Academy Press, Washington, D.C.
- Large Scale Networking Next Generation Implementation Team. 1998. *Next Generation Internet Implementation Plan*, Second Printing. National Coordination Office for Computing, Information, and Communications, Arlington, Va., February.
- Lindberg, Donald A.B., and Betsy L. Humphreys. 1998. "Medicine and Health on the Internet: The Good, the Bad, and the Ugly," *Journal of the American Medical Association* 280(15):1303-1304.
- Miller, Thomas E., and Scott Reents. 1998. *The Health Care Industry in Transition: The Online Mandate to Change*. Internet Strategies Group, Cyber Dialog, Inc., New York.
- Mittman, Robert, and Mary Cain. 1999. *The Future of the Internet in Health Care: Five Year Forecast*. Institute for the Future, Menlo Park, Calif., January.
- Morgan, M.W., R.B. Deber, H.A. Llewellyn-Thomas, P. Gladstone, R.J. Cusimano, and K. O'Rourke. 1997. "A Randomized Trial of the Ischemic Heart Disease Shared Decision-Making Program: An Evaluation of a Decision Aid," *Journal of General Internal Medicine* 12(1):62.
- Morrison, J. Ian. 1999. "Healthcare in the New Millenium: The Promise of the Internet," Presentation at Internet Health Day II: Health Care in Transition-Preparing for an Interactive Future. New York, October 12.
- Nash, Sharon. 1999. "The Doctor Is Online," *PC Magazine Online*, July 14. Available online at <www.zdnet.com>.
- National Science and Technology Council, Committee on Technology, Subcommittee on Computing, Information, and Communications R&D. 1999. *Information Technology Frontiers for a New Millennium*. Supplement to the President's FY 2000 Budget, National Coordination Office for Computing, Information, and Communications, Arlington, Va., April.
- National Telecommunications and Information Administration (NTIA). 1999. *Falling Through the Net: Defining the Digital Divide*. U.S. Department of Commerce, Washington, D.C.
- Rosenberg, Matt. 1999. "Popularity of Internet Won't Peak for Years," *Puget Sound Business Journal*, May 24. Available online at <<http://www.amcity.com/seattle/stories/1999/05/24/focus.html>>.
- Science Panel on Interactive Communication and Health (SCIPICH). 1999. *Wired for Health and Well-Being: The Emergence of Interactive Health Communication*, Thomas R. Eng and David H. Gustafson, eds. Office of Disease Prevention and Health Promotion, U.S. Department of Health and Human Services, Washington, D.C., April. Available online at <<http://www.scipich.org>>.
- Siwicki, Bill. 1999. "Applying the Internet in Health Care," *Health Data Management* 6(3):38-48.
- Smith, K.A., and R.B. Mehnert. 1986. "The National Library of Medicine: From MEDLARS to the Sesquicentennial and Beyond," *Bulletin of the Medical Libraries Association* 74(4):325-32.
- U.S. Department of Commerce, Economic Statistics Administration. 1999. *The Emerging Digital Economy II*. Washington, D.C., June.
- Vickery, D.M., T.J. Golaszewski, E.C. Wright, and H. Kalmer. 1988. "The Effect of Self-Care Interventions on the Use of Medical Service Within a Medicare Population," *Medical Care* 26:580-588.

NOTES

1. For an illustration of the role of health applications in motivating federal programs to develop national information infrastructure, see IITF (1994).

2. The Internet also has applications in support of clinical research (e.g., clinical trials), but these applications are not investigated in great detail in the report.

3. Others have also noted the importance of organizational and policy issues in influencing the rate of adoption of Internet technologies in health applications. For example, see Lindberg and Humphreys (1998).

4. For example, systems to transfer large medical image files between sites can be designed in different ways. Some systems demand high network bandwidth because there is little preprocessing of images and little attention to representative workflows; others rely more on preprocessing, which reduces network bandwidth requirements.

5. For a more detailed history of MEDLARS and MEDLINE, see Smith and Mehnert (1986).

6. These data are from the Bureau of Economic Analysis as presented in U.S. Department of Commerce (1999), the Appendix to Chapter III.

7. For an example of a Web site enabling communications with insurers, see AT&T (1999).

8. Managed care plans integrate insurance and delivery of care—functions otherwise provided by separate entities. Most managed care plans now pay care providers some form of discounted fee for services rendered, although some still pay a fixed fee based on the number of patients enrolled in their care.

9. Integrated delivery systems combine entities related to the provision of health care and may have relationships with health insurance plans. Such organizations typically include a range of different facilities, from major hospitals to local clinics, so they can provide a continuum of care.

10. For an example of Internet-based quality indicators and managed care data exchange, see Halamka and Hughes, 1998.

11. For example, researchers at Intermountain Health Care in Salt Lake City, Utah, have developed a system that provides clinical guidelines in real time to physicians who use the electronic medical record system. One study indicated that use of the system improved from 30 to 70 percent the percentage of diabetic patients with safe blood-sugar levels. It is estimated that the clinical guidelines have saved the organization \$10 million, or \$2,000 per patient, through improved clinical decision making (see Gillespie, 2000).

12. As noted by the Science Panel on Interactive Communications and Health (1999), self-care books provided to members of health maintenance organizations and Medicare beneficiaries have been shown to reduce office visits and specialty referrals (Vickery et al., 1988); systems to help patients prepare for office visits have been shown to improve treatment outcomes for chronic diseases (Greenfield et al., 1985); computer access to support groups and decision guidance has been shown to help women with breast cancer and patients with AIDS (Gufstafson et al., 1992, 1993, 1994); and shared decision-making tools have been shown to improve health outcomes while reducing the use of surgery and other high-cost medical procedures (Barry et al., 1995; Morgan et al., 1997).

13. These metrics are not independent of each other. For example, a high packet loss rate is likely to lead to low throughput because lost packets must be retransmitted, and the complete message cannot be reassembled until all packets are received.

14. Quality of service is distinct from reliability, which refers to the likelihood that a service remains available at all times. A network may be highly reliable in the sense that it is always possible to obtain connectivity to a given destination, but the same network may lack any assurance of performance (QoS, as defined here).

15. This level of isolation can be achieved even if there is some physical sharing at the

very lowest layer of the protocol stack; for example, the transmission links of the two different networks might share a physical fiber. At the same time, the separation is only as good as the trust of the user in the service provider. A simple misconfiguration of a router could connect a third-party link to a private network. In addition, the service provider has full access to the data carried over a private network.

16. As a result of recent consolidation in the insurance industry, for example, care providers now work with policies established in large corporate headquarters that are greater distances away, and standards for reducing the administrative burden on providers can no longer be set at the state level.

17. For more information, see <<http://www.ccic.gov>>.

18. The formal specification of the NGI program reverses the second and third items in the list above. The order of presentation is changed herein for stylistic purposes and to highlight that the development of testbed networks is just one element of a much broader-based program.

19. The term "telemedicine" refers to the delivery of health services when distance separates the care provider and patient (see Institute of Medicine, 1996). This construction recognizes that a range of different interactions are possible, from videoconferencing at the one extreme to the use of the telephone or text e-mail at the other. Indeed, the most prevalent uses of telemedicine today are not video-based but involve the use of asynchronous store-and-forward systems to exchange still images across networks. Other applications include telephone- or Internet-based systems for monitoring patients in their homes.

20. The study committee visited with the researchers at NASA Ames Research Center as part of this project. A summary of that visit is contained in Appendix A of this report.

21. It is expected that 25 more sites will be added to this testbed in FY00.

22. The vBNS is a nationwide network that supports high-performance, high-bandwidth research applications. Launched in 1995, it is the product of a 5-year cooperative agreement between NSF and MCI WorldCom. Approximately 100 research institutions, chosen through a peer-review process, will be connected to the network. It currently connects 92 institutions.

23. For additional information on Internet 2 and UCAID, see <www.ucaid.org>.

24. Belfast Telegraph Online 10/26/99 as summarized in "Internet 2 Gets Ready to Operate," *Edupage*, November 1, 1999.

25. This information was obtained from the Abilene Web site at <<http://www.ucaid.edu/abilene/>>.

26. For additional information on these networks and the evolution of the Internet more generally, see Chapter 7 in Computer Science and Telecommunications Board (1999).

2

Health Applications of the Internet

Many health-related processes stand to be reshaped by the Internet. In clinical settings, the Internet enables care providers to gain rapid access to information that can aid in the diagnosis of health conditions or the development of suitable treatment plans. It can make patient records, test results, and practice guidelines accessible from the examination room. It can also allow care providers to consult with each other electronically to discuss treatment plans or operative procedures. At the same time, the Internet supports a shift toward more patient-centered care, enabling consumers to gather health-related information themselves; to communicate with care providers, health plan administrators, and other consumers electronically; and even to receive care in the home. The Internet can also support numerous health-related activities beyond the direct provision of care. By supporting financial and administrative transactions, public health surveillance, professional education, and biomedical research, the Internet can streamline the administrative overhead associated with health care, improve the health of the nation's population, better train health care providers, and lead to new insights into the nature of disease.

The capability of the Internet to support these applications depends on whether the relevant technical needs are met and whether the operational aspects of the systems involved are understood and manageable. As with any information technology system, the technical requirements depend heavily on the specific characteristics of the individual systems—the number of anticipated users, degree of real-time interaction desired, number of simultaneous sessions that must be supported, and so on.

Many of these factors, in turn, are influenced by considerations other than network performance. These include organizational competencies, changing preferences and expectations of consumers and care providers, reimbursement policies for different health services, availability of complementary technologies, and laws. The confluence of so many factors confounds attempts to predict viable future applications of the Internet in the health sector.

This chapter presents a broad overview of the types of applications that the Internet can support in consumer health, clinical care, financial and administrative transactions, public health, health professional education, and biomedical research. It draws on a series of site visits by the committee (these visits are summarized in Appendix A) and other briefings to the committee to examine applications that have been deployed and that are still in the early stages of conceptualization. The chapter attempts to assess the technical capabilities demanded of the Internet in terms of bandwidth, latency, security, availability, and ubiquity (as defined in Chapter 1). Specific technical information is presented where possible, but because of the nascent nature of many Internet applications in the health sector, often the most that can be offered is a qualitative assessment. Accordingly, a ranking scale is used to assess the importance of each technical dimension to each class of applications. These dimensions are ranked on a scale of one to four, with one plus sign (+) indicating little importance relative to the other dimensions and four plus signs (++++) signifying the most importance. The chapter also identifies organizational- and policy-level issues that will influence the way the Internet is deployed in different health applications and notes, where applicable, other technologies that must be developed to make certain applications feasible. Specific technical, organizational, and policy issues are addressed in subsequent chapters of the report.

CONSUMER HEALTH

Consumer health is one of the areas that could be most dramatically reshaped by the Internet. Consumer health refers to a set of activities aimed at giving consumers a more pronounced role in their own health and health care, ranging from the development of tools for self-assessment of health risks and management of chronic diseases, to home-based monitoring of health status and delivery of care. This area is similar to public health (discussed later in this chapter) in that it aims to provide consumers with the information and tools needed to improve their health, but it is less concerned with the detection of regional outbreaks of disease and is not part of government-based reporting structures. The Internet could become a significant enabler of consumer health initiatives in that it pro-

vides an increasingly accessible communications channel for a growing segment of the population. Moreover, in comparison to television—also a widely available medium for reaching consumers—the Internet offers greater interactivity and better tailoring of information to individual needs. These capabilities may lead to significant changes in consumer behavior (e.g., cessation of smoking, changes in diet) that could greatly improve health.

Ongoing trends in health care are likely to reinforce the shift toward consumer-oriented health information. Since the mid-1960s, patients have been encouraged to take a more active role in their own health care, and care providers have recognized the value of engaging patients to participate more meaningfully in their own care. Furthermore, attempts by care providers and managed care plans to streamline services and cut costs have shortened hospital stays, increasing the need for patients and their families to understand how to provide care themselves. Greater emphasis is being placed on preventive care, which requires consumers to understand health risks and the effects of different behaviors (such as smoking and dietary habits) on their health. These trends heighten the need for consumers to have access to reliable health information and open channels of communication to care providers and other health professionals.

Consumer health initiatives that rely on the Internet reflect, and could even drive, significant changes in the structure of the health care industry. Concurrent with changes in the economics of the health care delivery system, the duration of a medical consultation is steadily declining, and the availability of practitioners for substantive discussions between visits is decreasing. Continuity of care is increasingly disrupted as patients change care providers in response to changes in their health insurance plans. These trends favor consumers who are well informed and autonomous. Consumer health initiatives attempt to involve patients more actively in care-related decision making and enable them to exercise greater control over their health. Indeed, the Internet could change the culture of health care from one in which patients are viewed as recipients of care to one in which they are partners in care. Eventually, they may be able to use the Internet to access and update their personal medical records or receive care in their homes.

Consumer-Oriented Health Web Sites

Over the past few years, leading providers of health information have identified the Internet as an effective medium for reaching large numbers of health consumers. The most visible aspect of this recognition is the explosion of Web sites geared to consumer health issues (Table 2.1). These sites are dedicated to the diagnosis and management of diseases, the

TABLE 2.1 Examples of Commercial Health-Related Web Sites

Site	Content
Americasdoctor.com	Offers free, private chats with physicians. Also sells health-related items.
Betterhealth.com	Covers various aspects of physical and emotional health. Includes expert advice, feature articles, and support groups.
Discoveryhealth.com	In conjunction with Intellihealth, offers disease-related information, health news, online prescription ordering, and risk-assessment services.
drkoop.com	Offers health information and more than 120 chat groups with advice from a physician. Also allows consumers to check for drug interactions. Plans to add capabilities for consumers to keep track of their medical histories and medical expenses.
Healtheon/WebMD.com	Started as a subscription service for doctors but has a free consumer site that offers health news and information, a physician directory, and condition-specific support groups.
InteliHealth.com	A joint venture between Johns Hopkins University Hospital and Health System and Aetna U.S. Healthcare that offers health news; access to the Johns Hopkins health library, drug databases, and journal abstracts; and catalog items for sale.
Mediconsult.com	Focuses on patients with chronic ailments, offering information on 60 common conditions. For a fee, consumers can pose questions for a physician.
Medscape.com	Offers original, peer-reviewed reports and journal articles organized by specialty and intended for both health professionals and consumers. A dedicated consumer site is under development.
OnHealth.com	Aimed primarily at women, has an alliance with drugstore.com for pharmaceutical purchases and offers guides that rate the health quotient of communities nationally.
Thebody.com	An AIDS and HIV information Web site aimed primarily at the homosexual community.
Thriveonline.com	Features alternative medicines, diet, and exercise tips.

SOURCES: Carns (1999); Nash (1999).

promotion of various healthy lifestyles, and interventions to prevent the onset of disease. The formats range from mailing lists to interactive Web sites, chat sessions, or compilations of online resources. One recent survey suggested that consumers use these sites to gather information on diseases, medications, and nutrition, as well as to find care providers or participate in support groups (Table 2.2).

TABLE 2.2 Primary Health Activities for Consumers to Conduct Online

Activity	Percent of Respondents
Research an illness or disease	62.1
Look for nutrition and fitness information	20.0
Research drugs and their interactions	11.6
Look for a doctor or hospital	3.7
Look for online medical support groups	2.3

SOURCE: *USA Today* (1998).

The network capabilities required by consumer health Web sites are not especially demanding today, but the requirements could grow over time. Most sites offer text and limited graphics, which do not require significant bandwidth, but the availability of greater bandwidth—especially in the local loop—could enable the design of more sophisticated sites offering educational videos for downloading over the Internet. Security requirements are also minimal because personal health information is generally not exchanged on these sites. Protection is needed for financial transactions related to the purchase of health products, but this requirement is no different than that for other e-commerce applications. Similarly, consumer health Web sites do not demand exceptional reliability because they are unlikely to be used for applications in which lives are at stake. However, consumer health Web sites may drive the need for improved privacy-enhancing technologies. The information sought by consumers on the Internet, and the purchases they make, can reveal much about personal health concerns and problems. To prevent organizations from compiling profiles of their health concerns, consumers may demand greater anonymity in their Web browsing and purchasing and tighter restrictions on the ways in which organizations can use information about their habits.

A larger issue is the need for tools to help consumers find information of interest and evaluate its quality. The sheer volume of health information available on the Internet can be overwhelming. For example, a simple Web search for “diabetes mellitus” can return more than 40,000 Web pages,¹ and some 61,000 Web sites contain information on breast cancer (Boodman, 1999). To sort through this volume of material, consumers need effective searching and filtering tools that can identify and rank information according to their needs and capabilities and present it in a form that they can understand, regardless of educational and cultural background. Consumers also need a way to judge the quality, authoritativeness, and provenance of the information. The Internet enables anyone

to publish information, so filtering and credentialing become more important. A recent study found that 6 percent of the 400 sites containing information on a form of cancer called Ewing's sarcoma contained erroneous information, and many more were misleading. Sites contained different (and often incorrect) estimates of basic information such as survival rates (Biermann et al., 1999).

Several initiatives are already under way to evaluate the quality of health information on the Internet. The Department of Health and Human Services' Scientific Panel on Interactive Health Communication calls for disclosure statements on Web sites to make it easy for consumers to evaluate the source and authority of information resources (SCIPICH, 1999). Other efforts focus on systems for classifying health Web sites according to metrics such as accuracy, timeliness, completeness, and clarity.² With these evaluations, standard search engines could provide consumers with a measure of trust in the information they are retrieving—at least to the degree that they trust the organization performing the content labeling. The World Wide Web Consortium, for example, has created a system called the Platform for Internet Content Selection (PICS), which can help users control the types of information retrieved from the Internet.³ To accommodate different perspectives on health and health care (e.g., alternative as opposed to traditional medicine), a wide variety of organizations could rate health Web sites. Additional research may suggest ways of automating the evaluation process, perhaps using metrics such as the number of pointers to, or users of, a given site as indicators of the site's effectiveness (as some search and referral engines are currently doing). Technology could also be used to help prevent alterations of the site's rating to assure consumers that an evaluation was indeed performed by the stated third party. This function requires cryptographic authentication technologies that are currently available but have not yet been widely deployed for this purpose.

E-mail Between Patients and Providers

The Internet can also be used to facilitate electronic communications between patients and care providers, typically in the form of electronic mail (e-mail). To date, e-mail has been used only sporadically between patients and providers, but it is of growing interest. It could prove to be an effective mechanism for improving care and lowering costs because more frequent communications might enable better tracking of a patient's progress or eliminate the need for an office visit. This premise has yet to be tested rigorously in clinical settings, and a number of technical and nontechnical issues need to be resolved (Mandl et al., 1998).

Bandwidth and availability are not issues in the near term because most messages currently consist of text only and are not used for time-critical communications. The most pressing technical issue is security. Most e-mail exchanges between patient and provider involve discussions of personal health information, which must be suitably protected from breaches of confidentiality and, to a lesser extent, alteration. Most e-mail is not encrypted during either transmission or storage, and its point of origin is not authenticated. It is therefore much easier to forge an e-mail message than a clinician's note or telephone call.

Several approaches are available for improving the security of e-mail exchanges. Secure Sockets Layer (SSL) encryption, which is commonly used to encrypt e-commerce transactions (see Chapter 3 for a description of the technology), can be used to protect communications between a user's personal computer and the electronic mail server. Other protocols, such as Pretty Good Privacy (also described in Chapter 3), can be used to protect communications as they move across the network between the sender and recipient. User authentication can be enhanced through the use of nontrivial user names and passwords or more secure forms of authentication, such as those based on public key encryption (also described in Chapter 3).

The more daunting barriers to patient-provider e-mail are institutional policies for confidentiality and for integrating e-mail into work flows. Most e-mail systems are without even the most basic protection of the confidentiality of message contents. Mail received at the place of work is, by law, fully accessible to the employer. One study showed that patients are hesitant to use e-mail from work to communicate about their health for fear that employers or insurance companies might use the information in ways that affect them personally (Fridsma et al., 1994). To avoid the risk of having messages discoverable at a place of work or other sensitive locations, individuals can store their e-mail files on the server of a trusted third party and/or encrypt messages for storage, but rules regarding disclosure still need to be developed.

Health care organizations are also concerned that e-mail might overload care providers with yet another task in the context of increased clinical and administrative burdens. There are related concerns about the liability of providers if, for example, they miss a subtle but (in retrospect) irrefutable and important question or comment in a patient's electronic note. Many organizations have yet to establish policies regarding the quality of service, such as a maximum time to respond or even acknowledge receipt, that patients can expect from e-mail with providers. Another important concern is economic: there are currently no mechanisms for paying providers for what could be as taxing or time-consuming a clinical

activity as any in-person clinical visit. Furthermore, no policies and procedures have been developed for incorporating e-mail into electronic patient records. As a consequence, decisions made on the basis of e-mail information are at risk of having no documented basis in the record.

Safe and effective use of e-mail for clinical discussions between patients and providers will require the development of policies to govern its use. These policies will need to address issues of confidentiality, data integrity, authentication, timeliness, and the appropriateness of the use of e-mail for different kinds of discussions. In some cases, telephone or face-to-face conversations may be considered a more appropriate form of communication. These policies will need to be articulated to all consumers and also embodied in the e-mail user interfaces so that health care consumers can have realistic expectations about the use and safety of clinical e-mail.

Online Health Records

The Internet is emerging as a medium for giving consumers direct access to their personal health records. Historically, care providers have maintained voluminous records of patient encounters within their organizations, documenting dates and times of consultations, diagnoses, lab results, prescriptions, and more. These records are maintained and largely controlled by care providers, although patients have the right, in some states, to review their records and propose amendments as necessary. In the past two years, however, a number of new Web sites have begun to allow consumers to store their own health records online.⁴ The potential benefits of these sites are many. With them, consumers can create comprehensive, longitudinal records that capture information about the care received from different organizations over an extended period of time. Consumers can use these records to help monitor and evaluate their health status, and they can grant access, if they wish, to different providers for purposes of care. Many sites provide some sort of override feature that enables care providers to gain access to a patient's records in an emergency situation—something that is much more difficult to do if the records are not stored online.⁵

Like e-mail used for clinical purposes, Web-based medical records require considerable attention to security to minimize the risks of inappropriate disclosure. Personal medical records must be protected against inappropriate disclosure, both to outsiders who attempt to break into the system and to those who operate and maintain the Web sites. Most existing services use SSL encryption to protect data communications between users and the host Web site and a combination of user names and passwords (transmitted securely over the Internet) to authenticate end users.

Systems operating with user identification and passwords can provide reasonably—but not fully—secure access to many types of applications. If online records become more widely used in the provision of care, then it may be advisable to enhance the robustness of user authentication, perhaps with public key encryption systems and user certificates (see Chapter 3). The PCASSO system being developed by Science Applications International Corp. (SAIC) and the University of California at San Diego, for example, uses public key encryption and a challenge-response token, as well as a password, to protect patient information at a far higher level than is possible with SSL.⁶

Other technical requirements will be modest in the near future unless online patient records become more complex and more widely used in the provision of care. At present, most online medical records consist primarily of text and demand little bandwidth for fairly rapid downloading. If such records begin to include medical images (e.g., X rays, computed tomography (CT) scans, and mammograms), then much higher bandwidth would be needed for timely downloading (see the section on medical images below). Similarly, reliability requirements are not high because online records are still supplements to, as opposed to replacements for, the records maintained by provider organizations; an inability to access an online record is unlikely to interfere with the provision of care. If online records become more widely used and more complete than providers' records, then reliability could become more of a concern. Scalability is not an issue, either, because records are not needed simultaneously by multiple users.

Ubiquity of access to the Internet is a significant consideration in the development of online medical records because it would ensure that all consumers could keep such records and that those records could be accessible from a large number of unpredictable locations, such as a consumer's home or office, a care provider's office, or an ambulance responding to an emergency. A number of business and policy issues need to be resolved as well. Organizations that store online health records will need to develop policies that balance the need for privacy and security against the need for ready access to records by patients and eventually by care providers and perhaps insurance companies, researchers, and others. Rules may also be needed to govern organizations' use of the online records they maintain. Under what conditions will they be able to provide consumers with recommendations about necessary medical tests or possible drug interactions? To what extent should they be allowed to mine patient records for information that might lead to direct marketing efforts? Under what circumstances should records be made available to public health agencies and researchers?

Patient Monitoring and Home Care

The Internet offers the opportunity for improved monitoring of consumer health and, potentially, provision of in-home care through video-based consultations with care providers (discussed in the Clinical Care section, below) and control of medical equipment (e.g., pacemakers and dosimeters) deployed in the home. The goals of such activities are to assist in the early detection of potential health problems, ranging from heart attacks to congestive heart failure and diabetes, and to reduce the need for clinical intervention and costly hospital stays.⁷ Remote consultations to the home may be most useful for monitoring patients with ailments such as congestive heart failure and end-stage liver disease. These applications do not require video imagery; the provider simply listens to heart and lungs, taking vital signs and pulse oximetry. In-home care is consistent with existing trends in the health care industry. Since 1975, the number of home health agencies has grown from 2,300 to almost 8,500, while the number of hospital beds per 1,000 enrollees has declined from 51 to 28.⁸ Similarly, the number of patients receiving home care nearly tripled between 1982 and 1994. These trends reflect, in part, attempts by health insurers and health management organizations to reduce the costs of care associated with long hospital stays.⁹

To date, few attempts have been made to monitor patients at home. Most efforts have focused on chronic conditions, such as diabetes, asthma, and congestive heart failure, for which well-established protocols exist for home care. The devices used for monitoring are minimally modified copies of devices used in hospitals. Little effort has been made to develop or distribute small devices that mimic the functionality of much larger hospital counterparts with automated quality control and calibration and remote polling and configuration by authorized care providers. Almost none of these devices is as portable or easy to use as a standard pager. In part because of these limitations, home monitoring has not grown as much in popularity as have consumer information on the Web and patient-provider e-mail.

In January 2000, however, Medtronic Inc. announced plans to work with IBM Corp. and Microsoft Corp. to develop a system that will enable heart patients with implanted pacemakers, defibrillators, and experimental cardiac-pacing and -monitoring devices to transmit cardiac data over the Internet to their cardiologists. Eventually, care providers may be able program the devices over a secure Internet connection without requiring patients to visit their offices. Developers of the system posit that it will result in fewer office visits and hospitalizations, thereby lowering costs while improving patient monitoring and care, but a means of charging for the monitoring service has not yet been devised. Medtronic hopes that its

secure Internet system will find utility beyond cardiac patients, perhaps allowing patients with implanted drug pumps to have their doctors change the drug regimen remotely over the Internet (Burton, 2000).

Continued advances in computing and communications technologies could enable more widespread deployment of home-based health monitoring systems. For more than two decades, the feasible density of transistors on an integrated circuit has been increasing by a factor of 10 every 7 years. Memory densities have increased even faster, gaining an order of magnitude every 6 years. As a result, medical devices such as stethoscopes, glucometers, and electrocardiogram monitors already can be equipped to support Internet connections and deployed to consumers at low cost. Over time, computing and communications capabilities will probably be incorporated into a number of other devices that could serve as sources of health information, whether bathroom scales or exercise equipment. If a house is networked, then it would be possible to use a personal computer to connect and control a number of medical monitoring devices. Although the number of homes with conventional local area networks (LANs) is small (mainly because of the high cost of wiring a house appropriately and the disruption involved), Ethernet-like connectivity can be provided to any room in a house through devices that are either wireless or attached to the existing telephone or electric wiring.

Indeed, advances in microelectromechanical systems (MEMS) devices, combined with those forecast in microelectronics, biosensors, and biomaterials, could lead to revolutionary changes in therapies, delivery of medication, and monitoring and alerting systems for the elderly and those with chronic conditions. Devices already on the market, such as pacemakers, wireless stethoscopes, and blood sugar monitors, could be augmented with networking capabilities. High-resolution digital video cameras that are acquired by consumers for recreational or other purposes might become useful in health care applications.

Home-based monitoring is unlikely to require high-bandwidth connections from homes to the Internet because individual messages tend to be small. In demonstration projects, however, investigators have had to work hard to ensure that all participating patients had uninterrupted access to even modest bandwidth, often contracting with the local cable or telephone company to hook up a specific home. The installations, connectivity, and subsequent support costs have accounted for a large portion of the cost of the monitoring efforts. Bandwidth is a more significant issue for provider organizations, which will need to ensure that their facilities can handle the aggregate load of monitoring numerous devices (e.g., if hundreds of thousands of patients with congestive heart failure are monitored at home). At this point, it is difficult to estimate the aggregate bandwidth needed by providers of monitoring services because it is

not clear how many patients would be monitored simultaneously or by the same server. The load on the network might be reduced if monitoring hardware reported only summary data and any anomalies detected, unless detailed raw data were requested. Home-based monitoring would require high reliability to ensure that data can be regularly and routinely transmitted and high levels of security to prevent alteration of data as they transit the network.

Other factors are equally or more important to the evolution of home-based monitoring. Even modest monitoring efforts will not be effective unless mechanisms are deployed to enable care providers to review the monitored data, identify worrisome outliers, and respond in a timely way. The need for oversight of such large numbers of patients at home could result in the emergence of a new category of ancillary health professionals. Furthermore, the effective use of such large amounts of monitored data will require automated data reduction and intelligent data analysis techniques. For some populations (e.g., patients with diabetes or congestive heart failure), this approach could enable fine-grained medical oversight that could result in improved short- and long-term outcomes. But, if used inappropriately, it could also afford vast opportunities for unnecessary and unwanted intrusions into the privacy of all health care consumers.

The benefits of home monitoring cannot be fully realized unless reimbursement is provided for virtual home visits and remote monitoring. In addition, policies for protecting the confidentiality of data gathered in this way will have to acquire the force of law if abuses are to be prevented. Even the strongest cryptographic methods cannot prevent the subversion of a system by parties with strong financial interests in breaching patient data confidentiality. The challenges that must be overcome to provide this level of surveillance appear to be more nontechnical than technical, and they include issues of organizational structure and reimbursement rather than networking capabilities.

Beyond the use of the Internet for home monitoring is the possibility of using it to modify home medical devices remotely. After a remote consultation or review of home monitoring data, a care provider might, for example, want to change the setting of a threshold on a patient's pacemaker, alter the parameters for a programmable insulin pump, or increase the dose delivered by an infusion pump for an oncology patient. Such capabilities are already used to control spacecraft and other remote equipment and could have a large impact on health care, especially in rural environments. Although remote control of such equipment will be unnecessary (or unnecessarily paternalistic) for some patients, it might be appealing in cases involving disabilities or simply for the sake of convenience.

The control of remote medical equipment would pose a number of challenges for the Internet—or any other control network. Although bandwidth requirements would be minimal because the commands would likely consist of short messages, the requirements for security and availability would be extremely high. Data would need to be protected from intentional and unintentional corruption to ensure that commands are transmitted as intended. High levels of authentication would be needed on both ends of the connection to ensure that the appropriate equipment is being manipulated and that only authorized personnel send modifications. The network would have to be protected from denial-of-service attacks that could prevent the receipt of update information.

Technical Requirements for Consumer Health Applications

The technical capabilities needed to support consumer health applications of the Internet are modest, largely because the systems developed to date have had to rely on the existing Internet infrastructure. Early experimentation with more advanced systems that provide real-time video connections between care providers and patients (or their parents) at home demonstrates the increased demands that consumer health could place on networking resources. The discussion below reviews the technical needs for consumer health applications with respect to bandwidth, latency, availability, security, and degree of access. As noted at the beginning of this chapter, the importance of each capability is indicated on a four-point scale, with one plus sign (+) indicating limited needs and four plus signs (++++) signifying an important need.

Bandwidth ++

Consumer health applications vary considerably in the bandwidth they demand. The retrieval of information from health-related Web sites demands little bandwidth on the consumer end, but the potentially large volume of requests made of any particular site could drive up the aggregate bandwidth requirement on the information provider's side. Access to patient health records could demand somewhat greater bandwidth than is typically available today or significantly greater if records include enhanced content, such as medical images or videotapes of telemedicine consultations.

Latency +

In general, applications that support consumer health do not require the instantaneous delivery of information, so the latency requirements of

the Internet are not great. In some patient-monitoring applications, timeliness is a concern, but delays of a few seconds would not threaten a patient's well-being. Latency could become more of an issue if online medical records became the norm and care provider organizations needed timely access to them for purposes of treating patients. In many instances, however, records could be uploaded from remote sites in advance of scheduled appointments, and latency would be a significant issue only in emergency situations.

Availability ++

The need for network availability differs significantly among consumer health applications. The Internet is already sufficiently available for the distribution of health information to consumers and for exchanges of e-mail between patients and providers. Somewhat greater availability would be needed for remote monitoring and remote control operations, although most home monitoring devices and medical equipment could be designed to buffer enough data to overcome short lapses of connectivity. Home monitoring and control will not become commonplace, however, until providers (and consumers) of such services receive guarantees that lengthy network outages will occur very infrequently.

Security ++++

Many consumer health applications demand high levels of security. Although this is generally not an issue with respect to the downloading of health information from consumer Web sites, access to online patient records demands confidentiality because such records contain personal information. The same is true for e-mail messages between patients and providers that contain personal health information. Data from remote patient monitoring devices also require security to prevent corruption (intentional or unintentional) during transit across the network or after storage. As described in greater detail in Chapter 3, both technological and administrative solutions are required to secure these types of consumer health information. For example, authentication technologies are needed to validate the identities of those requesting and transmitting data. Effective controls are needed to prevent users from accessing information about other consumers. Encryption technologies are needed to protect the confidentiality of data transmitted across the network and ensure its integrity. Policies will be needed to determine who can have access to consumer health information and under what conditions. Security requirements will grow as consumers use the Internet to store, retrieve, and update their personal health records.

Consumer health applications also raise the issue of online anonymity. Searches for online information can reveal a lot about consumers' health concerns, as can their online purchases of prescription and non-prescription pharmaceuticals. Given the sensitivity of some of these conditions, the demand for anonymous Web browsing and even anonymous e-commerce could grow. Consumers may also demand greater anonymity in e-mail to online physician services offered by some consumer Web sites. Whether anonymity is desirable from a social perspective—and under what circumstances (e.g., anonymous Web browsing may be more plausible than anonymous e-commerce)—is an issue for continued debate and discussion.

Ubiquity +++++

Key to the success of consumer health applications is widespread access to the Internet. As noted above, many consumer applications currently demand only moderate bandwidth and latency, meaning that standard modem access to the Internet, at 28.8 to 56 kilobits per second (kbps), may suffice. Additional bandwidth could be needed if online access to health records and downloading of educational videos become more popular and widespread and if online health records grow to include not just text but medical images and perhaps even videos. As discussed in the next section (Clinical Care), remote medical consultations to the home over the Internet could require bandwidth of 128 kbps or more in both directions—if such applications prove technically feasible and economically viable. The larger issue may be that of ensuring equitable access to health resources by different demographic groups. There are already considerable differences in access to health care in the United States; ensuring that differential access to the Internet along demographic lines does not exacerbate this imbalance could become an increasingly important issue, especially if the provision of health care moves online.

CLINICAL CARE

The Internet offers several avenues for augmenting the health care services in clinical settings. Remote video consultation, for example, could give consumers greater access to skilled health professionals regardless of geographic proximity. The use of the Internet to transfer medical images to expert interpreters could accelerate and improve the diagnostic process as well as reduce costs. Virtual reality tools could help surgeons plan medical procedures and improve their use of information during procedures. The use of the Internet to access and assemble health records could give a provider improved information for treatment purposes, regardless

of whether the patient is a regular client or a stranger. Each of these applications poses a range of technical challenges for networking researchers and other information technologists. In most cases, the applications have not yet been demonstrated on a scale sufficient to determine their medical efficacy or influence on costs of care. As the discussion below demonstrates, the use of the Internet in clinical care will be influenced by a range of technical, organizational, and policy issues.

Remote Consultation

Remote medical consultation has long been pursued as a means of overcoming the unequal distribution of clinical expertise. It is a method of offering expert consultations to patients in remote rural areas, for example, or underserved urban areas or prisons. Even where clinical expertise is available, but inconvenient for either the patient or the provider, remote medical consultations may be a cost-effective alternative to staffing multiple clinics with subspecialists. Remote consultations may also be useful to specialized service organizations that attempt to establish economies of scale for particular types of clinical service, such as the interpretation of radiological images (e.g., CT and magnetic resonance images), while also developing more effective bargaining units for health care contracting. These organizations, which are becoming more numerous, can benefit insofar as their reach is extended beyond their immediate geographical area, allowing them to serve a broader pool of consumers.

The network performance required for remote consultation is variable and depends on a number of factors, including (1) the resolution required in the transmitted signal or image to support diagnosis, (2) the timeliness with which data must be received and interpreted (e.g., whether the system is used for real-time consultation or asynchronous review), (3) the degree to which the data may be compressed, (4) whether the entire data set must be transferred or application-specific decisions can be made about which subsets to transmit, and (5) whether the transmission can be considered only on a point-to-point basis or as part of aggregate traffic. These factors vary significantly across different applications and operating modes. For example, psychiatric evaluations may be viable with video that has lower resolution than a cineo-angiogram, but the application needs to operate in real time rather than in a store-and-forward mode for review at a later time.

No conclusive studies have been done regarding the bandwidth needed for different applications; the results of research on this issue generally depend on the provider involved and the study structure. However, reasonable guidelines can be gleaned from experiments conducted to date. Practitioners at East Carolina University (ECU) in Greenville,

North Carolina, for example, have considerable experience with remote consultations, having conducted about 3,000 real-time consultations in 31 different specialties since establishing a telemedicine program in 1991 (see Appendix A for more information on the ECU program). The five most active specialty areas have been dermatology; cardiology; neurology; gastroenterology; and allergy, asthma, and immunology. Practitioners have found that the bandwidth needed for most real-time, video-based consultations varies from 128 kbps to 384 kbps, depending on the degree of resolution needed for diagnosis and the rate of motion in the video (Table 2.3).¹⁰

For some procedures, such as cineo-angiograms, echocardiograms, and gait analysis (Box 2.1), more bandwidth can be advantageous. Cineo-angiograms, for example, can be transmitted at 384 kbps, but 768 kbps produces better results. Cineo-angiograms are generally not performed in real time (because the source is film); therefore, they can be done in a store-and-forward mode, with bandwidth needs determined by the number of cases to be examined on a given day and the desired turnaround time. Adult echocardiograms are often done in real time and can be read adequately at 384 kbps, but pediatric cardiograms may require 768 kbps because the area being observed is so small. Extensive testing with echocardiography indicates that data rates in excess of 1.5 megabits per second (Mbps) probably do not contribute to an increase in clinical efficacy, but further investigation is under way.¹¹ For remote analysis of a

TABLE 2.3 Nominal Bandwidth Requirements for Different Telemedicine Applications

Type of Telemedicine	Needed Bandwidth ^a	Examples
High resolution, no motion	Store-and-forward	Radiology, dermatology, pathology
Medium resolution, low motion	128 kbps	Stethoscope, visual exams, psychiatric consultations, gastroenterology
Medium resolution, high motion	384 kbps	Cardiology, neurology, and emergency room consultations
High resolution, high motion	768 kbps	Cineo-angiography and echocardiograms
Very high resolution, high motion	Up to 2.5 Mbps	Gait analysis

^akbps, kilobits per second; Mbps, megabits per second.

SOURCE: David Balch, East Carolina University, personal communication, February 2, 1999.

BOX 2.1
Examples of Video Images Used in Medical Diagnoses

Cineo-angiograms are movies made from X rays taken in rapid succession while dye is injected into the blood vessels of the heart. The motion picture shows blood flowing through the heart and reveals blockages of the arteries of the heart and any abnormalities of the motion of the heart's pumping chambers.

Echocardiograms use high-frequency sound waves to generate images of the heart in motion—like pictures of submerged objects obtained with sonar. Echocardiograms reveal the size of the heart's chambers, images of the heart muscle contracting and relaxing, and information about heart valve function.

Gait analysis is the measurement and interpretation of how a person walks. It is useful in diagnosing problems affecting both nerves and muscles. Although often taken for granted, the act of walking is a remarkably complex sequence of balance and motion. Neurological problems such as Parkinson's disease and stroke may affect gait in distinctive and treatable ways. Muscle diseases such as muscular dystrophy and autoimmune disorders such as systemic lupus erythematosus may lead to walking difficulties that can be diagnosed and treated with the help of gait analysis.

patient's gait, bandwidth of up to 2.5 Mbps may be necessary, but this application has not been extensively evaluated.¹²

For the most part, remote consultation programs rely on dedicated networks—not the Internet—to provide connectivity between remote clinics and a centralized consulting facility. The ECU program, for example, uses an amalgam of microwave links, T1 lines, telephone lines, and integrated services digital networks (ISDN; see Chapter 3) for a variety of applications. The ECU program and other experimental programs, such as the National Laboratory for the Study of Rural Telemedicine at the University of Iowa, also make use of statewide fiber-optic networks for connectivity between some sites.¹³ Although costly, dedicated lines have been viewed as the most effective means of guaranteeing access to adequate bandwidth as needed. The Internet does not yet offer the quality of service needed for real-time video consultations. Some organizations, including ECU, have begun to shift their systems to the IP in anticipation of connections to the Next Generation Internet, but they will continue to rely on dedicated links until a more viable Internet-based infrastructure is available.

Continued advances in telecommunications infrastructure could cause

remote consultations to become less the province of a few sites equipped with specialized telemedicine rooms and more a routine component of the services offered by health plans. A number of integrated health care delivery systems have begun to experiment with remote consultations (typically over leased lines) to provide specialty services in outlying areas. If the Internet could support such capabilities, then remote consultation could become more common, perhaps even extending beyond regional boundaries. Indeed, an enhanced Internet could help extend teleconferencing to the home, enabling consumers to ask for videoconferences with care providers whenever health problems require immediate attention. Such capabilities could dramatically transform health care by eliminating many office visits.

Regardless of whether the patient is at home or somewhere else, remote consultations require sustained bandwidth in two directions—from the patient to the provider and vice versa. This stands in contrast to many high-bandwidth applications, such as entertainment, education, or scientific visualization, in which sustained access to high bandwidth is needed in one direction only, from a centralized distributor of content to a recipient. Clearly, the need for high bandwidth in two directions is not unique to health care; many businesses need bidirectional bandwidth to support collaborations between workers in different locations. However, remote medical consultations to the home demand that such capabilities be available from many locations, not just from corporate offices that might already lease a high-bandwidth access line to the Internet or a private, corporate network. The most likely users of remote medical consultations would be primary care providers and patients living in rural or remote areas, many of whom have limited access to high-bandwidth Internet connections (see Chapter 3).

The future of remote consultations will be influenced by a number of factors beyond network technology. Other technical challenges arise from the need for appropriate data acquisition devices to digitize the observations involved in the consultation. Although many medical devices have been instrumented to allow for remote control and data acquisition, few of them have achieved mass market acceptance. As a result, such devices tend to be confined to a few specially equipped rooms in institutions supporting remote medical consultation. Many more of the devices used as part of routine physical exams could be modified for data acquisition and control. They could then be deployed in patients' homes to facilitate home-based consultations with patients whose diseases require intensive monitoring and oversight. In the near term, these devices might be individually configured to work with a home computer and to relay information to a remote care provider over the Internet. Remote care providers might even be able to exercise some control over these devices, whether

adjusting their sensitivity or other operating parameters. Eventually, the devices could be designed to connect automatically to the Internet and be configured by a remote Web browser. Initially, this might be cost-effective for only small, high-risk populations, but remote consultations to the home could become more popular as the technology continues to evolve and costs decline.

Beyond technical challenges, a number of organizational and policy issues need to be resolved if remote consultations are to become more viable in the future. Health care organizations need to develop viable business models for remote consultations that meet the needs of different users. Will remote consultations bring in income directly, or will they generate revenues indirectly by channeling patients into a provider's health care system? East Carolina University, for example, has developed at least five different business models for its services, but only one—providing services to prison inmates—has proved profitable. The others typically operate with grants from federal agencies or are seen as experiments to broaden the reach of local provider organizations. Profitability is currently constrained by the fact that many health plans, including Medicare, do not yet routinely pay for remote consultations, although experiments are under way to examine alternative repayment schemes (see Chapter 5). Other issues, such as state-based licensure of health professionals, impede attempts to deliver remote consultations across state lines.

Medical Imaging

Closely related to the provision of remote medical consultations is the use of communications networks to transfer still medical images. This capability could enable care providers to retrieve digital images from an online repository (often referred to as a picture archiving and communications system, or PACS), send images to specialists for interpretation (a form of remote consultation), or receive information from emergency vehicles responding to a call. The potential benefits of such systems could include the following:

- *Improved management and use of medical images (i.e., reduced probability that images will be lost or incorrectly filed).* Physicians at the University of California at San Francisco (UCSF), for example, noted during the committee's site visit (see Appendix A) that before establishing their PACS, 15 to 20 percent of radiographic images were lost and hundreds went unread.
- *Improved quality of care through expert interpretation.* Five university medical centers, including UCSF, have established an expert radiographic

interpretation center, Telequest, which accepts images from a variety of clients, including rural clinics, and provides transcribed diagnoses. This experimental program makes expert interpretation more widely available throughout the country.

- *Reductions in the cost of radiological interpretation.* Centralization of expertise could reduce health care costs by obviating the need for individual provider organizations to maintain more expert radiologists than they can keep busy.

From a networking perspective, the challenge in remote imaging (sometimes called tele-imaging) is the size of medical images. Typical uncompressed radiographic files range from about 25 kilobytes (kB) for a nuclear medicine image to 50 megabytes (MB) for digitized mammograms, but multiple images often are needed, either to provide a complete view of the object of interest from different angles or to compare various images. Hence, the size of an uncompressed radiographic study can range from 1 to 2 MB for a nuclear medicine study to almost 200 MB for digitized mammograms (Table 2.4). The size of these studies is expected to grow as imaging technology advances; image resolution is expected to improve by a factor of 10 or more in cases such as cineo-angiography. As of early 1999 researchers at UCSF were working with digitized cineo-angiograms that were 60 MB in size and with intravascular ultrasound images that were 50 MB.¹⁴ High-resolution electron microscopes can produce individual images that are 2 MB in size, but such instruments are available only at a small number of research centers.¹⁵

TABLE 2.4 Nominal File Sizes of Common Medical Images

Image Type	Image Size (bits)	Images per Exam	Size of One Exam (MB)
Nuclear medicine	128 × 128 × 16	30-60	1-2
Magnetic resonance imaging	256 × 256 × 12	60	6
Ultrasound (color)	512 × 512 × 24	20-230	16-180
Digital angiography	512 × 512 × 8	15-40	4-10
Digitized electron microscopy	512 × 512 × 8	1	0.26
Digitized color microscopy	512 × 512 × 24	1	0.79
Computed tomography	512 × 512 × 12	40	20
Computed radiograph	2,048 × 2,048 × 12	2	16
Digitized X rays	2,048 × 2,048 × 12	2	16
Digitized mammography	4,096 × 5,625 × 16	4	184

SOURCE: Huang (1996, 1999).

The bandwidth required to transmit these images is determined by several factors, including the amount of time in which the image must be transmitted and the degree of compression that is allowable without degrading the image so much as to impair interpretation and diagnosis. Lossless compression techniques can reduce image size by a factor of 3 or 4; lossy compression techniques can reduce images by a factor of 10 to 20 without sacrificing diagnostic quality in some applications (Lou et al., 1997). Acceptable compression levels vary by application domain (e.g., teleconferencing versus radiology) and by intended user (e.g., radiologist versus primary care physician). With digital mammography, the maximum degree of acceptable compression is controversial because of concerns over the loss of detail.

In many applications, images can be sent, often as e-mail attachments, to a remote site for interpretation and diagnosis within 1 or 2 days. This technique does not place extreme demands on the network and has been used successfully by several organizations, even at low bandwidth. Several years ago, for example, Massachusetts General Hospital in Boston used regular voice lines at 9.6 kbps to receive radiographs and CT images from Saudi Arabia. The images were compressed at ratios of 20 to 1 and 10 to 1, respectively, so that exams could be transmitted in 20 minutes to 1 hour and interpretations could be provided in 1 to 2 days (Huang, 1996; K.J. Dreyer, Partners Healthcare System, personal communication, 1999). Increased use of this technique could demand greater bandwidth, however. A busy mammography center may perform 80 to 100 examinations per day. If all these images were sent out for interpretation, an average sustained throughput of almost 1 Mbps would be required without compression and almost 100 kbps with 10 to 1 compression. This admittedly high-end application is within the average performance capabilities of the Internet backbones but not of all Internet service providers.

A desire for faster turnaround in the interpretation of medical images could increase demands on networking resources, even if the volume of examinations transferred across the network is small. Faster networks would enable various types of service improvements. For instance, they could enable remote experts to provide faster diagnoses or second opinions to help referring physicians plan follow-up treatments while their patients are still in the office, or they could enable the remote experts to provide real-time advice, such as suggesting a need for additional images to aid in diagnosis before a patient leaves the mammography center. In the first example, sometimes referred to as teleconsultation, a response may be desired within 30 minutes or so; in the second example, sometimes referred to as telemanagement, a response may be desired in near real time. Given the size of the images to be transmitted and the possible need for reference sets, the bandwidth demands could be tremendous.

With lossless compression of 4 to 1, an entire mammography study would require 1.7 Mbps to be transmitted in 2 minutes to allow near-real-time interpretation. Such a capability would not be a necessity in most cases; expert reading of mammograms in real time is not needed for regular screenings, but it can be useful if potential abnormalities are discovered.

In some imaging applications, high-speed networks become less important if the data can be intelligently processed prior to transmission. For example, during the development of PACS for integrated health care delivery systems, users often specify that several care providers need to be able to simultaneously access uncompressed medical images from any location in the system within 2 seconds. Such characteristics can translate into a requirement for network bandwidths of 100 Mbps (typically a LAN) and specialized high-resolution monitors, which can make such systems costly (in the \$2 million to \$3 million range). With some attention to physician's schedules, however, systems can be developed that store images locally on a computer that the physician is likely to use, thereby reducing the strain on the network. Furthermore, systems can be developed that operate efficiently on 10 Mbps LANs and standard computer monitors. For example, most computer monitors cannot display a full-screen medical image at full resolution; if the maximum resolution of a clinician's video screen is $1,024 \times 768$ pixels \times 12 bits of gray scale (or 1.2 MB), then there is no use in sending all 128 MB of a digital mammogram study. Instead, a lower resolution image could be sent and additional detail could be requested on smaller portions of the image as the radiologist identifies areas of interest. A variety of commercial solutions are now available that enable the design of tele-imaging systems that appear highly responsive without greatly increasing network performance requirements (Box 2.2). Such systems tend to transmit only portions of a complete image at any one time and therefore place less stringent demands on network capabilities than full-screen, full-resolution systems. Continued evaluation will be needed to determine the relative effectiveness of these alternative designs in diagnosis medical conditions.¹⁶

Data security is also important to teleradiology applications. Both patient confidentiality and data integrity must be maintained during image transmission and storage. Confidentiality can be maintained through the use of a host of technologies for authenticating users, controlling their access to images, and encrypting transmissions (see Chapter 3). Data integrity can be maintained—an especially important function given the ease with which a digital image can be altered—through the use of technologies such as digital signatures, which are used in a number of e-commerce applications. Nevertheless, trade-offs need to be made between the level of protection of digital images and other considerations, including cost and ease of use (Huang, 1996).

BOX 2.2
**A Picture Archiving and Communications System for
Personal Computers**

The typical Picture Archiving and Communications System (PACS) provides multiple users with rapid access to imaging studies (e.g., computed tomography, magnetic resonance, and X-ray images). Imaging studies are typically 40 to 300 megabytes in size, and users want access in 2 seconds or less. Assuming 10 simultaneous users, the system must provide more than 100 Mbps of bandwidth. Although compression can reduce the bandwidth requirement, it can also result in the loss of subtle features that may be important to proper diagnosis. All these requirements can lead to costly systems that cannot take advantage of existing local area networks (LANs) in many health care organizations.

An alternative that has been pursued by Stentor, Inc., is a system that relieves pressure on network bandwidth by sending a lower-resolution image to a PC-based workstation and allowing the user to zoom in on particular areas. Images are represented as wavelet transformations (a mathematical encoding of an image based on wavelet functions). As the user zooms in on the image, the client machine asks the server for a block of wavelet coefficients, which, when received by the client, are transformed into an image that allows additional details to be seen. The update time depends on the size of the window being used, not on the level of resolution; a PC-based computer can transform a 2,500 by 2,000 pixel image in just 1.3 seconds. The system preserves network performance because it sends packets of 4 to 32 kilobytes rather than relying on dedicated network connectivity between the client computer and the image server. As a result, it can operate across slower LANs than is typical for PACs, allowing users to access full-resolution images from virtually any location within a health care facility.

SOURCE: Based on a presentation by John Huffman, chief technology officer, Stentor Inc., to the study committee, March 2, 1999, Washington, D.C.

Clinical Transactions

Several transactions form essential components of clinical care: (1) administrative functions such as referrals, practice management, and billing (addressed in the section below on financial and administrative transactions), (2) distribution of medical supplies and purchasing and inventory control, and (3) clinical functions, such as reporting of results from testing laboratories and interinstitutional communication of health information. While some progress has been made in the first two of these areas, progress in the third has been limited. Some health care organizations use a Web-based infrastructure for reporting laboratory results within an institution, but few integrate the laboratory data into the clini-

cal information system containing patient records. Fewer still use the Internet for exchanging patient records among affiliated or unaffiliated health care organizations. As a result, when patients arrive at a health care organization for the first time (perhaps after a referral or changing health plans) or visit the emergency room of a hospital they have not visited before, their medical records are either inaccessible or reduced to a photocopied or faxed subset of the paper record in another institution. This is often true of communications between departments even within the same institution because only a small minority of health care institutions have at least some form of enterprise-wide clinical information system.

Greater use of the Internet to facilitate exchanges of clinical information could improve the quality of care by making better and more complete information available to care providers. A recent report from the Institute of Medicine identified medical errors as the source of much unnecessary morbidity and mortality (IOM, 1999). By integrating the clinical transactions of all parties to health care delivery (hospitals, pharmacies, clinicians, insurance) across the Internet, there is a significant opportunity to detect and prevent such errors.¹⁷ Use of the Internet for transferring medical records would enable care providers to better treat patients who become ill or are injured while traveling or who have not previously been under their care.

Despite the lack of effort in this arena to date, the Internet appears to provide a viable medium for use by hospitals to share patient health records for the purpose of improving care. This capability was explored through the World Wide Web Electronic Medical Record System (W3EMRS), developed by researchers at Boston's Children's Hospital, Beth Israel Hospital, Massachusetts General Hospital, and the Massachusetts Institute of Technology. The original purpose of the system was to allow the sharing of clinical information across the Internet among emergency room clinicians at the three participating hospitals. The system was subsequently deployed for the sharing of birth data and perinatal maternal data among Brigham & Women's Hospital, Beth Israel Deaconess Medical Center, and Children's Hospital for the management of newborn infants with jaundice at Children's Hospital or one of its affiliated practices. It was also adapted for use within the seven affiliated hospitals of the Boston-area CareGroup; the system is accessible to all authorized clinicians and saves an estimated \$1 million annually by reducing the time spent searching for records, time needed to admit a patient, number of admitted patients, length of hospital stays, and time spent in training. The impact on patient retention and member attraction is projected to increase revenues by \$3 million to \$4 million per year (see Box 2.3 for

additional information about the W3EMRS system and its implementation within CareGroup).

The bandwidth needs for exchanges of clinical information vary with the size of the records to be exchanged, the number of records that are transmitted in a given period of time, and the timeliness with which records must be accessed. The size of a medical record transmitted electronically between sites can vary considerably, from as little as 1 kB to as much as several gigabytes if the record contains several medical images. In general, older and sicker patients have the largest records. Paper charts 3 to 4 inches thick and divided into several volumes are not unusual in hospitals, like those in CareGroup, that serve large numbers of such patients. The number of records transmitted also varies considerably.

BOX 2.3

Linking Medical Records Via the Web

The World Wide Web Electronic Medical Record System (W3EMRS), developed by several Boston-area health care organizations, demonstrates the capabilities of the Internet in supporting exchanges of medical records. The system enables participating organizations to request clinical information about patients under their care from other local hospitals that may have provided treatment in the past.

The main challenges in developing the system were (1) determining how best to integrate information from disparate, heterogeneous data sources at the various hospitals using different database management systems, data models, and vocabularies, (2) linking records from patients who receive care from a variety of facilities, often under slightly different names, (3) overcoming organizational obstacles to the sharing of clinical data, and (4) developing ways of working with long-standing legacy systems. Instead of developing a centralized data depository containing records from all participating hospitals, the system designers instead chose to pull information from each of the hospitals' clinical information systems. Physicians use a Web browser to query the databases of participating hospitals so as to retrieve information about patients under their care. Responses from the individual hospitals are consolidated into a unified presentation that contains Web pages with information on demographics, problems, medications, allergies, notes, and visits. A key to the technical success of the W3EMRS architecture was the agreement by all participants on how to represent a core set of data (dubbed the "common medical record") in the same form within the HL7 data model.

Another major challenge in designing the system was addressing threats to security and confidentiality. System developers had difficulty identifying realistic threats to the system but realized that as long as some form of encryption was used across the communications links, it would be easier and cheaper for an attacker to break into the source sites to obtain patient information than to intercept

The average person will see a doctor as an outpatient three to four times a year, with Medicare patients making five to six visits a year and healthy adults making approximately two visits per year. Children visit doctors at the same rates as the elderly, but their records are smaller. Overall, 116 hospitalizations occur per year per 1,000 people, with 332 hospitalizations per year per 1,000 people for those 65 and older. Overall, 380 emergency room visits occur per year per 1,000 people. The timeliness required depends on whether access is needed in an emergency room situation or whether the record is to be transmitted to a specialist for an appointment at a later date.

Of greater concern is security. At issue are the confidentiality of clinical transactions—which tend to contain personal information and

and decrypt patient information as it traveled across the Internet. As a result, they decided that most of the security effort should be devoted to protecting the source computers from which the clinical data were extracted and the so-called agglutinator—which compiles records from multiple sites into a single file—rather than the communications links themselves.

System designers settled on a combination of organizational policies and technical solutions to protect data, establishing common confidentiality agreements among participating institutions and deploying technologies for authenticating system users, patients, and the source site.¹ The technical solutions recommended included secure (encrypted) communications and token-based authentication (in which the identity of users is validated by something they know, such as a password, and something they have, such as a physical device similar to a bank card). For example, in the CareWeb system, which is based on W3EMRS, authentication of the requestor's identity is verified by a user name, a personal identification number, and a code number generated by a hardware device that generates new codes every 60 seconds (the SecurID system from Security Dynamics, Cambridge, Massachusetts). Data transmissions across the Internet are encrypted, and digital signatures are used to confirm the veracity of the request. Audit trails are available at all participating institutions so that all accesses to health records can be investigated. To address privacy concerns, patients at each hospital are allowed to request that their health records not be shared with other institutions.

¹For information on the confidentiality agreements, see Rind et al. (1997).

SOURCES: Kohane et al. (1996); Isaac S. Kohane, director, Informatics Program, Boston Children's Hospital, presentation to the committee, September 13, 1998; and John Halamka, "The Value of Implementing a Secure Web-Based Medical Records Retrieval System," unpublished paper provided to the committee on December 17, 1999.

identities—and the integrity of the data transmitted. Confidentiality of the transmitted record can be addressed with encryption, as long as it is accompanied by strong methods for authenticating both the sender and receiver of the record. In pre-Internet (i.e., telephone or fax) communications, authentication is handled by the participants, who have some notions of how to verify that they have reached the correct party. These already shaky notions are ineffective in Internet communications. Internet protocols such as the hypertext transfer protocol (HTTP), which handles Web-based transactions, do have provisions for encrypting communications and validating users' identities, but effective mechanisms have yet to be deployed for distributing the tokens, certificates, or other technologies needed for authentication to all potential users of a system. Few organizations have adopted basic security procedures for protecting data integrity and data stored on computers that are accessible over the Internet.

System reliability is also an issue in clinical transactions. If health care organizations are to rely on Internet-based systems for access to patient records, then they must be assured that the systems will function properly 24 hours a day, 7 days a week. System outages of limited duration may be tolerable if records are being transmitted to a specialist or another institution in advance of a scheduled appointment, but they cannot be tolerated if access is needed on demand, such as for emergency room treatments.

A number of other challenges stand in the way of Internet use for clinical transactions. Perhaps the most daunting is the lack of agreement on data interchange standards and standardized vocabularies, or nomenclature, to describe clinical entities. Considerable support exists for the Health Level Seven (HL7) standards, which were developed for clinical transactions (Box 2.4),¹⁸ and the HL7 organization continues to improve the completeness of its HL7 data model to encompass more possible medical transactions (e.g., the current standard does not cover all drugs, X-ray studies, or nursing interventions or problems). Nevertheless, proprietary concerns and institutional inertia have led many vendors of clinical information systems and health care organizations to develop commercial applications and home-grown systems that are not compliant with the HL7 standard or idiosyncratically compliant with it. Furthermore, health care organizations have not been able to agree on a standardized vocabulary to use in describing different sets of clinical entities—despite significant support by the National Library of Medicine for multiple standardization efforts, including the construction and maintenance of a metathesaurus as part of the unified medical language system (UMLS).

There are many reasons why neither the data models, such as HL7,

BOX 2.4 **Health Level Seven**

Health Level Seven (HL7) is one of several organizations accredited by the American National Standards Institute that is developing standards for representing and communicating data related to health care. HL7 has focused on the clinical and administrative data generated within and across health care institutions. The data model and transactions standardized by HL7 have been widely adopted throughout the health care information industry for interchange between applications in a heterogeneous vendor environment. However, the data models used by most vendors remain proprietary and nonstandard. The HL7 data model (the most recent version, 2.3.1)¹ has been implemented in several ways: as a message-based format using an HL7-defined syntax, as a relational data model, and as an object-oriented data model. Most recently, the data model has been implemented within the framework of the Extensible Markup Language (XML) (World Wide Web Consortium, 1998). Because of the widespread adoption of XML as an interchange format for all kinds of Web commerce, this implementation is likely to have broad adoption.

¹See <http://www.hl7.org/library/standards_non1.htm#HL7> Version 2.3.1.

nor the vocabularies that are among the constituents of the UMLS became popular, but these are beyond the scope of this report. However, it is clear that one important reason is that there have been no sufficiently motivating arguments for data sharing across or even within institutions. Unlike billing transactions or pharmaceutical transactions, clinical transactions have only an indirect effect, at best, on the profitability of health care organizations. The health care industry is much less consolidated than the pharmaceutical industry, which has been more successful in deploying an interoperability standard. An additional inhibitor is the nature of clinical transactions, which tend to be more complex and varied than commercial transactions. The inertia with which clinical information systems have been deployed and accepted into practice has encouraged the development of consumer-driven health information systems in which third parties store and provide access to clinical data. Nonetheless, the functionality of these consumer information systems will be no better than that of their counterparts in health care institutions without the widespread adoption of standards for health data exchanges and the development of a robust means for authenticating users.

Use of the Extensible Markup Language (XML) for Internet-based

transactions may provide additional interoperability but could encounter similar barriers. XML is the universal format for structured documents and data on the Web. Like HTML, XML makes use of tags (words bracketed by “<” and “>”) and attributes (of the form name=“value”), but whereas HTML specifies what each tag and attribute means (and often how the text between them will look in a browser), XML uses the tags only to delimit pieces of data, leaving the interpretation of the data to the application that reads it. XML documents include an XML document type declaration, which contains or points to markup declarations that provide a grammar for a class of documents. This grammar is known as a document type definition (DTD). Therefore, before XML documents adhering to the HL7 data model can be created, a commonly accepted HL7 DTD must be ratified (Dolin et al., 1998). Without agreement on this common HL7 DTD, the exchange of clinical data across the Internet among health care systems will be significantly more cumbersome and will probably be further delayed.

New federal mandates are likely to encourage greater standardization that could facilitate use of the Internet for exchanges of clinical information. The Health Insurance Portability and Accountability Act (HIPAA, P.L. 104-191) requires the secretary of Health and Human Services to adopt standards for the electronic transmission of health information associated with the following transactions: health claims or equivalent encounter information, health claims attachments, enrollment and disenrollment in a health plan, eligibility for a health plan, health care payment and remittance advice, health plan premium payments, first report of injury, health claim status, and referral certification and authorization. The challenge will be to ensure timely compliance with HIPAA standards by the multitude of legacy applications—a task that will require at least as much effort as the Y2K remediation.

Furthermore, the emergence of Internet industries to host several of the aforementioned clinical transactions and their data repositories may provide the means for the widespread implementation of automated health transactions. However, hosting such applications has its own dangers. By slicing the space of clinical transactions into sharply demarcated segments (e.g. clinician documentation, laboratory reporting, medication ordering), there is a risk that important data relevant to patient care will become more dispersed and functionally unintegrated. Only by ensuring close adherence to HIPAA with a high degree of interoperability (i.e., adherence to data storage and communication standards) can this risk be abated.

Technical Requirements for Clinical Care

The technical capabilities required by clinical applications of the Internet are even more demanding than those required by consumer health applications because a number of factors converge. The need to protect the confidentiality of patient information is combined with the need for high bandwidth and low latency to support remote consultations; high availability is also required to ensure that patient records can be accessed when needed and that systems remain operational for the duration of a remote consultation. Although not all clinical applications of the Internet simultaneously stress each of these dimensions, the set of foreseeable clinical applications, taken as a whole, does.

Bandwidth ++++

Bandwidth requirements for clinical applications vary considerably, but many possible applications could demand high bandwidth. Remote consultations, for example, would require sufficient bandwidth for real-time video at rates approaching 1 Mbps for some types of diagnostic procedures. The transmission of large medical images could also require high bandwidth in some instances, to support the transfer of large numbers of images between an imaging center and a remote interpretation center or rapid turnaround of diagnoses from a remote specialist. Even remote access to electronic medical records could demand relatively high bandwidth to the extent that such records include images or video. In many cases, records (or images) could be downloaded in advance of the need to view them, although this technique would not be as effective in emergency situations.

Latency +++

Latencies across the Internet are adequate for many clinical applications, such as results reporting and downloading of most medical records (if these records are already stored remotely online), but applications like remote consultation would demand lower latencies to facilitate more natural interactions between participants. In virtual reality applications, low latencies are needed to create realistic spaces and interactions that are not distracting to users or participants.

Availability ++++

Because of safety and timeliness considerations in patient care, availability of the network is vital. Clinicians awaiting a lab result, radiologi-

cal examination, or connection to a patient's home cannot tolerate any unavailability of the network or the clinical applications running on it. Likewise, remote consultations will not be viable if network availability cannot be assured and connections are broken frequently. Neither care providers nor patients will tolerate delays, downtime, or lost connections in such applications.

Security +++++

Because patient information is so sensitive and its safety is paramount, the security of the network is vital for clinical applications. Without assurances that the confidentiality and integrity of patient data will be protected and that critical services will be available when needed, both the government and the public will resist the sharing of data across institutions. Security improvements will entail both technical measures (many of which would be deployed within specific applications) and the development of robust confidentiality policies to govern the disclosure of personal health information. Considerable technology is available for improving the security of many clinical transactions between established partners, but it is not widely deployed in health organizations (CSTB, 1997). More advanced security technologies, as outlined in Chapter 3, could improve protection, especially for data exchanges between unaffiliated health organizations and between consumers and care providers.

Ubiquity ++

Ubiquitous access will be important for many clinical applications if for no other reason than the dispersion of health care providers, many of whom continue to work in private practice or remote clinics. For remote consultations, distributed collaboration, and home care, the ubiquity of network services is essential. High-bandwidth (broadband) access could also be important for applications requiring the transmission of large images or real-time video.

FINANCIAL AND ADMINISTRATIVE TRANSACTIONS

The Internet is being evaluated as a medium for streamlining financial and administrative transactions in the U.S. health care system. Health care in the United States is financed largely by a network of so-called third-party payers—entities that insure and pay for health services but are not directly engaged in providing care. These entities range from government programs such as Medicare and Medicaid, which pay for the care of the elderly and impoverished, to private-sector organizations,

including traditional indemnity insurers, self-insured companies, and managed care organizations. The Internet could be used by providers to submit claims for payment or by individuals to enroll, disenroll, and change their coverage. Payers could quickly confirm an individual's eligibility for coverage and convey any changes to the health plans, which, in turn, could quickly relay the information to the person's designated providers. By accelerating these transactions, the Internet could reduce misunderstandings and disputes among parties, hasten payers' premium payments to plans and plans' payments to providers, and reduce administrative costs, which by some estimates constitute 30 percent of all health care expenditures in the United States. By one estimate, paper claims cost between \$2 and \$18 each to process, whereas electronic claims have costs measured in cents (McCormack, 2000)

Health care organizations have filed claims electronically for some time. Approximately 65 percent of the 4.7 billion claims submitted for payment by care providers and pharmacies in 1999 were submitted in electronic form. Hospitals and pharmacies have gone the farthest down this path, having submitted 85 percent and 89 percent of their claims, respectively, in electronic format in 1999, compared to just 43 percent for individual physician practices. Much of this progress is due to prompting from the Health Care Financing Administration (HCFA), which administers the Medicare and Medicaid programs, and from Blue Cross/Blue Shield organizations, each of which was expected to receive more than 80 percent of claims electronically in 1999. By contrast, HMOs and other commercial insurers, which together account for 44 percent of all health claims, were expected to receive just 18 percent and 45 percent of claims, respectively, in electronic format.¹⁹ Medicare has required its contracted managed care organizations to transmit electronically beneficiary enrollment, disenrollment, and correction data in batch mode to its data center (HCFA, 1999b). These managed care plans also obtain data on the disposition of their transactions and on plan membership and payments electronically.

Despite the trend to electronic formats, only a few provider organizations use the Internet to submit electronic claims or related transactions, and few payers are capable of accepting Internet-based transactions. Until recently, Medicare transactions, for example, have been conducted using telephone lines to access the data center. HCFA is replacing this system with the Medicare Data Communications Network, which will be accessed with Web browsers. HCFA's carriers for the conventional Medicare program, fiscal intermediaries, and most of the Medicare+Choice plans are using this network, and by July 1999, the remaining managed care plans were required to have made the transition. No such requirements apply to communications and data transfer between individual physi-

cians' offices and the HCFA carriers that process their claims. Medicare and other payers plan to increase the electronic transmission of data related to the quality of care and satisfaction of beneficiaries. Medicare, for example, foresees that its health plan management system will collect plans' quality-related data from the Health Plan Employer Data and Information Set (HEDIS) and the Consumer Assessment of Health Plans Survey (CAHPS) (HCFA, 1999a).

Private-sector organizations have also begun to experiment with Internet applications for financial and administrative transactions. Many such organizations see the Internet as a plausible means of achieving the long-held vision of seamless integration of information across organizations. Health organizations can assume that networking capabilities will be in place so they can concentrate their resources on higher-order functionality. The Internet may also make electronic claims submission practical for small group practices that cannot afford the hardware and staff needed for more conventional electronic data interchange (EDI) systems. The Regence Group of Seattle, Washington, for example, has developed a Web-based interface application called Network Data Express (NDEX) for determining beneficiary eligibility and making referrals. The system features claim status inquiries, provider directories, reference materials (such as the formulary), e-mail, and managed care data and reports. It processes about 20,000 transactions per month, doing the work of two to three full-time employees who would otherwise give the same information out by phone (see Appendix A for additional information on Regence's system).

Other efforts have been initiated at the state and regional level to promote health information exchanges. The Community Health Information Technology Alliance (CHITA) in Seattle, Washington, the Minnesota Health Data Institute, and the Affiliated Health Information Networks of New England, a project of the Massachusetts Health Data consortium, are three examples.²⁰ Such programs attempt to facilitate information exchange among a variety of organizations, including care providers, insurers, pharmacies and pharmaceutical benefits managers, accrediting organizations, and state health organizations. Considerable effort has been devoted to defining standards for data exchange and determining the types of data that must be exchanged for different transactions.

Security concerns have been a major impediment to greater sharing of information for payment and administration. Many such transactions—especially payment—contain sensitive information related to a particular patient's health, so their confidentiality must be maintained. Similarly strong requirements exist for data integrity. According to one information security officer interviewed as part of this project, the responsibilities of security officers in health care differ from those of their counterparts in

other industries: the applicable state and federal laws are different, and the privacy and security concerns are greater. At the same time, the health care industry is driven by economics, not privacy, so there is a need to balance cost effectiveness with security protections.

CHITA and the Foundation for Health Care Quality worked with the Massachusetts Health Data Consortium and Minnesota Health Data Institute on a three-state project focusing on electronic security. The goals were to determine how electronic security could be implemented affordably and to develop a business case for a community-wide, secure infrastructure for electronic business. The group worked with SAIC to develop a security and risk management plan for business-to-business health information networks. The plan identifies seven levels of increasing health care security, numbered 1 through 7. CHITA is working with Seattle-area health care organizations to implement level 6 security practices (HSL 6) within participating regional organizations.

HSL 6 supports remote access to a data repository but not direct remote access to the internal network of a health organization. It includes specifications for three network-based information services that are deemed essential to financial and administrative transactions: (1) authenticated, secure messaging, (2) authenticated, secure file escrow and transfer, and (3) authenticated, role-based access at the level of individual users. The security model has been developed and published (SAIC, 1998), and CHITA is in the process of identifying an organization that will function as a trusted intermediary to oversee a prototype implementation, followed by a wider pilot project in the region. Issues to be addressed include the identification of a certificate authority, which might be a non-profit organization, the state or federal government, or a private corporation.

HCFA is also a strong proponent of EDI but has prohibited the submission of payment information over the Internet owing to concerns about security and confidentiality. In January 1999, the agency revised its security policies to allow Internet-based transmission of information after it has been received from outside parties (HCFA, 1999a). HCFA is not allowing claims transmission over the Internet except by those Medicare contractors participating in an interoperability pilot of the HCFA Internet security policy. The pilot, which began in September 1999 and was scheduled to run through December 1999, tested e-mail, batch, real-time, and Web-based transmissions, while utilizing various authenticating and encryption technologies and including digital certificates through cooperating certificate authorities. Results of the pilot and accompanying recommendations were anticipated in February 2000. Depending on the findings and the cost-benefit analysis associated with the transmission of claims, HCFA will decide on whether or not and when to allow Internet

transmission on an operational basis.²¹ Acceptance of the Internet by HCFA could stimulate its wider use for submitting claims because HCFA processes a significant percentage of the nation's health care payments, and its acceptance could signal that such submissions can be handled securely.

Advances in the use of the Internet for financial and administrative transactions will be accelerated by the HIPAA of 1996 and by the 1997 Balanced Budget Act (P.L. 105-34). The HIPAA requires providers, provider organizations, payers, and clearinghouses to adopt uniform transaction standards, code sets, identifiers, and electronic signature standards for electronic transmissions of health care claims. In addition, HIPAA prescribes security standards for the protection of all electronic health information both within and between health care enterprises, and it gave Congress until August 1999 to pass comprehensive health privacy legislation. In the absence of such legislation, the secretary of Health and Human Services is to promulgate regulations to protect the privacy of personal health information (Harman, 1998) (see Chapter 5). Some health care organizations report that they have been slow to implement programs for Internet-based submissions of medical claims until the new regulations are in place. They fear that they may need to modify their systems after final regulations are promulgated or that Congress may pass legislation that supersedes them. The Mayo Foundation, for example, uses dial-up connections and leased T1 lines to submit claims electronically to government and commercial payers but will not implement an Internet-based system until security and privacy guidelines have been adopted by the federal government (McCormack, 2000).

A number of obstacles may further delay the widespread use of the Internet for financial and administrative transactions. While many large health care organizations are moving toward electronic billing and have Internet connections, many private practitioners (almost half of all physicians practice independently) lack Internet access and practice management systems for electronic billing. Furthermore, many practice management systems are not interoperable with the Internet, requiring users to download claims information into a separate Internet-based application, a process that adds unnecessary cost and complexity. Organizations large and small that have legacy systems operating for EDI and claims processing may also be slow to replace those systems with Internet-compatible systems, although in the long run they will need to modernize their systems. A lack of standards for electronic claims will continue to impede efforts at Internet-based exchanges, as many practice management systems use different formats for data, and payers often cannot accept data in multiple formats. Continued standards efforts, including those mandated by HIPAA and those under way in other regional collaborations, may

ease this concern, but the decentralized nature of the health care industry presents a significant impediment.

Technical Requirements for Financial and Administrative Applications

Bandwidth +

The bandwidth requirement for most financial and administrative transactions is modest. Most transactions consist of short, text-based messages. In some cases payers may request care providers to transmit large diagnostic images in support of claims for payment, but such images need not be delivered rapidly, so even in these cases, bandwidth is not a significant consideration.

Latency +

Latency is also not a significant factor in financial and administrative transactions. In some cases, such as checking on the terms of a patient's coverage for certain procedures and eligibility for reimbursement, timely responses are desirable. If an organization's servers are far too small and the system response time gets long, there could be problems, but those would be problems not with the Internet itself but with the individual nodes on the network. For many transactions, such as submission of claims for payment, latency is not an issue at all as the payment process tends to be slow. Granted, improvements in technology and processes could eventually allow for near-real-time review of claims and electronic payment for procedures, but response times and latency would not be a driving consideration in such systems.

Availability +++

The importance of availability in financial and administrative transactions is relatively high, depending on the specific use. The routine uses of payers require no more than a moderate level of availability. System outages could be compensated for by waiting to send a payment request. However, the preapproval of immediate care for beneficiaries—for example, the approval of pharmacy claims—would demand greater availability since there is greater need for a rapid response. Care providers and payers are unlikely to use the Internet for such transactions if it is not reliable.

Security +++++

Security of Internet communication of health-related data is a *sine qua non* of the Internet's greater use for financial and administrative transactions. Many of these transactions contain sensitive, personal information on the types of health care services that were provided to a particular patient or the diagnosis of a condition. Hence, they are almost as sensitive as clinical records. Both care provider organizations and payers also have strong incentives to demand that data integrity be maintained, to ensure that information is not corrupted during transit or when stored in computers attached to the network. As with clinical transactions, providing such security entails both technological mechanisms and confidentiality policies that govern disclosures of information by health organizations. Care providers and payers with established relationships can make use of existing technology for securing information during both transmission and storage, but more advanced technologies for authenticating users would enable transfers of information among a larger number of payers and providers.

Ubiquity ++

Requirements for ubiquity of access would be high if the objective is that all providers, including individual physician offices, routinely submit claims and quality improvement information, eligibility checks, and other information via the Internet. In the short term, claims administrators could use the Internet for communications with institutional providers, reducing the degree of ubiquity required.

PUBLIC HEALTH

Public health workers promote health and the quality of life by preventing and controlling the spread of disease, injury, and disability. Public health officials collect statistics on the occurrence of diseases, disseminate guidelines to health care practitioners and the public, fund research on ways to improve public health, and deliver health care to underserved populations. A number of these activities could be enhanced by an Internet that is better attuned to public health needs, that provides sufficient security to protect sensitive medical records, that is accessible to all public health workers and the public at large, and that remains operational even in times of natural or man-made disasters. Public health surveillance, in particular, stands to benefit from Internet-based transactions to assist in collecting data about the health of individuals, personal risk factors, and medical treatments, as well as data about potential

sources of disease and injury in the environment and resources that can be used to take effective action.

In recent years, attention has turned to making certain that public health officials at local, state, and regional levels have adequate connectivity and expertise to use the Internet for their work. Several reports have reflected on the need for new relationships and better collaboration between public health officials and individual health care providers.²² The National Library of Medicine (NLM), in conjunction with several other public health organizations, has initiated a program, Partners in Information Access, designed specifically to help public health officials gain access to the Internet and to relevant health information.²³ Since October 1998, 20 awards totaling just under \$1 million have been made for programs in 20 states. The goals of this program are fourfold: (1) to increase public health professionals' awareness of the services of the NLM, the Centers for Disease Control and Prevention (CDC), and the National Network of Libraries of Medicine (NN/LM), (2) to assist public health professionals in getting connected to the Internet, (3) to train public health officials in the use of information technology and information services, and (4) to increase awareness of public health information needs and resources among NN/LM members. Individual projects will attempt to provide modems and connections to Internet service providers for public health departments lacking such capabilities; support access to public health information and related biomedical topics via local medical libraries; survey the information needs of public health officials; and train public health officials to use the Internet and specific information resources, such as PubMed and CDC WONDER.

These efforts reflect the growing awareness of the linkages between public health and the care of individuals. Recent changes in the ecology and epidemiology of disease and the organization of health care delivery systems have led to a convergence of these two components of health care. For example, the AIDS pandemic forced many to realize that high socioeconomic status did not confer immunity from epidemic infectious diseases. Second, it has become increasingly clear that the major controllable causes of disease involve the traditional interests of public health: smoking, alcohol and drug abuse; injuries; and nutritional problems, including obesity. Finally, the advent of managed care and capitation has made payers responsible for protecting the health of populations. This convergence of public and private health interests represents a historical opportunity to bring public health thinking into the daily practice of medicine. It makes public health surveillance a more compelling application of the Internet. By some estimates, only about 10 percent of all early deaths in the United States can be prevented by medical intervention; population-based approaches could prevent up to 70 percent of them by

targeting underlying risks such as tobacco, drug and alcohol abuse, diet and sedentary lifestyles, and environmental, occupational, and infectious risk factors (McGinnis and Foege, 1993).

Public Health Surveillance

The public health system in the United States is hierarchically organized around community (city, county, or other local jurisdiction), state, and federal efforts. Each of these jurisdictions is chartered to collect different sorts of data and share them in different ways. The federal public health centers must recognize large-scale trends in the occurrence of disease and allocate resources to minimize the damage to the public health. Community public health offices must process information about individual patients and local outbreaks in order to recognize and respond to the needs of the community. For historical reasons, the three levels of public health monitoring and surveillance have developed very different organizations and communication mechanisms. However, there are fairly well defined communication points where the systems interact with one another. For example, physicians and medical laboratories must report the occurrence of certain conditions to local health departments, depending upon the reporting requirements. Certain conditions must also be reported to state public health offices, which in turn file reports with the federal CDC. Although the sets of data reported to CDC are uniform across the states (and updated regularly), each state and county health department can require that any condition it deems significant must be reported (rural counties, for example, have different interests from urban ones).

An important mechanism for collecting information of great significance to public health—and one that is ripe for the Internet—is automatic reporting by medical laboratories of test results for some communicable diseases, such as tuberculosis. Such systems promise both to improve reporting of adverse events and to lower the costs of collecting and maintaining such data. Testing laboratories are required to report certain diagnoses to their local health offices so that public health officials may ensure that adequate treatment is delivered and that spread of disease is contained. Currently, most such reporting is done on paper, with laboratory results being sent by mail or fax to the public health office.²⁴ Officials from the county public health office then follow up with the local physician and/or patient to investigate possible causes of the condition, paths of contagion, and needed interventions. This reporting system is fraught with errors and delays, as reports are transmitted in a range of forms (mail, fax, etc.), following different sets of rules, to different county offices. Not surprisingly, reports are often incomplete or are sent to the wrong

county health department because testing laboratories cannot determine in which county the patient resides.

The Internet offers a way to streamline this system, ensuring that reports are sent in a timely manner to the correct local public health office and to the state for analysis. The Washington State Department of Health, for example, has begun to develop its Electronic Laboratory Reporting System, to support Internet-based submission and notification of cases of reportable diseases within the state, which total about 100,000 annually. The system uses the Internet to allow testing laboratories to report conditions directly to the Department of Health, which forwards the information to the appropriate local health department. It is intended to hasten the filing of reports, reduce the burden of reporting for laboratories and health agencies, improve the state's ability to track disease outbreaks that cross county lines, and ensure that reports are transmitted to the correct county health office. The system takes advantage of the broad reach of the Internet to establish connectivity among the health departments and private testing laboratories. In preliminary tests with the Group Health Cooperative of Puget Sound, the system improved the rate of reporting of health conditions at both the state and local levels, especially for smaller counties whose paper-based reports were more prone to be lost or misdirected. The time to file a report with the counties improved moderately—to less than one day—while the time to transmit reports to the state improved dramatically—from a mean of 40 days with a paper-based system to just a day with the Internet-based system.²⁵

The Internet also offers unprecedented opportunities for planning and resource allocation at the community, state, and federal levels, potentially improving care and reducing costs. Especially in a setting of limited resources, mechanisms for identifying the need for resources and deploying them rapidly to affected populations are of critical importance. Automated systems for tracking the outbreak of diseases both acutely (on the scale of hours) and subacutely (on the scale of days to a week) would allow for dynamic allocation of resources, such as medicine, non-pharmaceutical medical supplies, donated organs, blood products, and even medical personnel, based on needs. Consider an outbreak of illness caused by a pathogenic bacteria contaminating hamburgers sold by a chain restaurant. Early detection of such an outbreak could lead to rapid notification of local and state public health officials so they could begin to track down the source of the infection. At the same time, pharmaceutical suppliers could be notified that an extra supply of certain types of antibiotics or rehydration intravenous fluids would be required in the region. Finally, hospital personnel could be alerted to the fact that these cases were appearing and could be briefed on the signs and symptoms to make them more prepared for emergency room visits related to the out-

break. At the federal level, information about these outbreaks could contribute to decisions on the cost-effectiveness of setting up new regulations, their enforcement, or their propagation and dissemination within the health enterprise. Clearly, application software must be developed to assist decision makers in allocating resources and in identifying and responding to trends in disease, but the Internet would provide the infrastructure necessary to gather the data upon which these decisions will be based.

Integrating Data Sources for Improved Decision Making

By allowing automated queries to disparate databases, the Internet could also help public health officials better integrate the available data to improve data analysis and health monitoring. Currently, a number of political and bureaucratic boundaries impede the use of the Internet for public health purposes. Most importantly, federal and state public health agencies are organized in vertically integrated, disease-specific systems. One rationale for this structure is that vertically integrated data and communications systems best serve the traditional public health functions for a given disease.²⁶ Thus, dozens of systems support individual diseases (such as AIDS) or disease groups (e.g., hospital-acquired—nosocomial—infections). The result is massive duplication, and a patient's clinical information could reside in several different systems that do not interconnect. The Internet could be a powerful technical tool (and political motivator) to realign these programs and allow better integration of data for monitoring public health. Doing so would require that public health offices and their databases be connected to the Internet and that mechanisms be put in place for protecting the security and confidentiality of data that contains personally identifiable health information.

Beyond integrating databases within the public health sector, the Internet offers the opportunity for public health officials to collect data from private sources that might be important in their surveillance efforts. School attendance records and sales of prescription drugs or nonprescription remedies could signal the outbreak of a disease in its early stages, before symptoms reach the level at which people visit a doctor. Indeed, the New York City Department of Public Health arranged to receive such data from one local drugstore chain to improve its surveillance activities, recognizing that abnormal sales of antidiarrheal medicines could indicate a wide-ranging but low-level epidemic of food poisoning or problems with the water supply. Being able to access such information quickly through the Internet could allow health care providers to respond rapidly to disease clusters and reduce the exposure of the population to disease. Much of this information is available today in electronic format, and with

proper protections for proprietary and confidential information, it could be made available to public health officials via the Internet.

Responding to Bioterrorist Attacks

How to detect and respond to a bioterrorist attack (e.g., an intentional release of poisonous gases or tainting of the public water supply) has become a growing concern for the public health community. The use of biological weapons by terrorists—even an individual terrorist acting alone—could inflict life-threatening illnesses on a large scale and, unlike explosions or chemical releases, could easily escape immediate notice. Many biological agents would not produce symptoms in their victims for days, weeks, or longer, and initial reports of illnesses might not appear unusual, delaying recognition of a widespread problem.

In the case of bioterrorist attack, each of the phases of the public health process would depend on a successful infrastructure: recognizing a trend, identifying the cause of the trend, formulating a strategy for responding to it, allocating resources for the response, deploying the response, and monitoring its success. Initial clinical reports, which might come from doctors' offices and emergency rooms over a large area, would need to be aggregated at a high enough level for a geographical pattern to emerge and a problem to be detected. Local public health officials in the affected areas would need to confer with one another to plan a coherent response to the attack and allocate resources to address immediate medical needs. Data would need to be provided to public health teams charged with identifying the pathogen and formulating and implementing a response. The ability to keep information from the public in order to avoid panic could also be important, depending on the situation.

The CDC found in a 1998 study that most local health departments lacked the capabilities to adequately detect and respond to a report of bioterrorism. It found that most such departments lack basic information and communications systems and cannot communicate reliably with CDC, state health departments, or emergency response agencies in a crisis. Half lacked Internet access, 20 percent lacked suitable computer capacity, and 70 percent lacked training in the use of electronic information technologies for conventional health purposes (CDC, 1998).

To remedy this problem, CDC is developing a national Health Alert Network that will facilitate the collection of information from testing laboratories, the sharing of information among public health officials, and consultations among them regarding needed responses.²⁷ A total of \$28 million was allocated to this task in FY99.²⁸ The network will use desktop personal computers and laptops connected to the Internet with sufficient bandwidth to handle the transfer of laboratory reports, interac-

tive collaboration among public health officials, and multimedia distance training. It will make use of public key encryption for secure communications and authentication. Because of its critical nature and the need for its continuous availability, the network will be designed with sufficient redundancy to provide backup operations in case of a link failure and disaster recovery plans to allow rapid restoration of service in case of other component failures.²⁹ Videoconferencing capabilities are seen as important, for they would allow public health officials to communicate more effectively during a crisis than they could with either text or audio alone. Mechanisms may be needed to accommodate (possibly by diversion) high volumes of traffic in an emergency.

Technical Requirements for Public Health Applications

Use of the Internet for public health surveillance would require technical advances in a number of areas. Of primary interest are ubiquity and security, but availability is also of concern. Other technical parameters, such as bandwidth and latency, are less important in most public health applications, although the desire for videoconferencing in widespread emergencies would increase the need for bandwidth and for low latencies to support real-time, interactive video. Solutions to these technical problems could greatly expand the use of the Internet in support of public health.

Bandwidth +

In general, the information transmitted for the purposes of public health requires relatively low bandwidth. Public health data rarely involve images or other large data objects, although videoconferencing among public health officials would require higher bandwidth from at least some computers and locations. Of course, there is also the potential for many data objects to be transmitted through the network, raising the bandwidth requirement by virtue of aggregated traffic levels rather than large individual files.

Latency +

Few of the applications of the Internet in public health are sensitive to small delays (i.e., of seconds to minutes) in the transmission of data, so that latency is less important.

Availability +++

For public health, the availability of the network is of moderate importance. Although short downtimes can normally be tolerated, the minute-to-minute monitoring of outbreaks of acute disease (especially in the case of bioterrorism) would not tolerate extended periods of network failure. If the Internet were to be used for detecting bioterrorist attacks, it would have to be reliable and resistant to hostile attacks (which could accompany a bioterrorist attack). Because data collection and aggregation take place continuously, loss of network might lead to loss of data and failure to respond in a timely fashion.

Security +++

The security of data on the Internet is of paramount importance to public health applications. Data reported by testing laboratories contain identifying information that is used by public health officials to map diseases and conduct interviews with affected patients. The public health system depends on the public's trust that sensitive health data are being used for the benefit of the public only. Such data must be protected both in transit and while stored in computers in public health offices. Sharing public health information at the community, state, and federal levels requires the development of advanced technologies for intelligently stripping data of identifiers so that personal identities cannot be reconstructed from the data. Although certain local public health functions (treatment and prevention of tuberculosis, for example) require knowledge of the patient and his or her home situation, it becomes less necessary to have identifying information at the state and federal levels, where general trends are of interest. Even with these technologies, it is critical to have technologies for authenticating data and users. Also important is the ability to protect sensitive institutional data and sensitive information relating to bioterrorist attacks. Such protection would require policies to determine who may access which data, as well as technologies to protect the confidentiality of the information and its integrity.

Ubiquity ++

The success of the public health system requires that reporting and surveillance networks have widespread connectivity that includes local (e.g., community) health departments, testing laboratories, and the provider organizations that order the tests. To serve the entire nation in a cost-effective, standardized way, it is critical that the public health information infrastructure extend to every community, state, and federal public

health agency. Information gaps would be a great burden on the nation, since they would require creation of a secondary, mostly redundant mechanism for data collection and dissemination. As in many other areas of information technology, a few exceptions threaten to make the entire enterprise too expensive. At the same time, the benefits of a ubiquitous network to community, state, and federal agencies would be substantial and would probably improve public health greatly. Public health organizations often run on tight budgets, so the cost of access to networking technologies must be reasonable. In any event, fewer distinct entities would probably need to be connected for public health applications than for consumer health, clinical care, or financial and administrative applications.

PROFESSIONAL EDUCATION

Despite advances in technology and the Internet, the education of health professionals is practiced much the way it has been for decades. Students of medicine, nursing, pharmacy, and allied health disciplines set out on a course of graduate and postgraduate education, with much of this training occurring in classrooms or lecture halls. The emergence of the Internet and Internet-based technologies has the potential to transform health professional education at all levels. Educational systems that were once teacher-centered and geographically limited can now become learner-centered and unconstrained by geography. If the Internet is to support this transformation, the demands on it will be substantial.

Graduate Education

Graduate education is provided by 124 accredited four-year medical schools in the United States in two phases: basic science education and clinical education. Basic science courses, such as anatomy, physiology, and pharmacology, are taught in a traditional lecture format supplemented by reading and hands-on laboratory sessions. Significant challenges exist in providing basic science education, including the large amount of information that needs to be transmitted, the fast pace of change in the information base, and a lack of tools that would allow students to index what they learn and to retrieve it later in their training. In contrast, clinical education uses different methods. Knowledge about the diagnosis, treatment, and care of patients is transmitted mostly using an apprenticeship model, whereby the student learns from taking care of patients under the guidance of more senior clinicians.

Significant efforts have been made over the last decade to make basic science education less didactic and more problem-oriented. These efforts

have led to new teaching methods and materials, some of which use computers and the Internet, and new courseware. The advent of online textbooks, journals, and interactive courseware shared across institutions could accelerate this trend so that students spend less time reading books and attending lectures and more time researching topics online. Another trend in basic science education is the use of sophisticated simulations to demonstrate anatomical or physiologic concepts. Such simulations are three-dimensional, color representations that can be rotated or otherwise manipulated. The bandwidth requirement for these applications is high, straining local networks, especially local access connections to students' homes. There are other significant barriers to the routine use of computers and the Internet in basic science education. Not all students have computers, and few campuses have network connections that allow them to gain access to the Internet from classrooms, libraries, or other campus facilities. Networking bandwidth and servers within the institution cannot always handle the dozens of students trying to access the same resources at the same time.

The Internet can also reshape clinical education to overcome some limitations of the apprenticeship model. First, supervising clinicians may not themselves be up-to-date on certain issues and thus are not always the best source of information on these issues. Second, the location of clinical education can limit the student's exposure to certain types of patients and diseases. For instance, certain infectious disease such as tuberculosis and AIDS are seen more frequently in urban hospitals. Students who do their clinical education in a rural setting might not be properly equipped to deal with such patients if they later practice in an urban setting. Structural changes in the health care industry may also serve to limit the diversity of health problems students gain exposure to during their clinical education. The rise of HMOs and specialty clinics makes it far harder for an intern to see a reasonable variety of patients and diseases just by working in a hospital. One effect of the rise of HMOs has been to shift the locus of care away from hospitals and toward local clinics and outpatient facilities. As a result, many of the patients interns see in a hospital setting have already been diagnosed in one of these other facilities, which will tend to limit the interns' experience.

Computer-based tools and the Internet can complement apprenticeship-based clinical education. Perhaps the best examples of such tools are those that allow students and clinicians to search and retrieve the latest medical literature over the Internet and use the evidence retrieved to guide clinical decisions. The process of incorporating knowledge from the medical literature into patient care decisions is referred to as evidence-based practice. Today, students can use the Internet to search MEDLINE, the bibliographic database containing millions of citations to the bio-

medical literature, to read abstracts of journal articles, and—in some cases—to download electronic versions of the original journal article. Similarly, new Internet-based systems are emerging that allow search and retrieval of textbooks, drug information, medical news, and patient education material. The impact of this trend on local networks and the Internet could be substantial. Assuming that some 70,000 medical students, 100,000 medical residents, and 150,000 students in allied health sciences (e.g., nursing, dentistry, pharmacy, public health) in the United States will regularly be accessing textbooks, journals, and other educational material from centralized repositories on the Internet, traffic could increase substantially not only on the Internet but on the LANs of medical education institutions.³⁰ Although the item bandwidth required to transmit a single data element might be low (e.g., 100 kB for an HTML journal article with graphics), the total bandwidth requirements could be much larger owing to the large number of users and the high frequency of usage. This may call for a better understanding of network management and of the trade-offs between increased network capacity and the local caching of data or its replication on other sites to reduce bandwidth needs.

The trend to evidence-based practice will probably continue. Clinical students will increasingly be expected to support their patient care plans with evidence from the medical literature. To do this, they will become even heavier users of online literature retrieval systems such as MEDLINE and electronic journals. Although bandwidth will become an even larger network issue because more users than ever will be using these resources more frequently, the bandwidth needs will not generally be as great as the bandwidth needs to support real-time video streams for other applications.

Another trend in clinical education will be to community-based education. Students who trained mainly in academic hospitals will spend more time training in community hospitals and rural clinics. This trend could be accelerated by Internet links between remote areas and academic medical centers. Using such links, students can discuss cases with preceptors (using audio- and videoconferencing), share clinical experiences with fellow students, and download educational material from university and other Web sites. Internet access in all small community and rural health settings would be important for the success of such communication.

The Internet could also allow clinical students to take greater advantage of simulations to learn about diseases and situations they would not otherwise encounter during their training years. These simulations could take the form of interactive, multimedia modules retrieved over the Internet at the time of need. Modules could include high-resolution graphics and images, streaming audio and video, and text. Similar multimedia content would be required for simulations that test student knowledge for purposes such as allowing advancement through the school

curriculum and granting a license to practice. Such simulations will require that the Internet and local networks have adequate bandwidth and, for interactive simulations, low latency. Already researchers are working on systems to allow the simulation of surgical techniques. These simulations combine three-dimensional imagery with haptic feedback that recreates the touch and feel of live surgery. Such systems require extremely low latency, on the order of a few hundred milliseconds per round-trip, to prevent users from perceiving an unnatural lag between the time they take an action and sense a response (Table 2.5). They also require the elimination of mismatches between different data sources: visual, audio, and haptic information need to be properly synchronized for a user to properly experience a virtual surgical system.

Continuing Education

Once clinicians are in practice, they are essentially on their own to keep their knowledge and skills up-to-date. They do so informally by reading journals and textbooks, by interacting with consultants, and by talking with peers. A formal process, designed to maintain and enhance clinician knowledge and skills, also exists and is referred to as continuing education (CE). CE credits are not a national requirement but are required by some states and subspecialty boards for licensure and board certification. For instance, 28 states require physicians to meet minimum CE requirements for licensure, and 9 specialties require it for board certification (AMA, 1996). The CE requirements vary, but they usually call for completing 150 hours of courses over a three-year period. Three of the eleven states with the largest concentration of physicians have no CE requirements (New York, New Jersey, and Illinois).

Traditional CE consists of a time-based system of credits that are awarded for attending conferences, workshops, or lectures. Typical CE

TABLE 2.5 Effect of Latency on Interactive Simulations

Round-trip Latency (msec)	Effect
25	User cannot detect a problem.
50	User detects a problem but cannot identify it as latency.
100	User recognizes a latency problem and can compensate for it.
200	User recognizes a latency problem but cannot compensate for it.

SOURCE: Colonel Richard Satava, Yale University, presentation to the committee on March 1, 1999, Washington, D.C.

courses are teacher-initiated, use passive educational models such as lecture, and are often sponsored by the health care industry. Systematic reviews of CE interventions have shown that traditional CE—short courses, conferences, and seminars—are largely ineffective in improving knowledge or health care outcomes (Davis et al., 1995). Two newer approaches, academic detailing (targeted visits by physician educators such as pharmacists) and computerized reminders, have, on the other hand, been found to have a positive effect on knowledge and outcomes. The general success of interventions such as computerized reminders suggests that knowledge delivered in the context of daily patient care and for the purpose of assisting in problem solving is where CE should focus in the future. If this suggestion is acted on, the Internet and systems that integrate patient data with general medical knowledge will probably play a central role in transforming postgraduate education.

The main trend in postgraduate education will be continuous (as opposed to continuing) education. Instead of being concentrated in a week's worth of off-site conferences, education will be provided using multiple modalities available at different times during daily practice. Although traditional CE classrooms and conferences will still exist, virtual conferences will become more common. Using the Internet, clinicians will be able to choose from libraries of video and audio lectures, interactive courseware, and live discussions among colleagues from around the world. For the educational tools that are not live, clinicians will have great flexibility where and when they use them. CE credits, once awarded for sitting in lectures, will be awarded based on time spent and information learned using these online resources. To make virtual conferences a reality for all practicing clinicians, the clinicians will need high-speed Internet access from their health care sites and from their homes.

In addition to learning in virtual conferences, clinicians will do much of their learning of new diagnostic and therapeutic measures in the context of daily patient care. This new learning modality will be fueled by two converging trends: (1) the emergence of patient records in electronic form and (2) the availability of medical literature over networks. In this new modality, behind every abnormal test result, unfamiliar diagnosis, or new drug in the electronic medical record will be a link to the best available knowledge on that topic. Instead of having to initiate a search for information when a question arises, the answer will be anticipated and a link to the answer created within the patient record. As they use this up-to-date knowledge at the point of need, clinicians will also be able to fulfill CE requirements, because the time they spend using the resources and the effect the knowledge has on the patient care process will be logged and reported automatically.³¹ Of course, such capabilities also raise issues of privacy. Will care providers be able to peruse outside information

sources and pursue learning opportunities without being monitored? Will the use of such resources be viewed positively (e.g., the provider is trying to expand his or her knowledge) or negatively (e.g., the provider does not understand some new procedure or diagnostic method)? Such issues will need to be addressed in order to ensure acceptance of these technologies (see Chapter 3 for a discussion of technologies to protect online anonymity). In this vision of integrated patient data and knowledge sources, computer networks will play a vital role. Because of the time-critical nature of the knowledge delivery during the patient care process, reliability of the network and information servers will be vital. Because queries posted to knowledge sources will be based on patient characteristics, security of the network will also be important.

Technical Requirements for Health Professional Education

Bandwidth +++

The bandwidth requirements for health professional education are moderately high. Whether large numbers of people frequently use low-bandwidth applications, such as literature searching, or infrequently use high-bandwidth applications, such as teleconferencing or simulations, bandwidth will be important and sometimes a limiting factor. The development of virtual classrooms and interactive surgical simulations could drive bandwidth requirements even higher.

Latency +++

In general, applications to support health professional education do not require instantaneous delivery, and so the latency requirements of the Internet are not great. However, interactive simulations (such as those for teaching surgical techniques) and conferences would suffer from long latencies.

Availability ++

In general, the availability of the network for health professional education is of moderate importance. Many educational activities are not as time-critical as patient care activities and can tolerate low-level data losses or occasional unavailability. However, as the Internet becomes more and more of a tool for education, as students and instructors come to rely on it more as a communication medium, and as education becomes more integrated with patient care activities, the need for availability will increase. As computer-generated reminders and links to external resources become

more closely integrated into medical records, availability will become more important.

Security +

For the most part, health professional education is based on public domain information, so the security requirements for the network are not great. However, with the interplay between patient data and medical knowledge required to support new modes of education, security will become critical. Tools for protecting anonymity may also become important to the extent that clinicians want to be able to consult online resources anonymously.

Ubiquity +++

For health professional education, the ubiquity of the network is of great importance. Improvements in clinician knowledge and clinical outcomes will partially depend on the extensive deployment of new learning techniques and tools. Without access to the Internet from potential sites of care delivery and from their homes, clinicians practicing in poor or remote areas will not be able to benefit from these new capabilities. As a larger number of medical students do internships in remote locations, the need for access will also increase.

BIOMEDICAL RESEARCH

Biomedical research attempts to understand the mechanisms underlying human health and disease. It ranges from basic investigations of the molecular details of biological systems to the study of clinical implications of new scientific findings. In basic biology, the work tends to focus on (1) the biological sequencing of DNA and proteins, (2) the three-dimensional structures of anatomical parts and biochemical molecules, and (3) the determination of metabolic pathways. Progress in biomedical research has recently been fueled by an explosion of biological data available for analysis, as evidenced by the growth in the number of DNA bases (chemical units) that have been sequenced, from next to none in 1982 to in excess of 3 billion in 1999. The Internet has been widely accepted within the biomedical community and greatly facilitates the research enterprise by helping integrate disparate databases for improved analysis, allowing linked simulations, and enabling remote control of biomedical research apparatus. Each of these applications poses a range of technological challenges.

Biomedical Databases

The most important reason for the adoption of Internet technologies within the biomedical community has been the development of publicly available databases containing biological information. Many major biological databases are available at no charge on the Web and offer rapid access and query capability (Table 2.6), and research laboratories are beginning to release primary data onto their Web sites so colleagues can use them for reanalysis or testing new hypotheses. Some databases are extremely popular: each day some 600,000 searches are run from 120,000 different addresses against PubMed, a Web-based service hosted by the National Center for Biotechnology Information (NCBI) with abstracts and some full articles from MEDLINE plus additional journals in the life sciences. These figures grew at an annual rate of 50 percent over the last 3 years.³²

Such high rates of use create a number of problems for the host sites. The network bandwidth required for any individual database request may be small (a few kilobytes of data per query), but the aggregate effect of this traffic on bandwidth going into the database server can be significant, overwhelming capacity. Because databases such as MEDLINE and GENBANK are intended to serve a multitude of users in a timely fashion, they are designed for individuals to use in an episodic manner; they cannot routinely allow companies or institutions to perform many queries in a short period of time, such as to run automated queries of a program that is searching systematically through the literature as part of some data-mining application. Commercial online databases have similar problems (Box 2.5). For this reason, many users prefer to obtain a local copy of the databases so that they can subject them to high levels of use without

TABLE 2.6 Examples of Online Databases of Interest to the Biomedical Research Community

Database	Content
MEDLINE	Index of the entire biomedical serial literature. Contains basic reference information since 1966 and limited full-text reference information from the National Library of Medicine. Contains more than 9,000,000 references from a list of 3,900 periodicals.
GENBANK	Database of DNA sequences.
SWISS-PROT and PIR	Database of protein sequences.
Protein Data Base	Three-dimensional macromolecular structures.
OMIM	Information about human genetic diseases.

BOX 2.5
Bandwidth Concerns of a Commercial Content Provider

Ovid Technologies, Inc., is an aggregator of information on science, technology, and medicine. The company provides centralized access to resources such as bibliographic databases and full text journals via individual CD-ROMs, servers located on local area networks, and the Internet. The entire Ovid database contains about 200 gigabytes (GB) of text and images, about 5 to 10 GB of which is updated each month. On average, each subscriber (there might be 100 to 250 individual users) downloads 300 MB of data per month, which translates to 70 kB of data per individual per day. Some large academic institutions have 100 concurrent users who download 30 GB of data per month. Usage rates triple every year.

Ovid Online is based in Utah and is connected to the public Internet via three T1 lines. This bandwidth is sufficient for most North American users, but some 13 large universities, consortia, and corporations have established dedicated connections to the database (typically via frame relay connections—see Chapter 3) with speeds of 56 to 256 kilobits per second (kbps). These dedicated connections also allow pharmaceutical companies to prevent competitors from finding out what information they are seeking. Other users have established hybrid connections in which certain databases (e.g., MEDLINE) are loaded onto a local server and updated weekly, but other large databases (e.g., those containing full text journals) are accessed via the Internet or a dedicated line.

Because of inadequacies in Internet bandwidth to locations overseas, the company has established additional servers in the United Kingdom, Sweden, Japan, Hong Kong, and Australia onto which Ovid products have been loaded via CD-ROM. The company recently contracted with a service provider that will provide guaranteed bandwidth and quality of service for transoceanic downloading of updates. Updating the entire database demands kilobytes per second of bandwidth, but it is difficult and expensive to update six to eight regions of the world simultaneously. Ovid would like more bandwidth and improved quality of service in order to conduct network-based updates. Alternatively, it could redesign its system to send just the 5 to 10 GB of changed information. The trade-off is a matter of cost as much as of technical capability.

SOURCE: Based on a presentation by William Detmer, vice president, Ovid Technologies, Inc., to the committee on March 2, 1999, Washington, D.C.

monopolizing public resources.³³ Doing so also allows companies to use the databases without fear that their searches will be watched by competitors. Knowing the kinds of information that companies are searching for can yield clues about the projects they are working on. Local replication would be unnecessary if trusted security services were available across the Internet that could guarantee that queries and results from a Web site remained anonymous and confidential or if the servers could support

individual use as well as heavy automatic use by computer programs employing data-mining techniques.

While increasing the bandwidth into biomedical research databases is one way of alleviating bottlenecks, the rate-limiting factor in some systems is the computational server, not the communications bandwidth. The NCBI, for example, uses a T3 line for connectivity to the Internet, which provides it with 45 Mbps of bandwidth. As of mid-1999, NCBI was utilizing only about one-third of that capacity.³⁴ The computational server has become the bottleneck, because NCBI is receiving more requests to compare large data sets with one another and with data sets provided by users. The volume of requests for large data set comparisons is still small, but NCBI has developed some governors to limit requests from particular sites so that other users can access the system.

Network limitations also pose difficulties for users of biomedical databases. Some algorithms that act upon databases require that every single element of the database be compared with every other element. Thus, if there are 1,000,000 entries in the database, then 1,000,000,000,000 possible comparisons must be computed, requiring very fast computation if solutions are to be found in reasonable amounts of time. In many cases, investigators transfer their data sets to remote supercomputing sites (such as the sites sponsored by the National Science Foundation in San Diego and Illinois) so that the processing will not be slowed by data transfer rates over the Internet. Such remote processing has limitations, especially in providing real-time feedback to the researchers.

A data source commonly requiring extensive computational analysis is the output of high-resolution imaging devices. High-resolution images containing millions of elements cannot be transferred rapidly enough to allow researchers to manipulate them in real time. Thus, accurate visualization requires that the image be rendered by a local computer with sufficient computational capabilities or by the efficient transfer of information from an image server to a display device on the Internet. Furthermore, the size of the databases makes transfers slow, and replicating them makes it hard to keep them current during a computation.

This process can also be facilitated by improved networking capabilities. Replication of databases would not be necessary if researchers had higher bandwidth networks that could transfer a terabyte (TB, or 10^{12} bytes) of data in a few minutes. But doing so requires networking capabilities of tens of gigabits per second (downloading a 1 TB database in 10 minutes demands a network capable of 13 Gbps). Short of such bandwidth, techniques for rapid streaming of data would allow simulations to "pretend" that the data is already local, even though it is being streamed from a remote database. This capability is difficult to achieve routinely

today. Although it is possible, it is not generally implemented, and local replication could be easier.

Linked Simulations

Some biomedical investigations require multiple simulations to be run simultaneously. For example, attempts to understand the physiology of vision might require simulations of both macroscopic and microscopic behaviors, including the quantum mechanics of photoreceptors, the molecular dynamics of macromolecules that respond to light, the population dynamics at the cell membrane as it signals the detection of photons, and the neural network of cells that convey these signals to the brain. These models are all highly interactive and need to share data with one another. The output of one model must be fed into another model. The amount of data transferred between simulations may be large or small, but the effects they induce on subsequent levels of simulation can be substantial. Each of these simulations may itself require significant computation.

Technologies that would facilitate distributed simulations include those for creating uniform techniques for accessing disparate data sources (static, preexisting data, as well as dynamically created data) on the Internet. Many important biomedical questions can be answered only by querying multiple databases, extracting subsets of data, and combining them to determine the final answer. A major software innovation that promises to make biomedical researchers more effective and efficient will be the development of intelligent software agents that assist the investigator in understanding what data are available, what they mean, and how to use them to test new hypotheses. As biomedical researchers perform experiments, such technologies could transfer data directly into a database using the Internet, making them available to other collaborating researchers or computational processes simultaneously. Researchers are creating software to monitor the progress of long, complicated experiments and to alert investigators to unanticipated irregularities in the data or the progress of data collection. As technologies are developed for representing the set of interests for a biomedical researcher (an "interest profile"), intelligent agents could scour the Internet for data of interest and relevance to the researcher, based on this profile. These agents could scan newly published biomedical literature, the publicly accessible Web sites of other scientists, and other Internet information resources, bringing the most relevant sources to the attention of the researcher or abstracting and summarizing them in a manner that is most relevant. Some early examples of this technology have long been available through various journals and online services that notify users when items that match their personal profiles (based on keywords) are published.

The existence of databases on the Internet enables automated (or semiautomated) data mining to extract new principles from data. Data-mining techniques often use statistical associations between variables to postulate relationships that have not been appreciated previously and then go to the available databases seeking evidence to support or refute the association. For such software agents to be effective and reliable, the databases need to be accessible continuously. Even though each individual software agent might not require very high bandwidth, a network experiencing multiple agents operating for millions of individuals will have a large aggregate requirement for bandwidth.

Remote Control of Experimental Apparatus

The Internet provides a means for remotely controlling some of the expensive experimental equipment used in biomedical research, including electron microscopes, DNA sequencing facilities, gene chips for analyzing the expression of nucleic acid or protein sequences, nuclear magnetic resonance spectrometers, and X-ray crystallographic radiation sources.³⁵ In such systems, investigators send samples of interest to device operators, who load the samples and prepare the equipment. The investigators can then run their experiments remotely, specifying the desired magnification, controlling the focus and field of view, and retrieving images as desired. Such systems have proven especially effective in instances (pathology, for example) where the desired information could not be gathered from a set of still images but called for moving the sample and changing the magnification of the microscope (Wolf et al., 1998).

The ability to remotely control experimental equipment offers several benefits. First, it could help make unique or expensive equipment available to a larger number of researchers. Just as networking has opened up the nation's supercomputer resources to the broader research community, so it could open up specialized facilities to the biomedical research community, thereby improving utilization rates. Second, remote access could reduce travel costs associated with experiments. Because sophisticated equipment is scarce, researchers often travel from their home institutions to remote locations to use it, which consumes both time and money. Moreover, the development of appropriate methods for specimen preparation and analysis is often an iterative process that is difficult to complete in a single visit to the laboratory. Remote access to instruments and computation could allow researchers more control over specimen preparation, data collection, and image processing without subjecting them to the time limits of a visit. Long-term studies that require multiple sessions could also be made more practical. Third, the networking of experimental apparatus could allow research results to be more easily

shared among collaborating researchers or displayed to a classroom of students for educational purposes. Most implementations of remotely controlled equipment to date send imagery back to the researcher via a Web site. Any researcher with a password can view the results, and some systems are being developed to allow collaborators to hand off control of the equipment during the course of an experiment.

Simple telermicroscopy systems create images that can be transmitted across the Internet with little difficulty. One system developed in Germany generated full-screen images measuring $1,024 \times 768$ pixels with 8 bits of gray scale, for a total file size of 786 kB. These images could be transferred uncompressed over a 28.8 kbps modem in less than 4 minutes. Using standard JPEG compression, this same image could be transferred in 20 seconds. In experiments with the system, overall response times were dominated by image compression times rather than by delays across the Internet. Indeed, the researchers in Germany were able to reduce response times to 2.5 to 4 seconds across a local area network; the times did not differ significantly when the microscope was operated through a direct Internet connection from other sites in Europe (Wolf et al., 1998).

Nevertheless, for higher resolution images the Internet can introduce significant lag times, especially when multiple images must be retrieved. For example, the National Center for Microscopy and Imaging Research (NCMIR) at the University of California at San Diego houses a state-of-the-art 400 kilo-electron-volt (keV), intermediate-high-voltage electron microscope (IVEM) that can be used to create three-dimensional images from multiple two-dimensional images via a technique known as electron tomography.³⁶ The slices required for three-dimensional reconstruction are $1,024 \times 1,024$ pixels, with 16 bits of precision per pixel, for an image size of 2 MB. A typical data set consists of either 61 or 121 images, depending on experimental requirements (a total of 121 or 242 MB). During peak periods, three to four such tomographic data sets might be acquired in a single day, generating up to about 1 GB of raw data. The intermediate image-processing tasks can easily quadruple that storage requirement and the final tomographic volumes can alone easily exceed 400 MB. A new, high-resolution camera with an image dimension of $2,560 \times 1,960$ 14-bit pixels has boosted data storage requirements by a factor of nearly five.³⁷

The Collaboratory for Microscopic Digital Anatomy (CMDA) is building an infrastructure for allowing researchers to use NCMIR's IVEM and other imaging instruments from a remote site for the purposes of investigating their biological specimens and analyzing the three-dimensional structure using tomography.³⁸ Early experiments with remote operation of the NCMIR found that the Internet was too slow to allow visual guidance of the microscope. As a result, researchers were forced to rely on a

digital survey of the specimen—consisting of a large mosaic of low-magnification images—to guide the process. Features on this survey calibrate the spatial coordinates for remote image acquisition. Researchers examine the survey with specialized software and issue requests to the microscope to image certain areas, create image mosaics, or collect a series of tilted images for tomographic reconstruction.

Improvements in information infrastructure and the anticipated availability of high-speed networks led researchers involved in CMDA to develop a video-based controller for the IVEM that can run on any Java-enabled Web browser. The video controller displays optical and stage parameters for the microscope, the command being executed, and a live video image of the specimen being examined to allow more natural, interactive control. Researchers can adjust the focus, brightness, stage position, and magnification of the microscope, and they can acquire and view high-resolution images of the specimen. Control can be traded among multiple researchers participating in a session, all of whom can view the images. Users can individually set the size of images transmitted to them and the amount of image compression in order to match the speed of their Internet connections to the frame rate desired. During sessions, video streams are generated for 1 to 4 hours.

In experiments conducted to date, simple commands to the microscope were processed in less than 1 second; automated commands for focus and exposure setting were performed in approximately 30 seconds. For users with conventional network connections, video streams were compressed using JPEG algorithms to create grayscale images varying in size from 3 to 12 kB per frame. With these connections, the system performed at a maximum rate of 8 frames per second (96 kB/sec), but average performance was more often in the 3 to 5 frames per second range. Higher bandwidth connections can allow the transmission of full-screen digital imagery to researchers. Such images require approximately 36.5 Mbps of bandwidth.³⁹ In April 1999, researchers were able to use a combination of the vBNS and other networks to allow remote operation of the microscope from Osaka, Japan. High-resolution images were acquired and transmitted in as little as 45 seconds but currently require 36.5 Mbps (before intraframe compression).

The visually guided system has led to a dramatic improvement in remote use of the microscope. Researchers are now able to scan their specimen, find areas of interest, and capture high-resolution images with ease and great precision. The current network infrastructure is adequate for low-resolution, low-frame-rate video, which leads to increased control of the microscope. At times, frame rates are still too slow or latencies fluctuate too much (i.e., there is too much jitter) to provide the level of interactivity required for operations such as manual focusing from a dis-

tance. Higher speed networks and new transport protocols are needed for high-resolution video at full—and constant—frame rates. NCMIR is on the vBNS and expects that a growing number of its collaborators will also join the network or other high-speed networks being developed under the Next Generation Internet or Internet 2 initiatives (see Chapter 1). With the higher transmission rates available on these networks, visually guided control may become more feasible. Use of MPEG compression may also allow higher frame rates to be transmitted over more conventional network connections.

Security, availability, and ubiquity of access are of less concern than bandwidth in the remote control of experimental apparatus but are still important. Security is important for ensuring the integrity of data returned to the investigator and, depending on the nature of the experiments, for maintaining the confidentiality of the data once collected. Reliability is of interest to the extent that researchers want to ensure that the network is available at the time they have been assigned for their experiment. Ubiquity of access is of less concern because experimental apparatus will be used by a small number of highly specialized researchers, most of whom have Internet access through their institutions. However, the system would be better for educational purposes if smaller educational institutions could download images or observe ongoing experiments remotely.

Publication on the Internet

Biomedical research depends on first creating a hypothesis about the world, then designing and running an experiment to test this hypothesis, and finally collecting and analyzing data to determine if the hypothesis is supported or refuted. Because hard-copy publication is so expensive, the scientific community has compromised by publishing papers that present the primary data in summary visual form and that describe the methods used to collect the data, as guarantees that the reported results are accurately reported. With the growth of the Internet, it now becomes possible to consider publishing all scientific data (in its raw form or after some processing) on the Internet for sharing and analysis by other scientists. This forms the basis of the E-Biomed proposal advanced by Harold Varmus, former director of the National Institutes of Health (NIH), for the NIH to house copies of publications and associated primary data sets for the life sciences (Varmus et al., 1999). The physics community has permitted the submission of primary data sets for many years, and certain types of biological data are being released routinely at the time of manuscript publication (for example, DNA sequencing data (GENBANK) and macromolecular structure data (PDB)). The advent of new experimental

technologies (e.g., gene expression arrays, or “DNA chips,” which record the level of expression of a gene product within a cell at a particular moment in time) that produce massive amounts of data makes it attractive to consider large-scale Internet publication of these data sets.

These data sets could be even more useful to the scientific community if they were linked with other data sources to create a grid of related biological information. By having otherwise disconnected data types linked together, computer programs could propose scientific hypotheses based on the data in one set of databases and then test them based on the data in another set of databases. The NCBI has already created a repository of roughly 10 databases that link biographic information, data on genetic sequences and structure, and data on human diseases. Other technologies are being developed for the similar linking of data (e.g., SRS, BioKleisli, and KEGG).⁴⁰

The support of large-scale deposition, storage, and retrieval of primary biomedical data on the Internet calls mainly for availability and security, with moderate emphasis on bandwidth. It is critical that the data be reported accurately on networked resources and that the creators of the data be identified and authenticated; it is also critical that all data be captured and available, in order to avoid losses of valuable scientific data. Latency and ubiquity are less important, since the retrieval of these data is often asynchronous with their collection and is done by specialized researchers.

The success of online publishing of biomedical research findings (both primary data and the conclusions drawn from them), and the much larger audience that such publications may draw, could strain the existing model for scientific peer review. Peer review is essential to ensuring the validity of information published by the scientific community, but too many documents are released for public consumption for them all to be reviewed. Methods will be needed to track which documents have been read, reviewed, and revised by authors in response to critiques and which have not. Many social issues remain to be resolved (e.g., control of publication and dissemination, interaction between peer groups and publishers, and the very definition of peer groups), but technologies are also needed to support the outcome of the social negotiations. For example, methods are needed for (1) providing an enduring stamp of approval for documents on the network so that those that have been reviewed can be identified securely, (2) allowing peer groups to be defined and maintained, (3) searching the Internet to retrieve documents of interest, and (4) validating the authenticity of online documents by, for example, digital watermarking. In this context, it is important to note that the idea of peers can be generalized beyond the current idea that they are a group of scientific investigators from a particular field. Already, other groups have emerged

that may wish to provide a stamp of approval, including disease-specific activist groups, consumer groups, political groups, and others. There is no reason technologies cannot be used by all these groups to label and distinguish documents of interest to their members, using their own criteria.

Collaboration Among Researchers

The Internet could also prove to be a useful medium for enhancing collaboration among biomedical researchers in different locations. The remote control of experimental apparatus is one example of this capability, but others are also possible. For example, envision the following scenario:

Biomedical researchers in three distant cities are interested in the structure and biological function of a new transporter protein whose structure has just been reported in a journal as a result of the Human Genome Project. They believe this newly discovered transporter is expressed in abnormal amounts in a debilitating disease that affects many older individuals within the population. The researchers individually have studied various aspects of the biochemistry associated with this particular disease but think their work could be advanced considerably if they could collaborate with one another. Use of the Internet and specialized network-aware molecular modeling software could enable them to carry on their collaborative research from afar. They could conduct a virtual meeting from their respective offices using the Internet and specialized conferencing and interactive modeling software. Each scientist could display and interactively manipulate three-dimensional molecular models on his or her local workstation as well as the remote workstations of the other collaborators. By using the workstation mouse, one collaborator could, for example, point out the putative binding site on the protein while another suggests a small molecule that he or she thinks might be good at inhibiting the function of this protein. Together, the scientists function as a group and can accomplish much in a short time.

Such scientific collaborations are common and convenient when they take place within an institution, but when the participants are far apart, schedules must be coordinated and travel arranged to a single location. With enhanced Internet services and software, such collaborations could

be performed at a distance as well. If the scenario described above is extended to a larger group, where one of the participants is an instructor and the others are students, the ability of the students to question the instructor interactively (e.g., to use the mouse to point to a portion of the protein in the above example and ask why this portion of the protein doesn't contribute to binding) adds an extremely important quality to the educational experience: the ability to enter into dialog with the instructor.

These kinds of applications would require that the Internet provide sufficient bandwidth to enable real-time multimedia communication among participants. To the extent that participants need to engage in real-time manipulation of biological images, the network would also need to support low latencies. Both distant scientific collaborations and interactive distance learning could benefit substantially from multicast protocols that allow sending network packets to multiple destinations simultaneously and efficiently. In fact, any time multiple recipients are involved, multicast protocols may substantially reduce the impact computer applications have on the network.

Another form of remote collaboration is virtual conferences. A critical element of scientific progress is the ability of scientists to gather at conferences to share new ideas, the latest results, and the latest theories. It is widely recognized that in addition to the formal proceedings at such conferences, the conversations that take place in side rooms are often just as critical for ensuring scientific progress. Thus, there would be some advantage in allowing remote participants not only to attend formal presentations but also to make contacts with their colleagues and have private conversations. Whether the cost of enhancing the Internet to provide such capabilities would exceed the benefit is not yet clear. Building in an infrastructure for ubiquitous real-time videoconferencing would be very expensive. Today, a researcher can attend a remote conference using technologies like RealVideo that produce quite good sound and passable video over the Internet, and when this is coupled with a shared whiteboard or shared applications, there is a good approximation to being there—except for the real-time interaction. Latencies across such networks are typically a few seconds, but that should not keep remote participants from listening to the speaker and viewing their slides. Some systems allow questions to be sent by e-mail or an electronic whiteboard, also with some time delay.

Biomedical research is an international enterprise, and language is still a barrier to communication. Although English is recognized as the dominant language for scientific communication, there are still some applications (especially for informal collaboration) where support for multilingual interactions would accelerate progress. Indeed, language

translation capabilities could be of great help in the consumer health and clinical care arenas. One of the main reasons for the poor access to health care in this country as it becomes increasingly diverse is the number of non-English-speaking persons encountering an English-only-speaking health care system.

Clinical Research

Clinical research involves both clinical trials to establish the efficacy of a drug or a device and the subsequent monitoring of the effectiveness of a product in general (rather than in controlled circumstances) after it has entered into widespread usage.⁴¹ Additional elements of operations management and organizational policy also have heavily clinical research overtones. The Internet can contribute to a number of these activities, as manifest in clinical trials. As computer-based health records become more widely available, health services researchers will likely use them to explore dimensions such as effectiveness and patient satisfaction via the Internet. The Agency for Healthcare Research and Quality, as well as the NIH (and NLM), is likely to become more interested in the potential of the Internet to achieve better quality outcomes and cost management.

Clinical trials are an essential activity in the creation and testing of new drugs and devices for medical diagnosis and therapy. The U.S. Food and Drug Administration requires careful and statistically valid testing by human volunteers before it gives marketing approval. With the mapping of the human genome and the rise of pharmacogenetics, clinical research and clinical trials could become even more prevalent. Knowledge of the availability of clinical trial opportunities, and guidance to conduct them in a timely and accurate fashion, present a significant knowledge distribution and management challenge for which the Internet is a useful infrastructure. Clinical research in human health and disease, such as that supported by the NIH via federal grants and contracts, has similar information management requirements. The Internet provides the capability to enroll patients, validate eligibility, collect data, and disseminate results to and from widely distributed urban and rural sites. Internet-based clinical trials may be extremely important to progress on a number of rare diseases that require large populations of patients in order to make clinical research feasible.

In the area of clinical research and clinical trials of drugs and devices, a growing number of companies and academic centers are using the Internet to recruit volunteer participants. Pursuant to the FDA Modernization Act of 1997, a congressionally mandated national clearinghouse and directory of clinical research studies for serious diseases is being developed as an Internet-accessible resource by the NLM in collaboration

with other federal health and science organizations. There is interesting work at the National Cancer Institute (NCI) on cancer trials using networked information facilities and proposals to mount collaborative national (and international) databases for other clinical trials that might reduce cost or increase effectiveness. Commercial companies are building and making available similar "one-stop-shopping" information resources for patients interested in participating in clinical studies. Since clinical research requires detailed compliance with complex diagnostic and treatment schedules (called clinical protocols), there are both commercial and academic efforts under way to develop detailed, participant-specific protocol guidelines that can be transmitted from a central data management unit via the Internet to participating clinical investigators. Encounter-specific guidance and secure data capture via wide-area computer networks promise to improve the speed with which clinical trials can be completed, as well as to reduce errors of omission and commission in the conduct of clinical research. Current estimates indicate that each day of delay in introducing a new drug to the marketplace costs pharmaceutical companies \$1 million in lost revenues (CyberAtlas, 1999).

Security is an extremely important technological consideration in clinical trials. In addition to concerns about the privacy of patients involved in the trials, there will probably be significant commercial interest in some of the resulting data sets, making security and control of the raw data a serious consideration. Tools will need to be in place to authenticate the source of information, protect the confidentiality of information collected, and protect its integrity. Ubiquity of access is important to the extent that it will allow researchers to draw upon larger population bases for their studies. Depending on the protocol for the trials, access at a physician's office or public kiosk may or may not suffice, and in some situations, access may be needed from the home.

Technical Requirements for Biomedical Research

Bandwidth +++++

The bandwidth requirements for many biomedical research applications are high. Teleconferencing and high-resolution, real-time transfer of images (during remote instrument manipulations, for example) have very high requirements for bandwidth. There is also a trend in the research community toward increasing dependence on the Internet for communicating data and scientific models. It is impossible to predict the long-term needs of biomedical research, but it is likely that the needs for bandwidth will increase as researchers invent new methodologies for the large-scale collection of data about entire genomes, organisms, and com-

munities of organisms. These data may be collected at points all over the world at very high rates. Aggregated traffic back to individual research centers could be very high.

Latency +++

In general, biomedical research is not a time-critical enterprise. There are exceptions, of course, such as the use of the Internet to drive biomedical research instruments (as, for instance, in remote telemicroscopy), where feedback is critical for positioning samples or for adjusting the settings of the instruments. Large distributed simulations also require low latency to improve the speed of their calculations.

Availability ++

For biomedical research, the availability of the network is of moderate importance. Research efforts are not often time-critical and can tolerate low-level losses of data or network unavailability. Obviously, long stretches of such poor performance would be unacceptable, but the needs for availability are not as great in this domain as they might be in clinical care or business applications. Nonetheless, as the Internet plays an ever larger role in research (that is, as it becomes the primary means for accessing primary data, publications, and professional colleagues), it is likely that availability will become more important and even mission-critical for the biomedical research enterprise. Most importantly, only if they perceive an available Internet will reticent adopters of Internet technologies embrace these technologies fully.

Security ++

For the most part, biomedical research deals with public domain information, so the security requirements for the network are not stressed. Since most studies can be done on aggregate data in which no individual patient is identified, issues of privacy are not paramount. If the research deals with patient information (clinical or genomic), however, then security requirements of the Internet jump to the highest levels.

Ubiquity ++

For biomedical research, the ubiquity of the network is not a critical factor. Most major medical centers and research institutions have network connectivity and are motivated to maintain first-class resources to

support their investigators, making the issues of universal access less relevant. One exception to this might be an epidemiological study in which data are collected from people over the Internet. In that case, the network would need to be accessible to all patient populations of relevance to the study.

SUMMARY

Internet applications promise to improve the quality of, and access to, health care while simultaneously reducing its costs. Realizing these applications requires overcoming a number of technical and nontechnical obstacles. For example, quality of service across the Internet must be improved to provide the bandwidth and latency required for applications such as video consultations and remote surgery. Reliability must be improved to ensure that failures of critical network connections occur only infrequently and impose minimal consequences, especially where human life is at stake. Security capabilities must ensure the confidential transmission of health information across the Internet while vouching for the integrity of the information. Access controls must take into account the different access privileges of different kinds of health care workers. And, to achieve its most far-reaching effects, all care providers and patients must have access to the Internet. Additional detail on these needs is provided below. Chapter 3 goes on to examine technical challenges in further detail, while Chapters 4 and 5 provide additional insight into the organizational and policy issues that must be resolved.

Bandwidth

High bandwidth is important for a number of health applications, especially those relying on the transmission of real-time video or large medical or biomedical images. Beyond high bandwidth for specific data-intensive applications there is a need for high aggregate bandwidth to support a high volume of moderately data-intensive applications, such as transfers of large medical records. But bandwidth is not the most important capability for all health care applications. Many consumer health and public health applications, for instance, can currently be supported by the bandwidth available on today's Internet. Bandwidth is particularly important in a number of biomedical research applications, especially in the rendering of three-dimensional images of biomedical structures. It could also be important in professional education, where it would support a virtual reality system for simulated surgeries and other forms of training.

Latency

Certain highly specialized health applications, such as remote control of experimental equipment or simulation of surgical procedures for educational purposes, require much lower latency than is available on today's Internet. However, many other health care applications, such as searching for information on the Internet, do not require instantaneous delivery of information and therefore will not be adversely affected even by the latency of today's Internet.

Availability

Because health care can be a life-and-death matter, the availability of many Internet applications related to its provision and the network across which these applications run is paramount. If time-critical information is not available for decision making because data have been lost in transfer, then the safety and quality of patient care can be compromised. Although some health care applications might have lower requirements for network reliability, the most demanding applications still require a higher level of availability than most consumer applications. If health care organizations are to use the Internet for important patient care tasks—whether retrieving medical records, accessing decision support tools, or conducting telemedicine sessions—they need to know that the network will be available a large percentage of the time.

Security

Because of the highly personal nature of health information and the detrimental effects inappropriate releases of such information could have on social standing, insurance eligibility, and employment, the level of protection required for some health information is extremely high. Such protection must be afforded by security protocols embedded in the relevant applications and in the computers connected to the Internet, as well as in the network itself. It will be as much a matter of the rules governing appropriate releases of information as it will be of technical security mechanisms, such as encryption. Equally or perhaps more important from a quality-of-care standpoint is the need to protect the integrity of data and software and the availability of critical services.

Ubiquity

The continuing trend toward patient empowerment is being fueled by the greater access of patients to general and personal health informa-

tion. The Internet is already playing a large role in improving access to this information, but unfortunately not all Americans are able to benefit. Socioeconomic status and geographic location are still strong determinants of whether a person has access to the Internet. If it is a societal goal to give all persons access to Internet-based health care information and services, then near-ubiquitous access to the Internet will be required.

Use of the Internet in support of health care financial and administrative transactions, public health, professional health education, and biomedical research presents a number of technical challenges that rival those presented by the provision of health care (Table 2.7). Security is of utmost concern in financial and administrative uses, as well as in public health, both of which require access to health records containing patient-specific information. Availability is, in general, of lesser concern than in other health care applications of the Internet, if only because human life is not immediately at stake. Nevertheless, financial and administrative transactions, public health, and biomedical research all require high degrees of system availability—especially public health, where the network would have to continue to function even in the wake of a large-scale disaster. Ubiquity is important in all these applications, although fewer people would need access to the Internet for non-care-related activities than for those directly related to health care.

Beyond these demands for technical capabilities, applications of the Internet in health care financial and administrative transactions, public health, professional education, and biomedical research demand attention to a number of organizational and policy issues. Most importantly,

TABLE 2.7 Relative Importance of Technical Needs of the Internet by Health-Related Applications

Application	Bandwidth	Latency	Availability	Security	Ubiquity
Consumer health	++	+	++	++++	++++
Clinical care	++++	+++	++++	++++	++
Financial and administrative transactions	+	+	+++	++++	++
Public health	+	+	+++	+++	++
Professional education	+++	++	++	+	+++
Biomedical research	++++	+++	++	++	++

NOTE: Plus signs (+) denote the relative importance of the technical feature within the designated application area. The scale ranges from a single plus sign, which denotes minimal importance, to four plus signs, signifying great importance.

organizations engaged in these health-related activities need to recognize the value of the Internet for their missions. Second, they need to develop standards for information exchange, identifying the data elements of importance and agreeing on a standardized vocabulary for describing data and a standardized format for exchanging data. Third, organizations will need to ensure equitable access to Internet resources. This issue may be of greatest importance in the educational arena, where schools have begun to mandate the purchase of laptops by students but have found that some students lack high-bandwidth connectivity from their homes or off-campus work locations. These issues are explored in greater detail in Chapters 4 and 5 of this report.

REFERENCES

- Affiliated Health Information Networks of New England. 1999. *Leading the Way to Health Information Exchange in the Electronic World*. Massachusetts Health Data Consortium, Waltham, Mass., April.
- American Medical Association (AMA). 1996. *Continuing Medical Education Directory*. AMA, Chicago, Ill.
- Baker, D.B. 1998. "PCASSO: Providing Secure Internet Access to Patient Information," *SAIC Science and Technology Trends II*. Science Applications International Corporation, San Diego, Calif.
- Biermann, J. Sybil, G.J. Golladay, M.L. Greenfield, and L.H. Baker. 1999. "Evaluation of Cancer Information on the Internet," *Cancer* 86(3):381-390, August 1.
- Boodman, Sandra G. 1999. "Medical Web Sites Can Steer You Wrong," *Washington Post*, August 10, Health Section, p. 7.
- Burton, Thomas M. 2000. "Medtronic to Join Microsoft, IBM in Patient-Monitoring Venture," *Wall Street Journal*, January 24, p. B12.
- Carns, Ann. 1999. "www.doctorsmedicinesdiseasesgalore.com: Today's Cybercraze Is Any Web Site Devoted to Health or Maladies," *Wall Street Journal*, June 10, p. B1.
- Centers for Disease Control and Prevention (CDC). 1998. *Strengthening Community Health Protection Through Technology and Training: The Health Alert Network*. CDC, Atlanta, Ga.
- Chand, G., B.C. Breton, N.H.M. Caldwell, and D.M. Holburn. 1997. "World Wide Web-Controlled Scanning Electron Microscope," *Scanning* 19:292-296.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1997. *For the Record: Protecting Electronic Health Information*. National Academy Press, Washington, D.C.
- CyberAtlas. 1999. "Online Healthcare Market Looks Energized." Available online at <http://cyberatlas.internet.com/big-picture/demographics/article/0,1323,6061_153701,00.html>.
- Davis D.A., M.A. Thomson, A.D. Oxman, and R.B. Haynes. 1995. "Changing Physician Performance: A Systematic Review of the Effect of Continuing Medical Education Strategies," *Journal of the American Medical Association* 274(September 6):700-705.
- Dolin, R.H., W. Rishel, P.V. Biron, J. Spinosa, and J.E. Mattison. 1998. "SGML and XML as Interchange Formats for HL7 messages," pp. 720-724 in *Proceedings of the AMIA Symposium*, Bethesda, Md.

- Fridsma, D.B., P. Ford, and R. Altman. 1994. "A Survey of Patient Access to Electronic Mail: Attitudes, Barriers, and Opportunities," Paper presented at Eighteenth Annual Symposium on Computer Applications in Medical Care, Washington, D.C., October 15-19. See <http://smi-web.standord.edu/pubs/SMI_Abstracts/SMI-94-0524.html>.
- Goedert, Joseph, 1999. "Electronic Claims Growth Sputters," *Health Data Management* (September):84-86.
- Harman, J. 1998. "Topics for Our Times: New Health Care Data—New Horizons for Public Health," *American Journal of Public Health* 88:1019-1021.
- Health Care Financing Administration (HCFA). 1999a. *HCFA Information System Security Bulletin Handbook*, Bulletin 98-01, Baltimore, Md., January.
- Health Care Financing Administration (HCFA). 1999b. "Telecommunications Requirements: Migration of Medicare Managed Care Organizations (MCO) to the Medicare Data Communications Network and the Replacement of the RLINK Software," Operational Policy Letter No. 92 OPL99.092, U.S. Department of Health and Human Services, May 6. Available online at <www.hcfa.gov/medicare/op1092.htm>.
- Hripscak, G., P.D. Clayton, T.A. Pryor, P. Haug, O.B. Wigertz, and J. Van der Lei. 1990. "The Arden Syntax for Medical Logic Modules," pp. 200-204 in *Proceedings of the Symposium on Computer Applications in Medical Care*, R.A. Miller, ed. IEEE Computer Society Press, Los Alamitos, Calif.
- Huang, H.K. 1996. "Teleradiology Technologies and Some Service Models," *Computerized Medical Imaging and Graphics* 20(2):59-68.
- Huang, H.K. 1999. *PACS: Basic Principles and Applications*. Wiley-Liss, New York.
- Institute of Medicine (IOM), Committee on the Quality of Health Care in America. 1999. *To Err Is Human*, Linda Kohn, Janet Corrigan, and Marla Donaldson, eds. National Academy Press, Washington, D.C.
- Kohane, I.S., P. Greenspun, J. Fackler, C. Cimino, and P. Szolovits. 1996. "Building National Electronic Medical Record Systems via the World Wide Web," *Journal of the American Medical Informatics Association* 3(3):191-207.
- Lasker, R.D. 1998. "Challenges to Accessing Useful Information in Health Policy and Public Health: An Introduction to a National Forum Held at the New York Academy of Medicine," *Journal of Urban Health: Bulletin of the New York Academy of Medicine* 75(4):779-784.
- Lou, S.L., Edward A. Sickles, H.K. Huang, David Hoogstrate, Fei Cao, Jun Wang, and Mohammad Jahangiri. 1997. "Full-field Direct Digital Telemammography: Technical Components, Study Protocols, and Preliminary Results," *IEEE Transactions on Information Technology in Biomedicine* 1(4):270-278.
- Mandl, Kenneth D., Isaac Kohane, and Allan M. Brandt. 1998. "Electronic Patient-Physician Communication: Problems and Promise," *Annals of Internal Medicine* 129:495-500.
- McCormack, John. 2000. "Group Practices Find Their Way to the Internet," *Health Data Management* 8(1):46-53.
- McGinnis, J.M., and W.H. Foege. 1993. "Actual Causes of Death in the United States," *Journal of the American Medical Association* 270:2207-2212.
- Nash, Sharon. 1999. "The Doctor Is Online," *PC Magazine Online*, July 14.
- Resnick, Paul. 1997. "Filtering Information on the Internet," *Scientific American* (March):106-108.
- Reuters New Service. 1999. "Internet Could Organize Medical Records," July 27.
- Rind, D.M., I.S. Kohane, P. Szolovits, C. Safran, H.C. Chueh, and G.O. Barnett. 1997. "Maintaining the Confidentiality of Medical Records Shared over the Internet and World Wide Web," *Annals of Internal Medicine* 127(2):138-141.
- Rybowski, Lise, and Richard Rubin. 1998. *Building an Infrastructure for Community Health Information: Lessons from the Frontier*. Foundation for Health Care Quality, Seattle.

- Science Applications International Corporation (SAIC). 1998. *Security and Risk Management for Business-to-Business Health Information Networks*, Final Report, Three State Health Information Planning Project. SAIC, San Diego, Calif., June.
- Science Panel on Interactive Communication and Health (SCIPICH). 1999. *Wired for Health and Well-Being: The Emergence of Interactive Health Communication*, Thomas R. Eng and David H. Gustafson, eds. Office of Disease Prevention and Health Promotion, U.S. Department of Health and Human Services, Washington, D.C., April. Available online at <<http://www.scipich.org>>.
- USA Today*. 1998. "Health-Related Activities Conducted Online," July 10.
- U.S. Department of Health and Human Services. 1998. *Healthy People 2010 Objectives*. Draft for public comment, September 15, U.S. Department of Health and Human Services, Washington, D.C. Available online at <<http://web.health.gov/healthypeople>>.
- U.S. Public Health Service, Public Health Data Policy Coordinating Committee. 1995. *Making a Powerful Connection: The Health of the Public and the National Information Infrastructure*. July 6. Available online at <www.nlm.nih.gov/pubs/staffpubs/lo/makingpd.html>.
- Varmus, Harold, David Lipman, and Pat Brown. 1999. "E-BIOMED: A Proposal for Electronic Publications in the Biomedical Sciences," memorandum dated May 5. Available online at <<http://www.nih.gov/welcome/director/pubmedcentral/ebiomedarch.htm>>.
- Wolf, Guenter, Detlev Petersen, Manfred Dietel, and Ever Petersen. 1998. "Telemicroscopy via the Internet," *Nature* 391(February 5):613-614.
- World Wide Web Consortium. 1998. "Extensible Markup Language (XML) 1.0. W3C Recommendation," Report No. REC-xml-19980210, February.

NOTES

1. A search using AltaVista on July 29, 1999, returned 40,156 Web pages in response to the query "diabetes mellitus."
2. For an example of the criteria according to which health-related Web sites can be evaluated, see <<http://hitiweb.mitretrek.org/iq/onlycriteria.html>>.
3. Information on PICS is available online at <<http://www.w3.org/PICS/>>. See also Resnick (1997).
4. For example, a company named PersonalMD.com had stored the health records of 10,000 subscribers online free of charge as of July 1999. The company sends consumers a card with a personal access code that allows them to retrieve their records over the Internet or by a fax-back system (Reuters News Service, 1999). Another group, the Medical Registry, charges \$100 to retain medical information online, allowing customers to update it as often as they wish.
5. The Medical Registry, which was started by emergency room physicians, allows doctors to access a patient's record during an emergency by entering their Drug Enforcement Act number. Patients are issued a wallet card and alert bracelet containing the address of the Web site, the patient's password, and the phone number of a fax-back service that can access and download the patient's records.
6. For more information on PCASSO, see Baker (1998).
7. The National Heart Attack Alert Program is a federal effort that may lead to improved techniques for remotely monitoring patients. The program has the overall goals of, first, reducing morbidity and mortality from acute myocardial infarctions (heart attacks) through rapid identification and treatment and, second, heightening the potential for an improved quality of life for patients and family members. Remote monitoring and collec-

tion of patient vital signs is seen as one possible avenue for early detection of heart attacks and for getting patients into the health care system quickly. Information about the program is available online at <http://www.nhlbi.nih.gov/about/nhaap/nhaap_pd.htm>.

8. Data from Michael Kiensle, associate dean for Clinical Affairs and BioMedical Communications, University of Iowa College of Medicine, personal communication, July 12, 1999.

9. In-home monitoring with a video link offers benefits to patients, but not for diagnostic reasons. As one reviewer of an early draft of this report noted, the patient needs to see the care provider to address the problem of noncompliance, which often results when patients misunderstand instructions and take medications at the wrong time, in the wrong dosage, and so on. The way to improve compliance is to ensure that the care provider captures the attention of the patient while delivering instructions. Video can help ensure this happens.

10. At present, teleconsultations conducted across networks that use the IP require approximately twice the bandwidth of traditional point-to-point networks. The reasons are twofold: (1) Internet protocols impose some additional overhead functions that require bandwidth and (2) the devices used to encode video streams into IP packets (coder/decoders, or codecs) are much less efficient than their non-IP counterparts. But IP codecs are less expensive, in part because they carry less hardware compression, and next-generation IP codecs are expected to provide better performance and impose less of a penalty on IP-based systems.

11. East Carolina University recently received a grant from the National Library of Medicine to investigate these requirements.

12. Pending further study of the medical efficacy of higher bandwidth for teleconsultations, an upper limit on bandwidth for video consultations can be estimated by considering the need for broadcast quality video. A video display with 640×480 pixels that is refreshed 30 times per second and has 24-bit color demands 221 Mbps. With standard compression technologies, such as that of the Motion Picture Experts Group (MPEG), reductions of 90 to 1 are common, resulting in a need for 2.5 Mbps. Improved coding may lower this figure further. For transmission quality equal to high-definition television, which is just entering consumer production, 19 Mbps would be required. These figures represent the maximum bandwidth that remote video consultations could be expected to use, but, as the evidence collected by ECU and other practitioners indicates, much less bandwidth is sufficient in many applications.

13. Information on the National Laboratory for the Study of Rural Telemedicine at the University of Iowa is available online at <<http://telemed.medicine.uiowa.edu/index.html>>.

14. Anthony Chou, University of California at San Francisco, presentation to the committee, December 16, 1998.

15. As described later in this chapter, attempts are being made to make these specialized instruments available to a larger number of researchers through the Internet.

16. Stentor, Inc., has developed a system that can provide high-resolution images over lower-bandwidth networks by providing only portions of the overall image at any one time.

17. In addition to the lack of standardization of medical data models, there has been no widespread adoption of portable decision-support tools, despite the efforts of many in projects such as the development of the Arden syntax (see Hripcsak et al., 1990). The absence of sound, widely accepted automated decision-support tools that are integrated with each other and with Internet health transactions will undermine the capabilities of such tools to achieve the desired goal of medical error reduction. For example, if one set of Internet transactions attempts to optimize for medication orders and another set of Internet transactions attempts to optimize the ordering of procedures, several possibly dangerous

and/or expensive interactions between the two might occur. In a tightly integrated system, as compared to disparate and separate Internet-based systems, such interactions might be minimized. This situation suggests that a near-term challenge will be to ensure quality control and coordination among the many different Internet-born clinical transactions and to develop robust medical decision-support tools that can serve a wide range of institutions and patient populations.

18. In a survey of 153 chief information officers conducted by the College of Health Information Management Executives in 1998, 80 percent said they use HL7 and 13.5 percent planned to implement it in the future.

19. All claims data in this paragraph derive from research conducted for Faulkner & Gray's *2000 Health Data Directory*, as cited in Goedert (1999).

20. For additional information on these efforts, see Rybowski and Rubin (1998) and Affiliated Health Information Networks of New England (1999).

21. Further information on HCFA's pilot program can be obtained from either <www.wedi.org> or <www.afecht.org>.

22. For example, the U.S. Public Health Service released a report in 1995 describing the potential applications of the Internet in public health and identifying technical challenges to be addressed (U.S. Public Health Service, 1995). In 1997, the New York Academy of Medicine and the National Library of Medicine cosponsored a symposium on public health informatics that called for improved structures and assessment mechanisms for public health information (Lasker, 1998). Slide presentations of several symposium speakers are available at <<http://www.nlm.nih.gov/nichsr/nyam/nyam.html>>. The Department of Health and Human Services' document *Healthy People 2010* (U.S. Department of Health and Human Services, 1998) includes a section on objectives for improving the public health infrastructure. They include widespread access to the Internet and real-time, on-site access to public health data for public health workers and individuals. Section 14, objectives 5 and 6, is the most relevant example.

23. Participating organizations include the National Network of Libraries of Medicine, the Centers for Disease Control and Prevention, the Health Resources and Services Administration, the Association of State and Territorial Health Officials, and the National Association of County and City Health Officials.

24. Reports from physicians' offices and hospitals also tend to be reported on paper.

25. Jac Davies, Washington State Department of Health, presentation to the study committee, February 11, 1999, Seattle, Washington.

26. The traditional public health functions are surveillance, case identification, treatment, prevention, research, guidelines, education and feedback.

27. President Clinton's proposal for this program would also create a network of regional labs to provide rapid analysis and identification of select biological agents.

28. The Health Alert Network is part of a larger antibioterrorism effort that received \$158 million in FY99. Another \$72 million was proposed for FY2000, which would raise the total to \$230 million.

29. This information is derived from "Health Alert Network Architectural Standards," supplement to the Centers for Disease Control and Prevention Program Announcement No. 99051.

30. The Association of American Medical Colleges reports that total enrollment in full-time undergraduate medical programs in the United States was 66,900 in the 1997-1998 academic year. There were 99,099 residents being trained in clinical settings (primarily teaching hospitals). According to the quinquennial survey, approximately 242,000 students were enrolled in all health sciences programs during the 1996-1997 academic year.

31. The SHINE project at Stanford Medical Center is experimenting with providing CME

credit to physicians who request point-of-care information during patient interactions. Information on this program is available online at <http://shine.stanford.edu>.

32. These figures were provided by Dennis Benson at the National Library of Medicine in a personal communication dated February 11, 2000.

33. There have been laboratories whose access to NCBI/PubMed was suspended temporarily when usage rates climbed too high. One lab at Stanford lost access after a graduate student wrote programs that were downloading 3,000 abstracts per minute from the Web site. The scientific goals of this student were meritorious, but the resource was not built to sustain this use (Russ Altman, Stanford University, personal communication, December 22, 1999).

34. James Ostell, National Center for Biotechnology Information, presentation to the study committee on March 1, 1999, Washington, D.C.

35. Researchers at the University of Cambridge, the University of California at San Diego (see Box 2.4), and the University Hospital Charité in Berlin have all developed Internet-based systems for controlling experimental apparatus (Chand et al., 1997).

36. Electron tomography is a technique whereby three-dimensional structure is derived from a series of two-dimensional projections using advanced image processing steps. In the most common form, the specimen is tilted around a single axis and imaged at regular intervals. The IVEM at NCMIR is one of a few such instruments in the United States made available to the biological research community. Support for NCMIR is provided by the National Center for Research Resources (NCRR) of the National Institutes of Health (NIH).

37. This information is taken from a paper entitled "NCMIR's Collaboratory for Microscopic Digital Anatomy: A National Science Foundation National Challenge Project," which is available online at www-ncmir.ucsd.edu/CMDA/.

38. CMDA has already been used by researchers at Montana State University to collect data on synaptic organization in the sensory ganglia of the insect nervous system and by scientists at the University of Oregon studying neurotransmission (synaptic vesicle release) in vestibular hair cell synapses. Other users are studying the abnormalities in nerve cells in Alzheimer's disease, the structural relationships of protein molecules responding to calcium within nerve cells, and the three-dimensional pattern of branching of the dendrites in neurons that create a highly linked network of cellular communication.

39. In the longer term, it is hoped that digital video standards will give good resolution and smooth motion at 30 frames per second at much lower bandwidth.

40. More information on SRS is available at <http://srs.ebi.ac.uk:5000/>. Information on Biokleisli is available at <http://smi-web.stanford.edu/projects/helix/mis214/bdkowvldb95.pdf> (a paper). Information on KEGG is available at <http://www.genome.ad.jp/dbget/dbget.links.html>.

41. Clinical research lies at the juncture of clinical care, biomedical research, and public health but is somewhat distinct from each of these topics. It is described in the section on biomedical research in this report for reasons of editorial convenience and exposition.

3

Technical Challenges

Ongoing efforts to develop and deploy improved networking technologies promise to greatly enhance the capabilities of the Internet. Protocols for quality of service (QOS) could enable vendors to offer guarantees on available bandwidth and latencies across the network. Advanced security protocols may better protect the confidentiality of messages sent across the network and ensure that data are not corrupted during transmission or storage. Broadband technologies, such as cable modems and digital subscriber line (DSL) services, have the potential to make high-speed Internet connectivity more affordable to residential users and small businesses. In combination, these capabilities will enable the Internet to support an ever-increasing range of applications in domains as disparate as national security, entertainment, electronic commerce, and health care.

The health community stands to benefit directly from improvements in QOS, security, and broadband technologies, even though it will not necessarily drive many of these advances. Health applications—whether supporting consumer health, clinical care, financial and administrative transactions, public health, professional education, or biomedical research—are not unique in terms of the technical demands they place on the Internet: nearly all sectors have some applications that demand enhanced QOS, security, and broadband technologies. Nevertheless, particular health applications require specific capabilities that might not otherwise receive much attention. Use of the Internet for video-based consultations with patients in their homes, for example, would call for two-way, high-bandwidth connections into and out of individual residences, whereas

video-on-demand applications require high bandwidth in only one direction. Human life may be at risk if control signals sent to medical monitoring or dosage equipment are corrupted or degraded, or if electronic medical records cannot be accessed in a timely fashion. Even when no lives are at stake, the extreme sensitivity of personal health information could complicate security considerations, and the provisions of health care at the point of need—whether in the hospital, home, or hotel room—could increase demand for provider and consumer access to Internet resources via a variety of media.

This chapter reviews current efforts to improve the capabilities of the Internet and evaluates them on the basis of the needs of the health sector outlined in Chapter 2. Particular attention is paid to the need for QOS, for security (including confidentiality of communications, system access controls, and network availability), and for broadband technologies to provide end users with high-speed connectivity to the Internet. Also discussed are privacy-enhancing technologies, which are seen by many as a prerequisite for more extensive use of the Internet by consumers. The chapter identifies ways in which the Internet's likely evolution will support health applications and ways in which it may not. It gives examples of challenges that real-world health applications can pose for networking research and information technology research more generally. In this way, it attempts to inform the networking research community about the challenges posed by health applications and to educate the health community about the ways in which ongoing efforts to develop and deploy Internet technologies may not satisfy all their needs.

QUALITY OF SERVICE

Quality of service is a requirement of many health-related applications of the Internet. Health organizations cannot rely on the Internet for critical functions unless they receive assurances that information will be delivered to its destination quickly and accurately. For example, care providers must be able to retrieve medical records easily and reliably when needed for patient care; providers and patients must be able to obtain sustained access to high-bandwidth services for remote consultations if video-based telemedicine is to become viable. In emergency care situations, both bandwidth and latency may be critical factors because providers may need rapid access to large medical records and images from disparate sources connected to the Internet. Other applications, such as Internet-based telephony and business teleconferencing, demand similar technical capabilities, but the failure to obtain needed QOS in a health application might put human life at risk.

Compounding the QOS challenge in health care is the variability of a

health care organization's needs over the course of a single day. The information objects that support health care vary substantially in size and complexity. While simple text effectively represents the content of a care provider's notes, consultation reports, and the name-value pairs of common laboratory test results, many health problems require the acquisition and communication of clinical images such as X rays, computed tomography (CT), and magnetic resonance imaging (MRI). The electronic forms of these images, which often must be compared with one another in multiple image sets, comprise tens to hundreds of megabytes of information that may need to be communicated to the end user within several seconds or less. Medical information demands on digital networks are thus notable for their irregularity and the tremendous variation in the size of transmitted files. When such files need to be transmitted in short times, very high bandwidths may be required and the traffic load may be extremely bursty.

No capabilities have yet been deployed across the Internet to ensure QOS. Virtually all Internet service providers (ISPs) offer only best-effort service, in which they make every effort to deliver packets to their correct destination in a timely way but with no guarantees on latency or rates of packet loss. Round-trip times (or latencies) for sending messages across the Internet between the East and West Coasts of the United States are generally about 100 milliseconds, but latencies of about 1 second do occur—and variations in latency between 100 milliseconds and 1 second can be observed even during a single connection.¹ Such variability is not detrimental to asynchronous applications such as e-mail, but it can render interactive applications such as videoconferencing unusable. Similarly, the rates of packet loss across the Internet range from less than 1 percent to more than 10 percent; high loss rates degrade transmission quality and increase latencies as lost packets are retransmitted. Furthermore, because many applications attempt to reduce congestion by slowing their transmission rates, packet loss directly affects the time taken to complete a transaction, such as an image transfer, over the network.

Several approaches can be taken to improve QOS across the Internet, with varied levels of effectiveness. For example, Internet users can upgrade their access lines to overcome bottlenecks in their links to ISPs, but such efforts affect bandwidth and latency into and out of their own site only. They provide no means for assuring a given level of QOS over any distance. Similarly, ISPs can attempt to improve service by expanding the capacity of their backbone links. However, as described below, such efforts provide no guarantees that bandwidth will be available when needed and contain no mechanisms for prioritizing message traffic in the face of congestion. To overcome these limitations, efforts are under way to develop specific protocols for providing QOS guarantees across the

Internet. These protocols promise to greatly expand the availability of guaranteed services across the Internet, but their utility in particular applications may be limited, as described below.

Increasing Bandwidth

One approach taken by ISPs to improve their data-carrying capacity and relieve congestion across the Internet has been to dramatically increase the bandwidth of the backbones connecting points of presence (POPs).² Today's backbone speeds are typically on the order of 600 megabits per second (Mbps) to 2.5 gigabits per second (Gbps), but some ISPs have considerably more bandwidth in place. A number of ISPs today have tens of strands of fiber-optic cable between their major POPs, with each strand capable of carrying 100 wavelengths using current wavelength division multiplexing (WDM) technology. Each wavelength can support 2.5 to 10 Gbps using current opto-electronics and termination equipment. Thus, an ISP with 30 strands of fiber between two POPs theoretically could support 30 terabits per second (Tbps) on a single inter-POP trunk line.³ This is enough capacity to support approximately 450 million simultaneous phone calls or to transmit the 40 gigabyte (GB) contents of the complete MEDLARS collection of databases in one-hundredth of a second.

Even with this fiber capacity in the ground, most ISPs currently interconnect their POPs at speeds significantly lower than 1 Tbps—a situation that is likely to persist for the next few years. The limiting factors are the cost and availability of the equipment that needs to be connected to the fiber inside the POP. This equipment includes Synchronous Optical Network (SONET)⁴ termination equipment and the routers or switches that are required to forward packets between POPs. The SONET equipment is expensive—as are the routers and switches that connect to the SONET equipment—so ISPs have an incentive to deploy only enough to carry the expected traffic load. More importantly, routers are limited in terms of the amount of traffic they can support. As of late 1999, the leading commercial routers available for deployment could support 16 OC-48 (2.5 Gbps) interfaces, with a fourfold increase (e.g., to 16 × OC-192) expected to be deployable in the next 1 to 2 years. Terabit and multiterabit routers with a capacity at least six times greater than a 16 × OC-192 router are under development. Despite these increases in capability, routers most likely will continue to limit the bandwidth available between POPs for the foreseeable future. The commercial sector understands the need for faster routers and is addressing it, at least to meet near-term demands for higher link speeds. Additional research on very high speed routers may be justified to provide longer-term improvements in data-carrying capacity.

Increases in the bandwidth of the Internet backbone alleviate some of the concerns about QOS but may not completely eliminate congestion. Demand for bandwidth is growing quickly, and it appears that ISPs are deploying additional bandwidth just fast enough to keep up. Current traffic measurements indicate that some Internet backbone links are at or near capacity. Factors driving the growth in demand for bandwidth include the increasing number of Internet users, the increasing amount of time the average user spends connected to the Internet, and new applications that are inherently bandwidth-intensive (and that demand other capabilities, such as low latency and enhanced security). Nielsen ratings for June 1999 put the number of active Web users at 65 million for the month and average monthly online time per user at 7.5 hours, up from 57 million users and 7 hours per user just 3 months earlier.⁵ As an example of increasing bandwidth demands, medical image files that now contain about 250 megabytes (MB) of data are expected to top several gigabytes in the near future as the resolution of digital imaging technology improves.

Internet protocols further limit the capability of ISPs to provide QOS by simply increasing bandwidth. The Transmission Control Protocol (TCP), which underlies most popular Internet applications today, is designed to determine the bandwidth of the slowest or most congested link in the path traversed by a particular message and to attempt to use a fair share of that bottleneck bandwidth. This trait is important to the success of the Internet because it allows many connections to share a congested link in a reasonably fair way. However, it also means that TCP connections always attempt to use as much bandwidth as is available in the network. Thus, if one bottleneck is alleviated by the addition of more bandwidth, TCP will attempt to use more bandwidth, possibly causing congestion on another link. As a result, some congested links are almost always found in a network carrying a large amount of TCP traffic. Adding more capacity in one place causes the congestion to move somewhere else. In many cases, top-tier service ISPs attempt to make sure they have enough capacity so that the congestion occurs in other backbone providers' networks. The only way out of this quandary, apparently, is to provide so much bandwidth throughout the network that applications are unable to use it fully.

Applications that do not use TCP are not the solution, either, because they also tend to consume considerable bandwidth. Such nonadaptive applications are typically those involving real-time interaction, for which TCP is not well suited. Internet telephony is a good example of such an application. Although an individual call might use only a few kilobits per second, many current Internet telephony applications transmit data at a constant rate, regardless of any congestion along their path. Because these applications do not respond to congestion, large-scale deployment can lead to a situation called congestive collapse, in which links are so

overloaded that they become effectively useless. Furthermore, when these applications share links with TCP-based applications, the latter will respond to congestion to the point where they may become unusable. Short of deploying additional bandwidth in the Internet and replacing nonadaptive applications with adaptive ones, the primary approach to addressing this problem is either to equip routers with new mechanisms to prevent congestive collapse or to provide suitable incentives to encourage the development of adaptive applications.

More fundamental factors also limit the utility of increased bandwidth as a means of solving the QOS problem. Adequate bandwidth is a necessary but not sufficient condition for providing QOS. No user can expect to obtain guaranteed bandwidth of 100 Mbps across a 50-Mbps link; similarly, it is not possible to guarantee 10 Mbps each to 1,000 applications that share a common link unless that link has a capacity of at least 10 Gbps.⁶ The simple fact that Internet backbones are shared resources that carry traffic from a large number of users means that no single user can be guaranteed a particular amount of bandwidth unless dedicated allocation mechanisms are in place. In the absence of QOS mechanisms, it is impossible to ensure that delay-sensitive applications are protected from excessive time lags.⁷

In theory, ISPs could attempt to provide so much extra bandwidth to the Internet that peak demand could almost always be met and service quality would improve (a technique referred to as overprovisioning, used with some success in local area networks, or LANs). However, research indicates that overprovisioning is an inefficient solution to the QOS problem, especially when bandwidth demands vary widely among different applications, as is the case in health care. Overprovisioning tends not to be cost-effective for leading-edge, high-bandwidth applications—even those that can adapt to delays in the network. If the objective is to make efficient use of networking resources and provide superior overall service, then mechanisms that enable the network to handle heterogeneous data types appear preferable to the separation of different types of data streams (e.g., real-time video, text, and images) into discrete networks (Shenker, 1995).

A number of efforts are under way in the networking community to develop mechanisms for providing QOS across the Internet. The two main approaches are differentiated services (diff-serv) and integrated services (int-serv). Although they are very different, both attempt to manage available bandwidth to meet customer-specific needs for QOS. Both diff-serv and int-serv will enable greater use of the Internet in some health applications, but it is not clear that these programs will meet all the needs posed by the most challenging health applications.

Differentiated Services

Recent efforts in the Internet Engineering Task Force (IETF) have resulted in a set of proposed standards for diff-serv across the Internet (Blake et al., 1998). As the name implies, diff-serv allows ISPs to offer users a range of qualities of service beyond the typical best effort. The ISPs were active in the definition of these standards, and several are expected to deploy some variant of diff-serv in 2000.

Differentiated services do not currently define any mechanisms by which QOS levels could be determined for different communications sessions on demand; rather, initial deployment is likely to be for provisioned QOS that is agreed upon a priori. As a simple example, a customer of an ISP might sign up for premium service at a certain rate, say 128 kilobits per second (kbps). Such a service would allow the customer to send packets into the network at a rate of up to 128 kbps and expect them to receive better service than a best-effort packet would receive. Exactly how much better would be determined by the ISP. If the service were priced appropriately, then the provider might provision enough bandwidth for premium traffic to ensure that loss of a premium packet occurred very rarely, say once per 1 million packets sent. This would provide customers with high assurance that they could send at 128 kbps at any time to any destination within the ISP's network.

Many variations of this basic service are possible. The service description above applies to traffic sent by the customer; it is also possible to provide high assurance of delivery for a customer's inbound traffic. Similarly, an ISP could offer a service that provides low latency. It is likely that providers would offer several grades of service, ranging from the basic best-effort through premium to superpremium, analogous to coach, business, and first class in airline travel. A customer might sign up for several of these services and then choose which packets need which service. For example, e-mail might be marked for best-effort delivery, whereas a video stream might be marked for premium. Customers then would need to develop their own policies to determine which types of traffic flows would be transmitted at different QOS levels.

Although diff-serv is an improvement over best-effort services, it has several limitations that might preclude its use for some health-related applications. First, research has shown that simple diff-serv mechanisms (e.g., those that classify QOS levels at the edge of the network and provide differential loss probabilities in the core) can be used to provide a high probability of meeting users' QOS preferences for point-to-point communications (Clark and Wroclawski, 1997). However, in the absence of significant overprovisioning and explicit signaling to reserve resources, such guarantees are probabilistic, which virtually precludes absolute, quantifi-

able service guarantees. The QOS provided by diff-serv depends largely on provisioning of the network to ensure that the resources available for premium services are sufficient to meet the offered load. The provision of sufficient resources to make hard guarantees may be economically feasible only if premium services are significantly more expensive than today's best-effort service. Indeed, ISPs need to have in place some incentive mechanism (such as increased charges for higher-quality service) to ensure that customers attempt to distinguish between their more important and less important traffic.

A second limitation is that diff-serv can be offered most easily across a single ISP's network. Current standards do not define end-to-end services, focusing instead on individual hops between routers in the network. There are no defined mechanisms for providing service guarantees for packets that must traverse the networks of several ISPs. For many service providers, offering diff-serv across their own networks is likely to be a valuable first step—especially for providers with a national presence that will be able to provide end-to-end service to large customers with sites in major metropolitan areas. Services like these are also valuable over especially congested links, such as the transoceanic links; again, these types of services could be offered by a single provider. Nevertheless, there obviously would be great value in obtaining end-to-end QOS assurance even when the two ends are not connected to the same provider. To some extent, the diff-serv standards have laid the groundwork for interprovider QOS, because packets can be marked in standard ways before crossing a provider boundary. A provider connecting to another provider is in some sense just a customer of that provider. Provider A can buy premium service from provider B and resell that service to the customers of provider A. However, the services that providers offer are not likely to be identical, so the prospect of obtaining predictable end-to-end service from many providers seems considerably less certain than does single-provider QOS.

Third, diff-serv does not currently allow users to signal a request for a particular level of QOS on an as-needed basis (as is possible with the integrated services model, described below). Health care organizations have widely varying needs for bandwidth over time. For example, a small medical center occasionally might need to transmit a mammography study of 100 MB in a short time interval—creating a need for high bandwidth over that interval—but it is unlikely to need even close to that amount of bandwidth on average. Thus, a dynamic model of QOS would be preferable. Diff-serv does not preclude such a model; it simply provides a number of QOS building blocks, which could be used to build a dynamic model in the future. A variety of means for dynamically signaling diff-serv QOS are under investigation by networking researchers.

Finally, the diff-serv approach may not provide a means of differentiating among service levels with sufficient granularity to meet the QOS needs of critical applications, such as remote control of medical monitoring or drug delivery devices. In the interests of scalability, diff-serv sorts traffic into a small number of classes; as a result, the packets from many applications and sites share the same class and can interfere with each other. For example, a physician downloading a medical image could inadvertently disrupt data flows from in-house monitoring equipment if they are on the same network and share a diff-serv class. Although policing of traffic at the edges of the network helps to ensure that applications of the same class do not interfere with each other, it does not completely isolate applications. Stronger isolation, and thus a larger number of classes, may be required for some demanding applications.

Integrated Services

In contrast to the diff-serv model, int-serv (Braden et al., 1994) provides quantifiable, end-to-end QOS guarantees for particular data flows (e.g., individual applications) in networks that use the IP.⁸ The guarantees take the form of "this videoconference from organization A to organization B will receive a minimum of 128 kbps throughput and a maximum of 100 milliseconds end-to-end latency." To accommodate such requests, int-serv includes a signaling mechanism called resource reservation protocol (RSVP) that allows applications to request QOS guarantees (Braden et al., 1997).⁹ Int-serv provides a service model that in some ways resembles that of the telephone network, in that service is requested as needed. If resources are available to provide the requested service, then the service will be provided; if not, then a negative acknowledgment (equivalent to a busy signal) is returned. For this reason, int-serv already is being used in some smaller networks to reserve bandwidth for voice communications.

Several obstacles stand in the way of the deployment of int-serv across the Internet. The major concern is scalability. As currently defined, every application flow (e.g., a single video call) needs its own reservation, and each reservation requires that a moderate amount of information be stored at every router along the path that will carry the application data. As the network itself grows and the number of reservations increases, so does the amount of information that must be stored throughout the network.¹⁰ The prospect of having to store such information in backbone routers is not attractive to ISPs, for which scalability is a major concern.¹¹

Additional impediments arise from difficulties in administering int-serv reservations that cross the networks of multiple ISPs. Methods are needed for allocating the costs of calls that are transmitted by multiple

ISPs; new ways of billing users and making settlements among ISPs may be required. These are QOS policy issues, discussed below. It is clear that the management of reservations that traverse multiple ISPs, each with its own administrative and provisioning policies, will be quite complex. Solutions to these problems must address the possibility that one or more ISPs might experience a failure during the life of a reservation, resulting in the need to reroute the traffic significantly (Birman, 1999). Such concerns, if not successfully addressed, could slow or thwart the deployment of int-serv capabilities throughout the Internet.

Alternative Quality of Service Options

Given the difficulties of existing approaches, a promising avenue of research focuses on QOS options that lie somewhere between diff-serv and int-serv. The goal of such approaches is to provide finer granularity and stronger guarantees than are provided by diff-serv while avoiding the scaling and administrative problems associated with int-serv's per-application reservations. One such approach, which is being pursued in the Integrated Services over Specific Link Layers working group of the IETF, combines the end-to-end service definitions and signaling of int-serv with the scalable queuing and classification techniques of diff-serv.¹²

Another approach, referred to as virtual overlay networks (VONs), uses the Internet to support the creation of isolated networks that would link multiple participants and offer desired levels of QOS, including some security and availability features (Birman, 1999). This approach would require routers to partition packet flows according to tags on the packets called flow identifiers. This process, in effect, allows the router to allocate a predetermined portion of its capabilities to particular tagged flows. Traffic within a tagged flow would compete with other packets on the same VON but not with traffic from other flows. An individual user (e.g., a hospital) could attempt to create multiple VONs to serve different applications so that each network could connect to different end points and offer different levels of service. Substantial research would be required before a VON could be implemented. Among the open questions are how to specify properties of an overlay network, how to dynamically administer resources on routers associated with an overlay network, how to avoid scaling issues as the number of overlays becomes large, and how to rapidly classify large numbers of flows.

Quality of Service Policy

Because QOS typically involves providing improved service to some sets of packets at the expense of others, the deployment of QOS technolo-

gies requires a supporting policy infrastructure. In some applications, it is acceptable for QOS assurance to be lost for some short period of time, provided that such lapses occur infrequently. In other applications, however, the QOS guarantee must be met at all times, unless the network has become completely partitioned by failures. Work is needed to develop means of providing solid guarantees of QOS for critical information. Such mechanisms must scale well enough to be deployable in the Internet and will involve matters of policy (whose traffic deserves higher priority) as well as of technology.

In the int-serv environment, QOS policy is required to answer questions such as whether a particular request for resources should be admitted to the network, or whether the request should preempt an existing guarantee. In the former case, a decision to admit a reservation request might be based on some credentials provided by the requesting organization. For example, if a particular health care organization has paid for a certain level of service from its ISP and the request carries a certificate proving that it originates from that organization, then the request is admitted. In the latter case, a request might contain information identifying it as critically important (e.g., urgent patient monitoring information) so that it could preempt, say, a standard telephone call that previously had reserved resources.

It is difficult to predict all the possible scenarios in which policy information might play a role in the allocation of QOS. The RSVP provides a flexible mechanism by which policy-related data (e.g., a certificate identifying a user, institution, or application) can be carried with a request. The Common Open Policy Service protocol has been defined to enable routers processing RSVP requests to exchange policy data with policy servers, which are devices that store policy information, such as the types of request that are allowed from a certain institution and the preemption priority of certain applications. Policy decisions are likely to be complex, because of the nature of health care and the number of stakeholders involved in decision making. Accordingly, the design of policy servers, which are responsible for storing policy data and making policy decisions, would benefit from the input of the health care community.

Policy also has a role in a diff-serv environment. For example, if an institution has an agreement with an ISP that it may transmit packets at a rate of up to 10 Mbps and will receive some sort of premium service, then the question of exactly which packets get treated as premium and which as standard is one of policy. The institution may wish to treat e-mail traffic as standard and use its allocation of premium traffic for more time-critical applications. There may be some cases in which data from the same application will be marked as either premium or standard, depending on other criteria. For example, a mammogram that will not be read by

the remote radiologist until tomorrow might safely be sent by best-effort service, whereas one that is to be read while the patient remains in the examination room could be sent by premium service. Mechanisms for enforcing policy in a diff-serv environment are currently being defined at the IETF.

Multicast

Concurrent with efforts to implement QOS mechanisms for the Internet, attempts are under way to deploy multicast capability, which provides a means to make more efficient use of available bandwidth to simultaneously distribute information from one user to a number of specific recipients. Multicast stands in contrast to today's unicast delivery model, in which users communicate on a one-to-one basis, and to the broadcast model of radio and television, in which a single transmitter sends information out to a large number of unspecified recipients. Multicast makes large-scale multiparty conferencing possible in a way that makes efficient use of network bandwidth. It also provides an efficient mechanism for the distribution of streaming media to many recipients concurrently. At the same time, multicast presents new challenges regarding suitable pricing schemes and ways of protecting ISP networks from potential abuse. A related trend is the development of reliable multicast, which attempts to enable the reliable delivery of data from one source to many destinations, even in the presence of occasional packet loss in the network. A typical application is the timely delivery of financial information to many recipients. This technology is receiving a great deal of attention in the research and development communities and is likely to become mature within the next few years.

Both multicast and reliable multicast are likely to be useful in a range of applications, including health. Multicast technologies could be used to provide continuing medical education online through the real-time transmission of lectures over the Internet. It would allow for teleconsultations among geographically dispersed public health officials responding to a perceived public health hazard or bioterrorist attack. Multicast also could facilitate collaborative consultations among physicians, medical specialists, and patients.

Health care applications of multicast may emphasize different design and implementation features than would applications in other domains. For example, users in the health arena may be unlikely to create large multicast groups consisting of a single primary transmitter of information and millions of receivers, an approach more suited to the entertainment industry. Likewise, health organizations might not adopt the model of defense simulations, in which thousands of participants send information

to each other on a regular basis to report on their location, speed, and orientation.¹³ Instead, a health application might involve large numbers of small multicast groups featuring the formation of dynamic memberships to link collaborating physicians. The technical considerations that go into designing multicast protocols that can support large numbers of small groups may differ from those that can support smaller numbers of large groups. Health care applications need to be considered soon to ensure that suitable multicast capabilities are developed.

SECURITY

Security is a top priority for health applications of the Internet. Whenever personal health information is transmitted across the Internet or stored in a device attached to the network, precautions must be taken to ensure that the information is (1) available to those who need it, (2) protected against those lacking proper credentials, and (3) not modified—either intentionally or unintentionally—in violation of established policies and procedures. These three requirements—referred to as availability, confidentiality, and integrity—are of concern in most health care applications of the Internet, whether they involve the transfer of personal medical records between health care providers or a provider and a plan administrator, video telemedicine consultations, the reporting of information in a home monitoring situation, or the use of remote equipment in a biomedicine experiment (in which the data may be considered proprietary before formal publication).

Strengthening system security to better protect personal health information entails costs. Computing resources are needed to implement the changes, convenience may be compromised if system users are required to perform added steps (such as typing in additional passwords), and additional employees may be needed to monitor controls and investigate alleged violations of confidentiality. Furthermore, poorly designed security controls may actually impede health care delivery in emergencies by preventing or slowing access to needed information. It is therefore essential that system designers balance the potential costs of security mechanisms against their intended benefits—a process that requires assessing anticipated threats to personal health information, however dynamic and diffuse those threats may be.

To date, malicious attempts to sabotage the availability or integrity of electronic health information have been rare.¹⁴ However, the confidentiality of electronic medical records has on some occasions been compromised by individuals such as health care providers or administrators who have legitimate access to some aspect of an electronic record. Indeed, a previous study by the Computer Science and Telecommunications Board

(1997a) found that the most significant threats to patient privacy stem not from violations of established confidentiality policies or security practices but from the routine sharing of patient health information among care providers, administrators, public health officials, pharmacy benefits managers, direct marketers, and the like—most of which occurs without a patient's knowledge or consent. The effects of such compromises on the patient or consumer can range from embarrassment to loss of employment or loss of health insurance. In general, such concerns cannot be addressed by the use of security technologies but are the province instead of organizational policies and government legislation that set forth acceptable information-sharing practices.

Nevertheless, security technologies are increasingly important in an Internet environment. Health organizations have tended to rely on trust among health professionals to maintain the confidentiality of personal health information and have favored broad access to information (with some form of review of accesses) over strict controls. Connection to the larger, public Internet will require a new strategy. Health organizations contemplating the Internet as a source of networking infrastructure for critical transactions will need to be assured that they are protected against security risks while making their systems and data available to those who need them. The increasing interconnection of devices that monitor patients and devices that deliver treatment to them increases the possibility and potential consequences of attacks, particularly if the interconnections traverse the Internet. The very act of connecting health information systems to the Internet introduces a number of new vulnerabilities that malicious attackers can exploit. For example, data transmitted across the network can be intercepted and interpreted if it is not encrypted properly. Executable code downloaded from remote sites (or embedded in seemingly innocuous e-mail messages) can alter or destroy data contained in information systems. Denial-of-service attacks can limit system availability.

Elements of Security

To ensure availability, confidentiality, and data integrity, a system requires a number of supporting security functions, including the following:

- *Authentication* mechanisms to verify that people or systems are who they purport to be;
- *Access controls* to ensure that authorized entities can access and/or manipulate protected resources only in accordance with a predetermined set of rules and privileges;
- *Encryption* to protect data by scrambling it so that it cannot be read or interpreted by those without the proper decryption key;

- *Perimeter control* to manage interconnections between an organization's internal network and external networks (such as the Internet);
- *Attribution/nonrepudiation* to ensure that actions taken (such as sending or receiving a message) are reliably traceable and cannot be denied; and
- *Resistance* to denial-of-service attacks.

These functions are not necessarily performed by the network itself, but the need for them is heightened by the interconnection of information resources to a network, which can expand the scope of potential threats to information systems. In fact, the most appropriate place to perform these functions varies. For example, encryption by end hosts may be used to protect the confidentiality and check the integrity of application data. Other functions, such as authentication of routing updates, need to be supported by network elements such as routers. Security features operate at different network layers (Box 3.1), a variable that affects functionality.

BOX 3.1

The Open Systems Interconnection Model of Layered Networks

Packet-switched networks like the Internet are often described as using the Open Systems Interconnection model, which defines seven network layers:

1. The physical layer is concerned with the accurate transmission of bits over the physical media (fiber, cable, copper wire, etc.). The Synchronous Optical Network is an important physical-layer technology used for the long-distance transmission of both data and voice.
2. The data link layer is concerned with the delivery of packets between nodes. Point-to-Point Protocol is a well-known data link protocol.
3. The network layer handles the switching and forwarding of packets from one link to another, enabling large networks to be built. The most important network layer protocol is the Internet Protocol (IP). Devices that forward IP packets are known as routers, the basic building block of the Internet.
4. The transport layer handles end-to-end delivery of data between applications, often adding reliability. The Transmission Control Protocol, the best-known transport-layer protocol, provides reliable delivery of data for most Internet applications.
5. The session layer provides a variety of control functions used primarily by multimedia applications. Many applications bypass this layer.
6. The presentation layer handles issues of data formatting and standard representations of multimedia. Many applications bypass this layer.
7. The application layer provides application-specific functionality. Examples include the Hypertext Transfer Protocol, which is used for moving data across the World Wide Web, and the Simple Mail Transfer Protocol, the basis of electronic mail.

The implementation of security requires a mix of technological solutions and institutional policies and procedures. Institutional policies define the rules to be enforced; these rules then are allocated to administrative procedures, physical security measures, and technology. For example, organizations cannot simply deploy access control technologies to limit access to online data without first establishing rules to determine which users have the authority to view and alter information and under what conditions. In addition to implementing technologies such as passwords and smart cards to authenticate users, organizations also need to institute procedures for issuing, changing, and revoking passwords, and policies must be in place for disciplining offenders.

Today's Internet provides no security capabilities at the network level. The Internet was originally designed to facilitate information exchanges among mutually trusting entities (such as collaborating researchers), so information is encapsulated into packets that are passed through the network from node to node without encryption. Software programs called "sniffers" can be run on any node through which packets pass and can scan the contents of a message—even if the message contains sensitive health information or a user's password. As devices are attached to the Internet, their vulnerabilities tend to become more accessible. For example, weaknesses in operating systems can be exploited quickly over the Internet by individuals who wish to gain unauthorized access to resources. The large installed base of operating systems with known security flaws means that a large number of end points on the Internet are potentially vulnerable to attack.¹⁵ These shortcomings will become more significant as the demand for Internet-mediated health transactions grows and as the number of potential users increases.¹⁶

Several efforts are under way to improve the security of the Internet. Work continues on firewalls, which attempt to limit and control Internet-based access to an organization's computing resources and, hence, support the objectives of confidentiality, integrity, and access. In addition, new protocols and standards are being developed that will authenticate end points and provide greater confidentiality of information transmitted across the network. These advances promise to provide greatly increased security across the Internet, but continued work will be needed to ensure their deployment—and the deployment of the complementary capabilities needed to ensure their use. Primary among these complementary tools are certificate authorities (described below), which will help validate the identity of end users and which pose a number of technical and organizational challenges.

Firewalls

A firewall is a device that isolates one part of a network from another, typically isolating an organization's trusted internal network from an untrusted outside network, such as the Internet. The function of a firewall is to limit access to the organization's network from the outside by allowing only presumably safe traffic to pass through. Firewalls do so using one of four basic designs (Box 3.2) operating at three different networking layers. (This is significant because firewalls cannot prevent attacks launched at layers higher than those at which the firewall operates.) Typically, firewalls block all traffic from outside the organization's network by default and then are configured to allow specific, limited access. For example, a firewall might be configured to allow e-mail messages from outside the corporate network to pass through as long as they are destined to the appropriate mail server. Similarly, a firewall might allow access to an organization's public Web server but not to other Web servers

BOX 3.2 Four Basic Types of Firewall

Packet filters operate at the network layer, determining whether to pass or block individual packets based on criteria such as source address, destination address, and service requested. Pass/block criteria are not modified dynamically in response to the content of previous messages. Not all desired rules can be included in packet filters because the decision information (e.g., whether the message is a request or a response) may not be contained in the packet itself.

Circuit relays operate above the transport layer and pass or block entire conversations without examining the content of packets. They are generally considered more secure than packet filters but tend to require changes in application programs or user behavior, a characteristic that makes them unsuitable for some environments.

Application gateways operate at the application layer and typically use a separate program (or proxy) for each application so that different precautions can be taken for each one. A mail gateway, for example, might rewrite header lines to eliminate references to internal machines and keep a log of senders and recipients. Application gateways are well suited to sites that must authenticate outgoing messages and are considered highly secure.

Dynamic packet filters combine the packet filter and application gateway functions. Most packets are accepted or rejected based on information in individual incoming packets, but some packets cause modification of the rules that will be applied to subsequent packets.

SOURCE: Based on material in CSTB (1999).

that are used for proprietary, internal information. Thus, a company can maintain a public presence on the Internet without opening up its entire network to unlimited access by outsiders.

Despite their popularity, firewalls have several limitations. One problem is that they must be deployed at all places where an organization's network connects to the outside world. Although such networks often are designed to limit such external connectivity to one or two points, it is difficult to prevent unauthorized connectivity at other points. Individual users may establish dial-up modem connections to the Internet without the knowledge of the network administrators, creating a back door into the corporate network. If the network can be compromised at such a point, then firewalls provide no protection. Malicious intruders can gain access to the organization's network and the computers attached to it and may even acquire control of desktop computers.¹⁷ In addition, because each of the specific filtering functions of a firewall must be configured, the issue of correctness of configuration is critical; the mere presence of a firewall is not an assurance that any given set of protections has been established effectively.

Even properly configured firewalls in networks without back doors have limitations. First, many firewalls are constructed as applications sitting atop a standard operating system, so they are vulnerable to attacks against that system. A few firewalls use custom-built operating systems that avoid the vulnerabilities inherent in other systems, but these custom systems can introduce new vulnerabilities that have not been detected and remedied through the usual process of widespread use. Second, as noted earlier, a firewall cannot prevent attacks launched at layers higher than those at which the firewall operates. For example, a packet-filtering firewall cannot protect against attacks conveyed as e-mail attachments because it cannot understand attachment types and their risks. Third, firewall function is limited by the use of end-to-end cryptography. Firewalls cannot examine any fields in packets that have been encrypted or modify packets that must be cryptographically authenticated. Hence, many systems tend to decrypt at the firewall or use a layered approach in which some encryption or decryption takes place at the firewall and some inside it.

Fourth, firewalls tend to limit the external electronic communications of those behind the firewall. By filtering out (i.e., blocking) incoming messages that do not meet the specified access provisions, firewalls prevent internal users from receiving messages from some sources. Such a trade-off is inherent in the design of a firewall, which must use some fairly static rules to filter incoming messages. Fifth, firewalls increasingly are challenged by the advent of higher-speed Internet connections, which require the examination of packets at an ever-increasing rate. Sixth, and

perhaps most importantly, firewalls are effective only against external threats to an organization's data networks; they do not address threats posed by internal users who may intentionally or unintentionally violate security and confidentiality policies. Organizations that focus their security efforts too narrowly on external threats may not adequately protect themselves from insiders. In many industries, including health care, the insider threat historically has been of greater concern.

Security Protocols

As the scope and use of Internet applications grow, increased attention has been devoted to the development of protocols for user authentication and protection of messaging traffic. The protocols available today operate at a number of layers in the network. Whereas early protocols used in the military tended to provide link-level security, more recent protocols, such as Internet Protocol Security (IPSec), Transport Layer Security (TLS), and Pretty Good Privacy (PGP), tend to operate at the network, transport, and application levels, respectively. All of these are proposed IETF standards and all can be expected to see more widespread application. All use encryption as the basis for authentication and confidentiality.

Encryption Technologies

Encryption technologies generally are classified as either symmetric key systems (also called private key cryptography) or asymmetric key systems (also called public key cryptography), both of which may be used during the course of a single session between communicating parties. The two technologies differ in a number of respects, including the time it takes to encrypt and decrypt messages and ease of administration. As a result, they tend to be best suited to different types of applications. Symmetric encryption, for example, tends to work better when communicating parties have a preexisting relationship. Asymmetric systems, in contrast, work well between parties that have not communicated before, as in many electronic commerce applications; however, revocation of credentials can be more difficult than with symmetric encryption systems.

Symmetric encryption uses a single key to encrypt and decrypt data. Parties wishing to exchange information securely must ensure that they both have access to the key, meaning that mechanisms must be in place for distributing keys to pairs of users before secure transmissions can begin. This process sometimes involves physically distributing disks containing the key, but this approach is slow and can be used only if the communicating parties can be identified in advance. A number of mecha-

nisms, including online key distribution centers,¹⁸ have been established to facilitate the electronic distribution of symmetric keys on an as-needed basis. Many such mechanisms rely on asymmetric key cryptography to authenticate parties and distribute keys before they exchange sensitive information. In either case, care must be taken to ensure that keys are not divulged to others, are changed periodically, and are revoked as needed.

Asymmetric (public key) cryptography is an important component of Internet security because it provides a way in which strangers can establish the necessary set of shared information to support authentication and encryption. This means it could be useful in exchanges of patient health records between two unaffiliated hospitals. In an asymmetric key system, a given user (e.g., an individual or corporation) has a pair of keys, one of which is private (known only to the key owner) and one of which is public and may be shared with anyone (or posted in a directory). Data encrypted using the private key can only be decrypted using the public key; and data encrypted with the public key can only be decrypted with the private key. Private keys also can be used to support authentication in the following way: If a recipient can decrypt a message using the sender's public key, then the original message could have come only from the holder of the private key. Conversely, a user wishing to send a message to be read only by its intended recipient can guarantee confidentiality by using the recipient's public key to encrypt the message; the resulting data will be readable only by the owner of the private key. Asymmetric encryption systems require keys about 10 times longer than those of symmetric encryption systems and run considerably more slowly. For this reason, asymmetric cryptography is used only to authenticate the participants in an information exchange and to distribute symmetric keys, which then are used to protect messages exchanged during the remainder of the session.

Distribution of Encryption Keys

A major challenge in using asymmetric cryptography is the distribution of public keys. A person who wishes to use a certain public key for either encryption or authentication needs to know for certain that the key belongs to the appropriate entity; otherwise, authentication is not possible, and encrypted data may be read by an unintended recipient. This problem usually is handled by having some sort of certification authority (CA) issue certificates—digitally signed documents that state “public key X belongs to entity Y.” A certificate can be trusted only if the reader knows the public key for the CA. Thus, if a user obtains a single authoritative public key (that of the CA), any other entity that needs to prove

ownership of a certain public key to that user can do so by providing a certificate for the key.

One difficulty involved in issuing certificates is scale: Large numbers of certificates need to be given to all potential participants in various types of transactions. Health care would benefit greatly from the issuance of certificates to all health care consumers—the entire U.S. population. Such widespread deployment of certificates would enable consumers to gain authenticated access to sensitive health care information (e.g., lab test results) from any provider. Thus, a public key infrastructure (PKI) suitable for health care would need to have the capacity to operate on a scale of hundreds of millions of users. Good initial progress has been made in issuing certificates to Internet vendors of goods and services, but the process has not been extended to individual consumers. One way to make the process more scalable is to arrange CAs in a hierarchy, with the root CA certifying lower-level CAs that issue certificates to even lower-level entities, and so on down to the level at which certificates for individual users are issued. In this case, the process of proving that one is the owner of a certain public key may involve providing a hierarchical chain of certificates that leads from the root down to the owner of the key. The term PKI often is applied to the general problem of distributing keys and certificates to a large population in a scalable way.

The task of building a hierarchy of CAs presents its own challenges.¹⁹ Most notable of these is the establishment of consistent policies for issuing certificates. One CA, for example, might give certificates to anyone who requests one by e-mail, whereas another might require recipients to sign an affidavit and provide a birth certificate and passport before providing a certificate. In the system used to support PGP for secure e-mail, any user with a public key can issue a certificate, creating a “web of trust” among particular groups of users rather than a strict hierarchy. PGP also allows users to decide how well they trust a certain certificate based on who it is from and how many corroborating certificates the user holds. In a hierarchical model, if just one CA in the hierarchy uses weak procedures to establish an individual’s identity (e.g., issuing certificates to individuals without seeing them in person with positive proof of their identity), then the whole certification system is compromised, because anyone might be able to get a fake certificate from that CA. Such policy differences make it difficult for users and applications to interpret certificates and determine which ones meet their criteria of proof of identity, effectively undermining the goal of enabling broad deployment of asymmetric cryptography. One way to address this issue (proposed in the Privacy Enhanced Mail architecture but not widely deployed) is to have the top-level CA certify the policies used by lower-level CAs, so that weaker policies can be readily identified. However, it is clear that variability of policies among CAs will

add complexity to the system and is likely to weaken the overall level of trust that can be placed in public keys.

Another challenge is certificate revocation, which may be required if the owner of a public/private key pair believes that the private key has been compromised. Revocation typically is handled by a combination of expiration dates on certificates and revocation lists, which are published lists of certificates that are to be considered invalid, signed by the issuing CA. To be sure that a certificate is valid, therefore, it is essential to have the most up-to-date revocation list from the CA that issued the certificate. The use of expiration dates on all certificates ensures that revocation lists do not grow infinitely large, but it requires all users to undergo periodic recertification, thus increasing the workload on the CAs. An effective PKI must have an efficient means for disseminating up-to-date revocation lists.

The certificate model also raises issues of personal privacy. The CA models developed to date bind a public key to a particular identity, whether an organization or individual. The use of that key, therefore, can be linked to the activities of that organization or individual. Some work has been initiated on key-centric systems, in which names associated with public keys are not bound to a particular individual but have, rather, only local significance for the convenience of users. The Simple Public Key Infrastructure working group of the IETF is attempting to develop an Internet standard that incorporates these ideas, but no related commercial products are available.

Internet Protocol Security

Internet Protocol Security is an architecture and set of standards that provides a variety of services, such as encryption and authentication of IP packets, at the network layer (Kent and Atkinson, 1998a,b,c). IPSec can protect traffic across any LAN or wide-area-network technology and can be terminated at end systems or security gateways (e.g., firewalls). Now being standardized at the IETF, IPSec has been deployed initially in virtual private networks (VPNs) that use the Internet as the underlying medium but establish an encrypted tunnel across it. An encrypted tunnel can be created between a pair of IPSec gateways, which might be located, for example, at two geographically separated offices of a single company. Each gateway encrypts the data and sends them to the other gateway. The receiving gateway then decrypts the data before passing them on to the final recipient at the site. Because the data are encrypted using a key known only to the two gateways, the message cannot be read by anyone else while crossing the Internet. Furthermore, the receiving gateway can authenticate the data as having come from the sending site and not some

other source on the Internet. An attractive feature of this technology is that a configuration of one device at each site can protect traffic between those sites against eavesdropping and corruption. Protection does not extend beyond the ends of the tunnel.

Virtual private networks have several limitations. First, VPNs based on tunnels do not scale well. Increasing the number of participants in a VPN increases the number of points at which the system can be compromised and makes the process of key management much more difficult. Second, IPsec tunnels frequently require a priori knowledge of where connectivity will be required, because the gateways must be configured with appropriate keys (in the absence of an automated key management infrastructure) and routing information may need to be modified to force traffic to use the appropriate tunnel. This characteristic makes IPsec a viable alternative for information exchanges between organizations with well-established relationships but less effective for unexpected or transitory exchanges of information. In health care, VPNs might be useful for secure communications among elements of an integrated delivery system or between health care providers and the Health Care Financing Administration (HCFA), which processes Medicare claims, but they could not readily support exchanges of patient records between unaffiliated hospitals in an emergency situation.

Third, IPsec tunnels do not necessarily protect data during the entire transit between sender and receiver. Many enterprises encrypt data only between VPN gateways, which are devices that sit at the boundary between the public Internet and the corporate network and encrypt data traversing the Internet. This configuration simplifies the key management problem because it requires encryption keys for the gateways only—not for all the computers connected to them—and averts the need to modify end users' machines. At the same time, it leaves data unencrypted as they pass between the users' computers and their respective gateways, meaning that someone with physical access to the data lines could, in theory, intercept and read the messages. Some users note that, in practice, the expectation of security in a VPN can lull them into failing to encrypt data passed across it. Often, unencrypted data are sent over a LAN until reaching a VPN gateway; in other cases, as with frame relay, the data are not encrypted and are subject to misrouting.

Transport Layer Security

An alternative mechanism for providing encryption and authentication across the Internet is transport layer security, which is widely used across the Internet in the form of the Secure Socket Layer (SSL) system (Dierks and Allen, 1999). This technology is widely used for transmitting

sensitive information (such as credit card numbers) between Web browsers and servers. Transport layer security uses asymmetric encryption to authenticate the server (and, optionally, the client) and symmetric encryption to protect communications between the end user and the Web site. An organization that requires encryption for transactions processed through its Web site obtains a certificate from a CA (e.g., Verisign, CyberTrust, CertCo, or DST). When a user connects to the organization's Web site, the site provides its certificate. Modern Web browsers come equipped with the public keys of the major CAs so that the browser can verify the public key of the Web site and thereby authenticate it.

Secure Socket Layer can support encryption in both directions (to and from the Web site), but as commonly used today it provides authentication of only the host organization's Web site.²⁰ The bidirectional authentication function is generally not invoked; it would require the client (as well as the server) to have a certificate, which is not generally the case. As a result, users can readily verify the identity of the organization with which they are communicating, but the server site typically cannot verify the identity of the person using asymmetric encryption techniques. Existing transport layer security, therefore, has been used for credit card transactions in which authentication of the user is not performed cryptographically but rather by some other means (e.g., verification of card number, expiration date, and billing address or a name and password) in which the misidentification of a client has a known cost (e.g., unrecoverable accounts receivable). In the absence of client certificates, SSL is not well-suited to applications in which the costs of misrepresentation of identity cannot be quantified and in which secure passwords are not considered sufficient for authentication. In the health domain, such applications might include the delivery of health care via the Internet or real-time patient monitoring.

Although an individual can obtain a certificate in much the same way that a corporation running a Web site does, the issuance of certificates to many individuals requires methods that are more scalable than those available today and that ensure greater compatibility in the criteria used to issue them.²¹ Organizations that maintain a presence on the Internet generally find it possible to obtain a certified public key from one of the small number of CAs, but there is no infrastructure in place to enable individuals to obtain certificates on a large scale. Such problems are likely to inhibit the use of applications in which authentication of the consumer is as important as authentication of the vendor.

Several initiatives are under way to deploy CAs for health applications. In October 1999, Intel announced that it would work with the American Medical Association to provide digital certificates to physicians in the hope of enabling doctors to transmit information such as test results

to patients and other health care workers. Healthon/WebMD has agreed to provide the product to physicians, while other health Web sites, including WellMed and Franklin Health, intend to provide access to the product on the consumer side (Reuters, 1999).²² The Robert Wood Johnson Foundation has also awarded a \$2.5 million grant for a five-state HealthKey initiative that will explore ways to facilitate electronic exchanges of information among companies in the health sector while protecting the confidentiality of the data.²³ PKI is one of the solutions being considered.

The financial industry has found SSL suitable for many of its consumer-oriented activities, such as online banking and stock trading, but the limitations of this system in the health domain are apparent. At present, financial institutions rely primarily on certificates for server authentication and on passwords for client authentication. As users obtain more online accounts, they need more passwords. To help themselves remember all these passwords, users take many steps that reduce security, choosing passwords that are easy to recall (and easy for others to guess), reusing passwords already in place for other accounts (thus making all of their accounts vulnerable to the compromise of a single password), and writing down passwords (making them easier for others to find). Furthermore, financial organizations usually issue passwords by mailing them to the address of record for the account. This process introduces significant delay in password assignment, making the process inappropriate for health applications in which an emergency room physician may need access to a patient record at a remote hospital. The distribution process is also limited in that a mailed password can easily be intercepted by the wrong member of a household—a vulnerability that may have more serious consequences with health data than with financial information. Hence, the trade-offs that are acceptable for applications in the financial sector do not seem suitable for health care.

Confidential E-mail

Software to encrypt e-mail has been available for many years, but its use has been limited because, until recently, it had not been integrated well into standard e-mail applications. PGP is a well-developed collection of encryption software that is commercially supported and available free of charge for noncommercial use (Zimmerman, 1994). Although it supports a variety of functions, it is used most often to digitally sign and/or encrypt e-mail. To send an encrypted e-mail message, the sender must gain access to the public keys held by the intended recipient(s), who must have obtained PGP public/private key pairs in advance.²⁴ An e-mail message to be sent to several recipients is encrypted using a symmetric algorithm with a new (secret) session key. This key is then encrypted

using the public key of each of the recipients, and the results of this encryption are included in the header of the transmitted message. Software run by each recipient uses the recipient's private key to retrieve the session key.

The emerging standard for commercial e-mail systems is Secure/Multipurpose Internet Mail Extension (S/MIME). MIME is an extension to standard e-mail formats that supports the transmission of data, such as video and audio, that are not usually represented as ASCII text. S/MIME supports the transmission of signed and/or encrypted data in e-mail messages in much the same way that PGP does. However, S/MIME uses certificates based on an international standard called X.509 version 3 and generally embraces the hierarchical model for CAs, as opposed to PGP's webs of trust.

Access Controls

Access controls are an important element of computer systems security. They consist of a range of techniques for controlling the capability of active entities (such as computers, processes, or users) to use passive entities (such as computers, files, directories, or memory). For example, access controls can prevent certain users from viewing a particular database, modifying information that they can view, or running certain programs and functions. Because they operate at the level of bits, access controls can be used to permit users to access portions of an encrypted file while still protecting the overall confidentiality of the information. Access controls can operate at virtually all layers in a networked application—from the physical layers defining the communications medium itself through the application layer consisting of software programs—and can extend access privileges based on various characteristics, such as the user's identity or role in the organization (Box 3.3).

Health care poses an especially difficult challenge with respect to access controls because of the difficulty of balancing the need to protect confidential information from unnecessary dissemination against the need to ensure adequate access to enable the provision of care. It is often difficult in a clinical setting to determine a priori who should or should not be granted access to a particular patient's medical record. Many clinicians may interact with a patient during the course of treatment, and all are likely to need access to the medical record. During a typical hospital stay, numerous health care workers, including dietitians and pharmacists who must consider potential drug and food interactions, will need access to the record. A study by the Institute of Medicine (1997) identified 33 different types of individual users of patient medical records—including care providers, health plan administrators, researchers, educators,

BOX 3.3 Types of Access Control

Access controls can be based on labels, identity, capability, or location. In label-based access control, access decisions are made by comparing a label stored in the data with a label stored with the subject attempting to access the data—the subject's identity is irrelevant. In identity-based control, access decisions are made by comparing the identity of the subject with a list of individual users and groups of users who are allowed to access the object. As such, identity-based access controls depend on accurate authentication of the identity of the user.¹ In capability-based control, access decisions are made by checking a so-called permission ticket attached to the object. Location-based access controls are typically used in conjunction with the other three forms of access control to allow different levels of access depending on the user's physical location. For example, a terminal in a patient's hospital room might be configured to allow *any* authenticated care provider to access that patient's record, whereas remote access might allow providers to view records of only their own patients.²

These approaches can be combined to produce different forms of access control. In health care, many access controls are role-based, using a combination of label, identity, and capability-based controls. Role-based systems grant permission for a particular individual to access particular data or capabilities based on job functions. For example, administrative workers might be granted access to billing records but not to clinical databases containing medical records of patients. Physicians and nurses might both be granted access to the clinical database, but only to records of patients directly under their care. Similarly, distinctions might be made with regard to which care providers can view information and which ones can enter treatment orders and/or prescriptions. Roles can be defined broadly (e.g., physician, nurse, administrator) or narrowly (e.g., cardiologist, neurologist, intensive-care nurse, pediatric nurse, human resources administrator, billing clerk), depending on the granularity of the meaningful distinctions among workers in a particular organization and the level of differentiation required.

¹Some types of smart cards that are used for authentication contain information on user access privileges, providing a physical example of the close linkages between these two functions.

²Virtual private networks and firewalls can also be considered forms of high-level access control that limit user access to a network or a portion thereof. A firewall might use identity-based controls (such as the user's IP address), whereas a VPN might use a combination of identity-based control (using the identity contained in a certificate) and capability-based control (based on the possession of an encryption key).

accreditation boards, and policy makers—and 34 representative types of institutions. Each of these users needs different information, and their access privileges could be markedly different, compounding the difficulty of developing effective access controls and confidentiality policies (upon which access controls are based).

Also important are tools that support rapid access for authorized individuals. In some cases, military or commercial secrets may require protection over a span of decades or more (e.g., the formula for Coca-Cola or technical details of nuclear weapons). Patient records can contain data equally sensitive from the standpoint of the individuals concerned, requiring lifetime protection, yet they also must be available in emergency situations to a much larger class of people. Furthermore, many different users may have occasional permission to view or modify portions of records. For example, a major public benefit of increased automation of health care information systems should be the availability of the data for research studies. Unless the subjects of the studies can be assured that their information will be protected, through either the clear enforcement of appropriate release policies or the "anonymization" of the records, this benefit will be limited.

Because they need to ensure adequate access to health information in emergencies, many health care organizations either (1) routinely give physicians access to information on all patients within the organization's care or (2) routinely give them access to information on only those patients directly under the physician's care but provide emergency overrides to allow access to the records of other physicians' patients if needed. The organizations then rely on the use of audit trails to review accesses after the fact. Such audits are intended to deter the abuse of access privileges, but their effectiveness depends on the availability of tools for automatically detecting anomalous accesses to medical records. Some work is under way on such tools, but they are still in their infancy. More work is needed to develop sophisticated audit analysis tools that take into account expected usage patterns and auxiliary information, such as appointment schedules and referral orders, to more accurately identify potential violations of confidentiality policies.

The introduction of networking (e.g., the Internet) compounds the access control problem by facilitating the exchange of electronic medical records among different users of health information. Standard access controls can limit the use of information within a single organization but cannot control how information is used once transferred to another organization. Some work is under way to develop technologies that could control such secondary uses. Cryptographic envelopes and associated rights-management languages, such as those developed by IBM, Xerox, and InterTrust Technologies, enable content owners to send data in an encrypted form to users outside their organizations and to specify the actions that different users can undertake with protected data, perhaps allowing them to view it, for example, but not to print or redistribute it. Such tools may also allow auditing accesses of health records across institutional boundaries. Cryptographic envelopes are still relatively new and

were developed primarily to protect copyrighted material and ensure proper payment for viewing or copying it. Additional work is needed to extend this model into the health sector and develop rules for sharing information among different types of health care organizations.

Network Availability

Network availability is another essential element of information systems security.²⁵ Availability is the probability that the network (i.e., the Internet) will be operational at a particular point in time and accessible to those who need it. High availability is a key requirement for mission-critical and time-critical applications of the Internet, including many in health care. If the availability of the Internet is uncertain, then health care providers cannot rely on it for the provision of remote patient care or access to electronic medical records in the emergency room, although they may still be able to use it (with some degree of frustration) to submit bills and allow consumers to select physicians.

Network availability can be compromised by a number of factors, including hardware or software failures, operator errors, malicious attacks, or environmental disruptions (e.g., lightning strikes, backhoes cutting fiber-optic cable) that cause particular links or entire sectors of the network to fail. Availability is closely related to both QOS and security in that the failure of a link connecting two routers across the Internet can affect the capability of an ISP to meet its QOS guarantees, and security measures that protect a network from malicious attacks (whether actual network intrusions or denial-of-service attacks) can help ensure its availability. In addition, security measures such as ensuring software correctness and ensuring software integrity (avoiding viruses, worms, Trojan horses, etc.) will address some issues of availability. However, such mechanisms do not necessarily protect a network from accidental operator errors or physical damage or ensure network survival in the face of hostile attacks.

By virtue of its design, the Internet is reasonably resistant to many forms of failure. Its web-like interconnections among routers ensure the existence of multiple routes for channeling messages across the network. If one link fails, then traffic can be routed along an alternative pathway. In most cases, the network can converge on a new path that avoids the failed link within a matter of seconds, providing sufficient reliability for many, if not most, Internet applications. Nevertheless, service outages do occur. ISPs and Web hosting facilities operate their network infrastructures with cutting-edge technology but, despite the adequacy of the equipment and redundant fiber links, outages lasting for hours occur several times a year. The causes vary. Faulty routers can announce incorrect routes and

cause disturbances to propagate throughout a provider's network, for example, or an upgrade can cause unexpected problems when a network is large, despite extensive testing.

To counteract these problems, many end-user organizations maintain redundant links from different ISPs. To do so, an organization needs to run its own Internet routers and announce and manage its own routes across the Internet. The management overhead involved in deciding how to balance traffic between the links and when to switch all traffic to one link is significant and error-prone. Indeed, many of the difficulties that arise today are directly related to the complexity of this problem. The number of routes, the size of the resulting data structures, the inherently distributed nature of routing algorithms, and the constraints applied by administrative/business requirements all contribute to this complexity. A continued strong research effort is required to improve the reliability and performance of Internet routing protocols.

Additional efforts are needed to consider means of responding to disaster scenarios in which large portions of the network fail, resulting in major outages. Such disasters could be confined to the network, in which case mechanisms are needed for ensuring continued transmission of a variety of network traffic, or they could be more widespread—fires, earthquakes, or storms—and thus might call for ways of mobilizing health care resources despite widespread network outages. In both cases, mechanisms are needed to ensure adequate network availability for mission-critical applications and to handle high-priority traffic. Many policy issues would need to be addressed to help balance the networking needs of health care organizations against those of other critical communications. Additional work is needed in the area of survivability to ensure that the network can maintain or restore an acceptable level of performance during failure conditions by applying various restoration techniques.

Some work is under way in the Department of Defense to develop techniques for prioritizing network traffic in case of degradations that limit network capacity, but such work may need to be extended to consider the requirements of health care. As an example, in disaster situations, telephone service providers can block incoming calls to the affected region at their source (by providing a busy signal) so that limited link capacity can be used for more urgent outgoing calls. Congestion control on the Internet is problematic because no mechanisms exist to manage traffic based on user, connection, source, or destination. Routers do not store information about users and connections, and doing so would require significant memory and management mechanisms and would also raise a host of privacy concerns that would need to be addressed. Some newer algorithms for congestion control have been designed to work at the IP

level, but more research is needed, especially in the area of defining and enforcing flexible and varied policies for congestion control (CSTB, 1999).

BROADBAND TECHNOLOGIES FOR THE LOCAL LOOP

Before the health community and health care consumers can benefit from future Internet applications, they must gain access to sufficient bandwidth in their local connections to ISPs to handle the anticipated traffic loads. Health care organizations that intend to transmit detailed radiographic images to remote specialists for near-real-time interpretation, for example, will need Internet connections capable of transmitting hundreds of kilobits per second, if not megabits per second. Biomedical research institutes conducting distributed simulations will also need high-bandwidth connections. Organizations will meet many of these needs by leasing communications lines with the needed capacity. Alternatively, some organizations that provide content over the Internet and expect high demand for their services may attempt to offload some of their functions to third parties that can acquire the needed capacity, although there may be limitations to this model in health applications (see Box 3.4).

BOX 3.4 Web Hosting for Health Applications?

Web hosting is one approach taken by some organizations to overcoming bottlenecks in local access technologies. This term refers to an arrangement in which an organization hands off the management of its Web server to another party, such as an Internet Service Provider, while continuing to provide the content. The responsibility for maintaining performance, reliability, redundancy, and network bandwidth rests with the hosting company, which potentially can reap economies of scale by performing the same service for many Web content providers. Companies that host Web sites can contract for large amounts of bandwidth and spread the costs over many customers. Web hosting is popular in many business settings because of the challenges involved in administering a Web server and the cost of the high-bandwidth connections that are often required between a Web server and the Internet.

Despite its advantages, Web hosting is not well suited to many health applications. First, the sensitivity of the information may preclude its transfer to a third party in unencrypted form. Second, Web hosting works well when the Web server provides content that may not change very rapidly. This condition may not hold true in an environment in which health care companies endeavor to provide services over the Internet.

TABLE 3.1 Technologies for Wireline Connections to the Internet

Technology ^a	Maximum Data Rate ^b
DS-0	56 kbps
DS-1 dial-up	56 kbps to 1.344 Mbps
DS-1 private line (T1)	1.544 Mbps
DS-3 private line (T3)	45 Mbps
OC-3	155 Mbps
OC-12	622 Mbps
OC-48	2.5 Gbps
OC-192	10 Gbps

^aDS, digital service; OC, optical carrier.

^bkbps, kilobits per second; Mbps, megabits per second; Gbps, gigabits per second.

SOURCE: B. Davie, Cisco Systems, presentation to the committee on September 9, 1999, Washington, D.C.

Many businesses already connect to the Internet over dedicated lines at speeds ranging from 1.5 Mbps (T1 lines) to 155 Mbps (OC-3 access). The various possibilities are listed in Table 3.1. A few organizations lease OC-12 lines capable of transmitting 622 Mbps, but these tend to be limited to a few high-profile Web sites that expect large amounts of traffic. Alternatively, large organizations can subscribe to services that provide high-bandwidth access over a shared medium. Frame relay, for example, is a packet-switched connection typically sold at speeds ranging from 56 kbps to 1.5 Mbps; packets from various subscribers are mixed together across the underlying links. Frame relay is less expensive than a dedicated line but introduces some uncertainty about instantaneous capacity as other organizations contend for bandwidth. Asynchronous transfer mode (ATM) is another service that is starting to supplement frame relay as a switched alternative to leased lines. In addition to offering higher speeds (e.g., 155 Mbps), ATM offers a wider range of QOS options than frame relay does and was designed to handle mixed voice, data, and video.

Although leased lines, frame relay, and ATM are viable alternatives for a variety of institutional users, they are generally too expensive for residential users and small businesses (such as private practitioners). Leased lines can cost from hundreds to thousands of dollars per month, depending on the bandwidth provided, and a 56-kbps frame relay connection can cost \$150 per month after installation. The overwhelming majority of residential users today connect to the Internet using a modem connected to a conventional telephone line. Such connectivity is almost universally available—accessible from any location with a telephone line—but is limited in bandwidth. The fastest modem connections today are capable of providing 56 kbps, but many residential users connect at

28.8 kbps or less. These low connection speeds are suitable for many of today's health applications, such as downloading online information and participating in chat groups, but they can be bottlenecks for large files (whether text, video, or audio) and real-time video.

Future health applications of the Internet may demand more bandwidth to residential and small institutional users than conventional modems operating over telephone lines can provide. For example, if online health records become more widely used and begin to contain more medical imagery (e.g., X rays, CT scans, MRIs), then greater bandwidth will be needed to download the images quickly. If rural health clinics begin sending radiographic images to remote specialists for near-real-time interpretation, then they will need significant bandwidth. More significantly, applications such as video-based teleconsultations and teleradiology will require as much bandwidth out of the home or small office (upstream to the ISP) as into it (downstream into the home or office), because images will be transmitted in both directions.

This requirement for symmetry in upstream and downstream bandwidth allocation represents a significant shift from most current consumer Internet applications, which assume the majority of information will flow from the Internet to the consumer. Two of the more popular technologies currently available for broadband access in the local loop—modems using cable television lines and digital subscriber line (DSL) technologies (Box 3.5)—allocate bandwidth asymmetrically, providing more bandwidth downstream than up (Table 3.2).²⁶ Cable modems, for example, enable users to receive data at speeds as high as 10 Mbps but transmit data at only 384 kbps. Bandwidth is shared with neighboring residences (up to several hundred), so the exact amount of bandwidth available to an individual at a given point in time depends on his or her level of activity.²⁷ Most deployments of DSL technology support up to 1.5 Mbps downstream and 768 kbps upstream. Unless high-quality videoconferencing or distributed simulation games become popular with consumers and drive the need for downstream video, the health sector could become a driver for this capability.

It is possible to reconfigure DSL symmetries with a technology available today, discrete multitone (DMT). But this can be done only if an upcoming standards decision for very-high-speed DSL (known as VDSL) favors the DMT approach. Some VDSL technologies provide data speeds up to 26 Mbps downstream and 3 Mbps upstream over copper wire at distances of up to 3,000 ft. The use of DMT makes it possible to achieve more symmetric bandwidth allocations of up to 10 Mbps in each direction at distances of up to 5,000 ft. If this approach is not standardized, then asymmetric technologies will be locked in much more strongly and diffi-

BOX 3.5
Technical Alternatives for Wireline Broadband Access

Cable modems allow residential users to connect to Internet service providers over a hybrid fiber-coaxial (HFC) network that is currently installed to deliver cable television signals. Fiber-optic lines run from the cable operator's central facility (or head-end) to local neighborhoods (or nodes), where they connect into the existing network of coaxial cables that extend to individual homes. One or more of the 80 to 120 channels carried over a typical HFC network can be allocated to data transfers for Internet access. Each channel is 6 megahertz (MHz) and offers 30 megabit per second data rates shared with other users in the local neighborhood. The head-end assigns time slots to users on the same network and achieves utilization rates of approximately 85 percent. All transmissions are encrypted between the head-end and user, and encryption keys change frequently. Uplink capacity in cable systems is typically limited to roughly 384 kilobits per second by channel spectrum. Cable systems transmit information to the home in channels spread across the 50 to 700 MHz spectrum (50 MHz is the equivalent of channel 2, which is where cable companies were required to start for compatibility reasons). Data must be sent upstream over the 5 to 42 MHz spectrum, which offers lower data rates.

Digital subscriber line (DSL) technology runs over the existing twisted-pair copper wire that is used for telephony, but at the central office the twisted pair is connected to a DSL modem rather than a conventional telephone switch. The system uses multilevel coding of information to obtain high bandwidth across limited lengths of wire. A technique called Discrete Multitone in effect divides the bandwidth of the telephone wire into 256 separate subchannels, each with 4 kHz capacity. Because many service providers are interested in supporting video-on-demand for entertainment purposes, most of these channels have been allocated to downstream capacity, resulting in a proliferation of asymmetric digital subscriber line offerings, with roughly 6 MHz of downstream capacity and 600 kHz of upstream capacity. The signal coding is an international standard.

SOURCES: Milo Medin, Excite@Home Corp., presentation to the committee on December 17, 1999, Stanford, Calif.; Hawley (1999).

TABLE 3.2 Wireline Technologies for Residential Broadband Access

Technology	Downstream Data Rate	Upstream Data Rate
Dial-up modem	Up to 56 kbps	Up to 56 kbps
Asynchronous digital subscriber line	Up to 10 Mbps; typically 1.5 Mbps	Up to 768 kbps; typically 384 kbps
Cable modem	Up to 10 Mbps	384 kbps

NOTE: kbps, kilobits per second; Mbps, megabits per second.

cult to dislodge in the future. Many companies have strong vested interests in existing asymmetric technologies and might resist the use of DMT.

Another drawback to cable and DSL networks is that they are not currently accessible from all locations within the United States. Cable systems pass through approximately 80 million U.S. homes but tend to be concentrated in densely populated, not rural, areas. They also tend to pass through residential neighborhoods instead of business districts (an artifact of the focus on entertainment applications), a pattern that may impede the use of these technologies by some health facilities. In addition, some older cable networks cannot support cable modems. The deployment of cable modem service is accelerating, but cable companies are expected to focus on upgrading their infrastructure in densely populated areas, where the greatest revenue can be realized from high-speed data services.

The availability of DSL services and the amount of accessible bandwidth are highly sensitive to the distance of a residence from the central office and to the quality of the copper wiring. Asymmetric DSL (ADSL) services typically can support data rates up to 1.5 Mbps downstream and 384 kbps upstream over twisted pair up to 18,000 ft, which would reach almost 80 percent of U.S. households, according to the ADSL Forum. For example, one DSL provider offers services at 384 and 768 kbps in metropolitan areas only and expects to be able to reach 65 percent of the residences in those areas. The remaining 35 percent of residences would not be accessible with the current technology because of either their distance from the central office or the low quality of the telephone lines. Deployment in remote or impoverished areas is not likely to proceed quickly.

Another means of Internet access is wireless technologies, which (at least theoretically) could be used virtually anywhere and also ease the provider's burden of laying down wires, fibers, and cable. Low-speed wireless services (e.g., approximately 30 kbps) are currently available in only a few parts of the country but are likely to become more widespread. High-speed wireless is also likely to become an alternative for connecting to the Internet, initially for businesses but perhaps also for consumers who are not well served by cable or DSL. Local multipoint distribution system (LMDS) technology uses high-frequency microwaves for two-way communications at data rates of up to 155 Mbps. It operates within areas (or cells) 2 to 5 miles in diameter. Performance is limited by rain and by the need to maintain a line of sight between the transmitting and receiving stations.

Satellite-based systems using either geostationary or low Earth orbit (LEO) satellites boast maximum transmission speeds twice as fast as LMDS, 3 to 6 times faster than cable, and up to 12 times faster than DSL. Such systems are likely to cost hundreds of dollars per month for service

plus \$500 to \$1,000 for the antenna (Skoro, 1999). Geosynchronous systems are limited by significant propagation delays (i.e., latency), which may preclude their use in some interactive applications; furthermore, their data-carrying capacity is distributed among a large number of users. LEO systems overcome some of the problems with delay, but the satellites move fast and have smaller coverage areas, meaning that large numbers of satellites are needed to provide global coverage and techniques are needed to manage the handoff of connections between satellites. High-power transmitters are needed to achieve high data rates, which implies large antennas and/or high-frequency operation. At higher frequencies, signals degrade more quickly in rain and other adverse weather conditions.

Overall, the deployment of broadband Internet services has been slow, albeit increasing, in the United States. Only about 1 percent of all U.S. households with Internet access had broadband connections in 1999 (Clark, 1999). A number of factors are at play, including technology, economics (both the cost of building broadband networks and the costs of service), and policy. Most U.S. households have yet to subscribe to broadband services because these connections are not offered in their geographic area, or because they are too expensive or not viewed as useful. The spotty coverage and high cost of high-bandwidth access technologies mean, unfortunately, that those who could benefit most from the health care applications of the Internet—such as people in rural areas with limited access to medical specialists—are the least likely to have high-speed Internet access.

Work in many areas, both technical and policy-related, will be required to enhance network access for health applications. In some cases, technical work will be pursued by the computing and communications industries without the participation of the health community. Even so, by voicing its needs, the health community will help ensure that they are met. In other cases, the requirements of the health care community may motivate research. Again, the articulation of specific needs will be necessary, and participation in research may be needed as well. The following section identifies several needs that are of particular interest to the health care community.

PRIVACY-ENHANCING TECHNOLOGIES

A popular cartoon depicts a dog sitting in front of a computer monitor and is captioned, "On the Internet, nobody knows you're a dog." At one level, this statement is true—an ordinary Internet user can choose a cryptic pseudonym or screen name so that a typical e-mail recipient or chat room participant cannot easily identify the individual behind the

name. But, unless users encrypt e-mail before sending it, every router that forwards the message will be able to read it. Even if the message is encrypted, each router in its path knows the network address from which it was sent and its destination. The user's ISP knows the name and address of the individual who is paying for the service. If the user sends the message from a workplace, then the employer has the right to read it; even a free, public access system is not entirely safe because others may be looking over the user's shoulder. If the user browses the Web, then the Web server reached will very likely be able to learn a lot about the user's computer system, including the make, operating system, and browser. Accordingly, a user querying a database for information on a sensitive disease or condition might wish to take precautions.

There are powerful incentives for Web servers to monitor their visitors, because the data extracted have commercial value—they allow businesses to know which parts of their Web site are interesting to which visitors, thus supporting targeted advertising. Consumers may benefit from such advertising because they learn of new products in a manner that coincides with their tastes, but the implied lack of privacy can be a deterrent to the use of the Internet in certain health care applications. Patients express considerable concern about health information. To protect their privacy, some patients withhold information from their care providers, pay their own health expenses (rather than submit claims to an insurance company), visit multiple care providers to prevent the development of a complete health record, or sometimes even avoid seeking care (Health Privacy Working Group, 1999). The Internet may ease some such concerns because it enables consumers to find health information without visiting their care providers, and it may eventually allow them to seek consultation from, or be examined by, multiple providers in different parts of the country. But without additional privacy protections, a host of new companies could collect information about personal health interests from consumers who browse the Web, exchange e-mail with providers, or purchase health products online. Profiles of patients' online activities can divulge considerable information about personal health concerns. Patients have little control over how that information might be used—or to whom it may be sold.

Concerns about anonymity extend beyond consumer uses of the Internet. Care providers and pharmaceutical researchers, too, express concerns about the privacy of their Internet use. Some care providers wonder if their use of the Internet to research diagnostic information might be construed as a lack of knowledge in certain areas. If such information were tracked by—or made available to—employers or consumer groups, then it could hurt providers' practices. Pharmaceutical companies are concerned that the use of online databases by their researchers

may divulge secrets about the company's proprietary research. These concerns can be addressed in a number of ways, both technical and policy-oriented, but they need to be put to rest if the Internet is to be used more pervasively for health applications.

Some mechanisms are available that users can exploit to reduce their exposure to prying eyes on the Internet. Most attempt to protect the anonymity of users, so that the sender of a message or a visitor to a Web site cannot be identified by the recipient or the Web server. Encryption is the basic engine that underlies all of these mechanisms. Until now, most research on anonymous communication has been carried out informally and without specific attention to health care applications. Most of the existing mechanisms were designed and built in the context of the Internet, and the future development of Internet infrastructure may be intertwined with their use. The benefits and dangers of supporting anonymous communication mechanisms have been the subjects of recurrent (and appropriate) discussions.

Health care offers one of the most compelling cases for the benefits. Such mechanisms could make people feel safe in seeking out information about their own health problems, thereby leading to earlier diagnosis and better treatment. They could also be used to solicit reports about the spread of, for example, sexually transmitted diseases and other health problems that individuals may prefer to report anonymously. In addition, anonymous communication mechanisms can help users limit the capabilities of others to build databases of their behavior or can reduce the extent to which they are the targets of undesired commercial solicitations. For all of these reasons, it would be appropriate for the funders of health care research to support investigations into anonymous communication technologies for future Internet architectures.

Anonymous E-mail

Encryption of e-mail can prevent intermediaries in the network from reading the messages but cannot prevent them from knowing that the sender and receiver are communicating; likewise, it cannot necessarily hide the identity of the sender from the receiver.²⁸ The first mechanisms developed to support anonymous e-mail messages were called re-mailers. These mechanisms would permit a user to register a pseudonym with the server. Mail coming from the user would then be re-mailed by the server, which would strip out identifying material in headers and make it appear that the mail originated at the re-mailer. The re-mailer could forward replies sent to "pseudonym@remailer" to the registered user. For additional protection, the user could encrypt traffic sent to the re-mailer, so that a wiretapper with connections to the re-mailer's inputs and outputs

could not easily defeat the mechanism. The wiretapper still could probably identify which user was sending mail to which destination by looking at the timing and lengths of messages sent and forwarded by the re-mailer.

A single re-mailer remains a point of trust and vulnerability, because it knows the mapping between identities and pseudonyms. This vulnerability has been exploited in legal attacks: for instance, the operator of a widely used Finnish re-mailer discontinued operations when he found that, under Finnish law, he could be forced to reveal the identities of his subscribers. Some U.S. companies have revealed pseudonym-identity mappings when subpoenaed in civil cases.

To provide stronger protection, David Chaum (1981) proposed a network of re-mailers, called mixes. In this scheme, each e-mail message traverses a sequence of mixes and then is reencrypted for transit across each link. In addition, each mix collects a set of messages over a period of time and reorders the set before forwarding them, so that even an observer who could trace the sequence of arrivals and departures from all mixes would be unable to trace a message through the network. Ad hoc networks of re-mailers that incorporate some of these approaches are now operating on the Internet. A commercial service, Anonymizer.com, provides an anonymous re-mailing facility that permits a sender, free of charge, to specify a chain of re-mailers.

Protected Web Browsing

Because forwarding of e-mail does not require a real-time connection from sender to receiver, it is reasonably easy to protect sender anonymity, at least partially. Web browsing, because it depends on a reasonably prompt interaction between client and server, is more difficult to protect. The timing of message arrival and departure may make it obvious to an observer that two parties are communicating, even if the message contents and addresses are obscured. The problem of how to hide the identity of a user browsing the Web from a server that it accesses can be broken into two parts: first, how to prevent an eavesdropper from being able to trace the path of the traffic and, second, how to prevent the server from sending traffic over the path that causes the client (against the user's wishes) to reveal information that could identify the user.²⁹ Most of the techniques developed for protecting Web browsing have been, or could be, adapted to support anonymous e-mail and other functions (e.g., file transfer, news, VPNs) as well.

A straightforward approach to providing anonymous Web browsing uses a trusted intermediary, analogous to a simple re-mailer. The user forwards the universal resource locator (URL) of interest to the intermediary, which strips any identifying information from the requests, perhaps

even providing an alias, and forwards the request to the intended server. To the server, the request appears to have come from the intermediary.³⁰ The intermediary also forwards any data returned to the appropriate requester. Anonymizer.com, the Rewebber (formerly Janus), and Proxymate (formerly Lucent Personal Web Assistant) all provide services of this sort. The communication between the client and the trusted intermediary can be protected from simple eavesdropping by using SSL over this link. The Rewebber also supports anonymous publishing by providing encrypted URLs. A user wishing to retrieve data from an anonymous server obtains an encrypted URL for that server (this encrypted version may be freely distributed). The Rewebber then decrypts the URL, forwards the request to the hidden server, collects the reply, and returns it to the user.

Technology to hide the communication path, based on an enhancement of Chaum's Mix networks, has been developed and prototyped by the Naval Research Laboratory in its Onion Routing Project (Reed et al., 1998). This scheme creates a bidirectional, real-time connection from client to server by initiating a sequence of application-layer connections within a set of nodes acting as mixes. The path through the network is defined by an "onion" (a layered, multiply-encrypted data structure) that is created by the user initiating the connection and transmitted to the network. Only the onion's creator knows the complete path; each node in the path can determine only its predecessor and successor, so an attack on the node operators will be difficult to execute. This strategy also limits the damage that a compromised onion routing node can do; as long as either the first or last node in the path is trustworthy, then it is difficult for an attacker to reconstruct the path. All the packets in the network have a fixed length and are mixed and re-encrypted on each hop. In the event that the submitted traffic rate is too low to assure adequate protection, padding (dummy packets) is introduced. These defenses can be expected to make it extremely difficult to use traffic analysis to deduce who is talking to whom, even if an eavesdropper can see all links.

Onion routing needs a separate screening mechanism to anonymize the data flowing between client and server, so that the server is blocked from sending messages to the client that will cause client software to reveal its identity. Although the Onion Routing Project has implemented an anonymizing proxy to perform this type of blocking, a server can play any of an increasing number of tricks to determine the client's identity. Other projects, such as Proxymate, have specifically concentrated on hiding the identify of the client from the server and have devised more robust techniques for doing so than those developed under the Onion Routing project (Bleichenbacher et al., 1998). These techniques can be combined with onion routing to provide strong protection against both traffic analysis and servers that might try to identify their clients. A system for

protecting personal identity on the Web that appears to be closely based on onion routing is being offered commercially by Zero Knowledge Systems (1998) of Montreal.

A different approach has been prototyped by AT&T researchers in their Crowds system (Reiter and Rubin, 1998). Instead of creating a separate network of mixes to forward traffic, each participant in a Crowd runs a piece of software (called a jondo) that forwards traffic either to other nodes in the Crowd or to its final destination. In effect, when a member of a Crowd receives a packet, it flips a weighted coin. If the coin comes up heads, then the participant decrypts the packet and sends it directly to its Internet destination address. Otherwise, it forwards the packet to another randomly chosen jondo. The Web server receiving the packet can only identify the jondo that last forwarded the packet; it cannot deduce the packet's true origin. Return traffic follows the same randomly generated path in the reverse direction.

Anonymous Payment

In the physical world, individuals who do not want stores to track their purchases can pay cash. The standard approach for buying items on the Internet, however, is to use a credit card, which is guaranteed to reveal the purchaser's identity. Several schemes based on cryptographic mechanisms can enable anonymous payment over the Internet. Chaum (1989) pioneered research in this field, but the rise of e-commerce has triggered much additional work in recent years. The basic idea is to create the electronic analog of a coin—a special number. The merchant must be able to determine that the coin is valid (not counterfeit) without requiring the identity of the individual presenting it. Because computers can copy numbers so easily, a basic problem is to prevent a coin from being spent twice. Although Chaum and others have invented schemes that solve this problem and yet provide anonymity (at least for users who do not try to commit fraud), it has proven difficult to transfer these solutions into the world of commerce. Law enforcement authorities express concern over such technologies because of the potential for their use (or misuse) in money laundering and tax evasion.

Anonymous Data Released from Sensitive Databases

For many years, the U.S. Bureau of the Census has been charged with releasing statistically valid data drawn from census forms without permitting individual identities to be inferred. The problem of constructing a statistical database that can protect individual identities has long been known to researchers (Denning et al., 1979; Schlörer, 1981; Denning and

Schlörer, 1983). To limit the possibility of identification, statisticians have developed several techniques, such as restructuring tables so that no cells contain very small numbers of individuals and perturbing individual data records so that statistical properties are preserved but individual records no longer reflect specific individuals (Cox, 1988). Medical records have often been disclosed to researchers under the constraint that the results of the research not violate patient confidentiality, and, in general, researchers have lived up to this requirement.

Recently, researchers have shown how easily even data stripped of obvious identifying information (name, address, social security number, telephone number) may still disclose individual identity, and they have proposed both technical approaches to reduce the chance of confidentiality compromises and guidelines for future release policies (Sweeney, 1998). The benefits of having full access to relevant data for research purposes and the difficulty of rendering data anonymous without distorting it are likely to require a continuing trust between researcher and subject. An earlier report by the Computer Science and Telecommunications Board (1997a) discussed systemic flows of information in the health care industry and proposed specific criteria for universal patient identifiers. In particular, the report called for technical mechanisms that would help control linkages among health care databases held by different organizations, reveal when unauthorized linkages were made, and support appropriate linking. Because Internet connectivity greatly facilitates such linkages, it is appropriate to renew the call for research into such mechanisms in the present report.

CONCLUSION

As the discussion in this chapter demonstrates, ongoing efforts to enhance the capabilities of the Internet will produce many benefits for the health community. They will provide mechanisms for offering QOS guarantees, better securing health information, expanding broadband access options for consumers, and protecting consumer privacy. At the same time, the technologies expected to be deployed across the Internet in the near future will not fully meet the needs of critical health care applications. In particular, QOS offerings may not meet the need for dynamically variable service between communicating entities. Security technologies may not provide for the widespread issuance of certificates to health care consumers. And the Internet will not necessarily provide the degree of reliability needed for mission-critical health applications. Although much can be done with the technologies currently planned, additional effort will be needed to make the Internet even more useful to the health community.

One way to ensure that health-related needs are reflected in networking research and development is to increase the interaction between the health and technical communities. As researchers attest, most networking research is conducted with some potential applications in mind. Those applications are shaped by interactions with system users who can envision new applications. To date, interaction between health informatics professionals and networking researchers has been limited. By contrast, the interests of industries such as automobile manufacturing and banking are well represented within the networking community, in part because of their participation in the IETF and other networking groups. The health community may need to better engage these groups to ensure that health interests are considered.

BIBLIOGRAPHY

- Birman, K.P. 1999. *The Next Generation Internet: Unsafe at Any Speed?* Department of Computer Science Technical Report, Draft of October 21. Cornell University, Ithaca, N.Y.
- Blake, S., et al. 1998. *An Architecture for Differentiated Services*. IETF Request for Comment (RFC) 2475, December.
- Bleichenbacher, D., E. Gabber, P. Gibbons, Y. Matias, and A. Mayer. 1998. "On Secure and Pseudonymous Client-Relationships with Multiple Servers," pp. 99-108 in *Proceedings of the Third USENIX Electronic Commerce Workshop*, Boston, September.
- Braden, R., S. Shenker, and D. Clark. 1994. *Integrated Services in the Internet Architecture: An Overview*. IETF Request for Comment (RFC) 1633, June.
- Braden, R., L. Zhang, S. Berson, S. Herzog, and S. Jamin. 1997. *Resource ReSerVation Protocol (RSVP): Version 1 Functional Specification*, IETF Request for Comment (RFC) 2205, September.
- Chaum, D. 1981. "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Communications of the ACM* 24(2):84-88.
- Chaum, D. 1989. "Privacy Protected Payments: Unconditional Payer and/or Payee Untraceability," pp. 69-93 in *Proceedings of SMARTCARD 2000*, D. Chaum and I. Schaumuller-Bichl, eds. North-Holland, Amsterdam.
- Clark, D. 1999. "The Internet of Tomorrow," *Science* 285(July 16):353.
- Clark, D., and J. Wroclawski. 1997. *An Approach to Service Allocation in the Internet*. IETF Draft Report, July. Massachusetts Institute of Technology, Cambridge, Mass. Available online at <<http://diffserv.lcs.mit.edu/Drafts/draft-clark-diff-svc-alloc-00.txt>>.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1994. *Realizing the Information Future: The Internet and Beyond*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1996. *The Unpredictable Certainty: Information Infrastructure Through 2000*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1997a. *For the Record: Protecting Electronic Health Information*. National Academy Press, Washington, D.C.

- Computer Science and Telecommunications Board (CSTB), National Research Council. 1997b. *Modeling and Simulation: Linking Entertainment and Defense*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1999. *Trust in Cyberspace*. National Academy Press, Washington, D.C.
- Cox, L.H. 1988. "Modeling and Controlling User Inference," pp. 167-171 in *Database Security: Status and Prospects*, C. Landwehr, ed. North-Holland, Amsterdam.
- Denning, D.E., and J. Schlörer. 1983. "Inference Controls for Statistical Databases," *IEEE Computer* 16(7):69-82.
- Denning, D.E., P.J. Denning, and M.D. Schwartz. 1979. "The Tracker: A Threat to Statistical Database Security," *ACM Transactions on Database Systems* 4(1):76-96.
- Dierks, T., and C. Allen. 1999. *The TLS Protocol Version 1.0*. IETF Request for Comment (RFC) 2246, January.
- Ellison, Carl, and Bruce Schneier. 2000. "Ten Risks of PKI: What You're Not Being Told About Public Key Infrastructure," *Computer Security Journal* 16(1):1-7.
- Goldberg, I., and D. Wagner. 1998. "TAZ Servers and the Rewebber Network: Enabling Anonymous Publishing on the World Wide Web," *First Monday* 3(4). Available online at <<http://www.rewebber.com>>.
- Halabi, B. 1997. *Internet Routing Architectures*. Cisco Press, Indianapolis, Ind.
- Hawley, G.T. 1999. "Broadband by Phone," *Scientific American* 281(4):102-103.
- Health Privacy Working Group. 1999. *Best Principles for Health Privacy*. Institute for Health Care Research and Policy, Georgetown University, Washington, D.C.
- Huitema, C. 1995. *Routing in the Internet*. Prentice-Hall, Englewood Cliffs, N.J.
- Institute of Medicine (IOM). 1997. *The Computer-Based Patient Record: An Essential Technology for Health Care*, rev. ed. Dick, R.S., E.B. Steen, and D.E. Detmer, eds. National Academy Press, Washington, D.C.
- Jacobson, V. 1988. "Congestion Avoidance and Control," *Computer Communication Review* 18(4):314-329.
- Kent, S., and R. Atkinson. 1998a. *Security Architecture for the Internet Protocol*. IETF Request for Comment (RFC) 2401, November.
- Kent, S., and R. Atkinson. 1998b. *IP Authentication Header*. IETF Request for Comment (RFC) 2402, November.
- Kent, S., and R. Atkinson. 1998c. *IP Encapsulating Security Payload (ESP)*. IETF Request for Comment (RFC) 2406, November.
- Marbach, W.D. 1983. "Beware: Hackers at Play," *Newsweek* 102(September 5):42-46.
- Paxson, V. 1997. "End-to-End Routing Behavior in the Internet," *IEEE/ACM Transactions on Networking* 5(October):601-615.
- Perlman, R. 1992. *Interconnections: Bridges and Routers*. Addison-Wesley, Reading, Mass.
- Peterson, L., and B. Davie. 2000. *Computer Networks: A Systems Approach*. Morgan Kaufmann, San Francisco.
- Reed, M.G., P.F. Syverson, and D.M. Goldschlag. 1998. "Anonymous Connections and Onion Routing," *IEEE Journal of Selected Areas in Communication* 16(4):482-494.
- Reiter, M.K., and A.D. Rubin. 1998. "Crowds: Anonymity of Web Transactions," *ACM Transactions on Information Systems Security* 1(1):66-92.
- Reuters. 1999. "AMA, Intel to Boost Online Health Security," October 13.
- Schlörer, J. 1981. "Security of Statistical Databases: Multidimensional Transformation," *ACM Transactions on Database Systems* 6(1):95-112.
- Shenker, S. 1995. "Fundamental Design Issues for the Future Internet," *IEEE Journal of Selected Areas in Communication* 13(7):1176-1188. Available online at <<http://anaweb/www.lcs.mit.edu/anaweb/pdf-papers/shenker.pdf>>.
- Skoro, J. 1999. "LMDS: Broadband Wireless Access," *Scientific American* 281(4):108-109.

- Sweeney, L. 1998. "Datafly: A System for Providing Anonymity in Medical Data," in *Database Security XI: Status and Prospects*, T.Y. Lin and S. Qian, eds. Chapman & Hall, New York.
- Zero Knowledge Systems, Inc. 1998. *The Freedom Network Architecture, Version 1.0*. Available from ZKS, 3981 St. Laurent Blvd., Montreal, Quebec, Canada. December.
- Zimmerman, Philip. 1994. *The Official PGP Users Guide*, Technical report. MIT Press, Cambridge, Mass.

NOTES

1. Evidence of such latencies can be seen in data collected by the National Laboratory for Applied Network Research, which are available at <www.nlanr.net>.
2. ISPs typically have POPs in major urban areas; a large provider might have 30 or more POPs in the United States.
3. The 30 Tbps figure was calculated by multiplying the number of strands per fiber (30) by the number of wavelengths that can be transmitted over each fiber (100) and the capacity of each fiber at each wavelength (10 Gbps). A terabit is 10^{12} (one thousand billion) bits per second.
4. SONET is a standard developed by telephone companies for transmitting digitized voice and data on optical fibers.
5. See <<http://209.249.142.16/nnpm/owa/NRpublicreports.usagemonthly>>.
6. The 10 Gbps figure results from multiplying 10 Mbps by 1,000 applications (10 Mbps \times 1,000 = 10 Gbps).
7. For example, even if available bandwidth were 10 times greater than the average required, the load on certain links over short time periods could be large enough to impose large delays over those links.
8. IP is a connectionless, packet-switching protocol that serves as the internetwork layer for the TCP/IP protocol suite. It provides packet routing, fragmentation of messages, and reassembly.
9. Because of its reliance on RSVP, the int-serv model sometimes is referred to as the RSVP model.
10. With RSVP, the load on the router can be expected to increase at least linearly as the number of end points increases. Growth may even be quadratic—related to the square of the number of end points (Birman, 1999).
11. An example of a scaling issue for today's ISPs is the size of routing tables, which currently hold about 60,000 routes (address prefixes) each. Entries in the routing table consume memory, and the processing power needed to update tables increases with their size. It is important that such tables grow much more slowly than do the numbers of users or individual applications, making it infeasible to store RSVP information if it grows in direct proportion to the number of application flows.
12. The charter of the Integrated Services Over Specific Link Layers working group of the IETF is available online at <<http://www.ietf.org/html.charters/issll-charter.htm>>.
13. The Department of Defense has a long-standing interest in using multicast technology to support distributed simulations. See CSTB (1997b).
14. One of the more notorious cases occurred when the "414" group broke into a machine at the National Cancer Institute in 1982, although no damage from the intrusion was detected. See Marbach (1983).
15. Unix's Network File System (NFS) protocol, commonly used to access file systems across an Internet connection, has weaknesses that enable a "mount point" to be passed to unauthorized systems. Surreptitious programs called Trojan horses can be exploited to perform actions that are neither desired by nor known to the user.

16. Most U.S. health care providers continue to maintain patient records on paper, but current trends in clinical care, consumer health, public health, and health finance all indicate a shift to electronic records. Without such a shift, the health community's ability to take full advantage of improved networking capabilities would be severely limited. With such a shift, the need for convenient, effective, and flexible means of ensuring security will be paramount.

17. Tools such as Back Orifice can enable a hacker using the Internet to remotely control computers using Windows 95, Windows 98, or Windows NT. Using Back Orifice, hackers can open and close programs, reboot computers, and so on. The Back Orifice server has to be willingly accepted and run by its host before it can be used, but it is usually distributed claiming to be something else. Other such clandestine packages also exist, most notably Loki.

18. For a discussion of key distribution centers, see CSTB (1999), pp. 127-128.

19. For a discussion of some of the limitations of PKI systems, see Ellison and Schneier (2000).

20. It should be noted that when using SSL, data are decrypted the moment they reach their destination and are likely to be stored on a server in unencrypted form, making them vulnerable to subsequent compromise. A number of approaches can be taken to protect this information, including reencryption, which presents its own challenges, not the least of which is ensuring that the key to an encrypted database is not lost or compromised.

21. Whereas one organization may issue a certificate to anyone who requests one and fills out an application, another may require stronger proof of identity, such as a birth certificate and passport. These differences affect the degree of trust that communicating parties may place in the certificates when they are presented for online transactions.

22. Additional information on the Intel initiative is available online at <<http://www.intel.com/intel/e-health/>>.

23. Participating organizations in the HealthKey initiative are the Massachusetts Health Data Consortium, the Minnesota Health Data Institute, the North Carolina Healthcare Information and Communications Alliance, the Utah Health Information Network, and the Community Health Information Technology Alliance, based in the Pacific Northwest. Additional information on the program is available online at <www.healthkey.org>.

24. Users can do this, for example, by registering their public keys with a public facility, such as the PGP key server at the Massachusetts Institute of Technology.

25. Computer scientists generally consider system (or network) availability to be an element of security, along with confidentiality and integrity. As such, availability is discussed within the security section of this chapter. Other chapters of this report discuss availability as a separate consideration to highlight the different requirements that health applications have for confidentiality, integrity, and availability.

26. Cable modems and DSL services are typically not attractive to businesses, either because the number of connected hosts (IP addresses) is limited or the guaranteed minimum delivered bandwidth is low. In the San Francisco Bay Area, asymmetric DSL delivers anywhere from 384 kbps to 1.5 Mbps, depending on many factors. In other areas, DSL with 256 kbps/64 kbps down/up link speed costs approximately \$50 per month, but the costs skyrocket quickly to roughly \$700 per month for 1.5 Mbps/768 kbps down/up link speeds.

27. Quality of service mechanisms, such as integrated services, might help ameliorate contention for cable bandwidth, but only if the technology is widely deployed.

28. There is at least one way to hide the identity of the sender: All e-mail applications can be spoofed.

29. Intel Corporation introduced an identifying number into its Pentium microprocessors to help servers identify client machines in the hopes of facilitating electronic commerce. Public concern over the privacy implications of this capability caused the company to take the additional step of providing a means to prevent the number from being revealed.

30. This is essentially what a filtering firewall does: hides the identities (IP addresses) of those behind it.

4

Organizational Challenges to the Adoption of the Internet

Health care organizations have had to adapt to many changes in the world around them, from advances in diagnostic and therapeutic procedures to the emergence of administrative innovations such as managed care and the invention of new information technologies. The Internet represents a particularly profound change that will enable or force significant changes in organizational form and processes—a transformation as profound as any that have gone before. The Internet’s capability to empower consumers, support dynamic information exchanges among organizations, and “flatten” organizational hierarchies promises to result in new operational strategies, business models, service delivery modes, and management mechanisms. The changes will have such far-reaching implications that health care organizations need to start preparing now to adopt the advanced Internet applications that are expected to appear in the near future. Organizations need to evaluate the potential and implications of new Internet technologies, adapt them to local needs and conditions, minimize the risks associated with new product and service deployment, and plan to demonstrate the value of their efforts.

This chapter examines challenges to the adoption of Internet-based technologies by health care organizations. It attempts to identify fundamental impediments to greater use of the Internet that may be expected to persist for some time. The first section of the chapter provides a context for the analysis by reviewing the experiences of other industries that have achieved some success in changing business practices by adopting Internet technologies. The second section discusses in general terms why

more health care organizations should be interested in adopting the Internet: because it can advance their strategic interests. The third section deals with organizational barriers that hinder Internet use. The discussion there addresses internal and external factors, such as policy and technical barriers, that influence and constrain the form and extent of Internet use, as well as the range of uncertainties that inhibit decision making regarding the Internet. The last section addresses the importance of organizational leadership.

The chapter focuses primarily on Internet adoption by care provider organizations. To be sure, Internet technologies will need to be adopted by a number of different players in the health care arena, including consumers, physicians, and administrators. But because they are likely to bear many of the implementation costs and will have to address issues of acceptance by consumers and care providers, health care organizations are a suitable focus for this analysis.¹ The discussion recognizes that health care organizations come in many forms—community-based health facilities, managed care organizations (MCOs) and health maintenance organizations (HMOs), integrated delivery networks (IDNs), and nonacute facilities—each of which may adopt the Internet for different applications, whether management and administration, communications among health care professionals, consumer education, or patient management. The specific factors that facilitate or impede Internet applications will differ from one organization to the other, but the discussion that follows is broadly applicable because it identifies common challenges faced by a variety of health care organizations as they attempt to implement a range of applications serving different types of end users.

LESSONS FROM OTHER INDUSTRIES

Internet technologies offer a range of potentially useful applications to organizations in many different industries. Simple Internet applications such as electronic mail (e-mail) can facilitate communication within distributed multinational corporations. Related networked applications can simplify flows of information among elements of a single organization and among multiple organizations. Real-time teleconferencing technologies can support meetings involving individuals located in different cities. Direct capture of sales information can enable retailers to streamline the delivery of inventory and forecast purchasing patterns. New automation systems can allow for distributed management of supply chains, support of human resource functions, and exchange of contact and other sales information. Although the deployment of these systems is still in an early stage, Internet technologies appear to have enhanced organizational performance by lowering costs, increasing efficiency, dif-

ferentiating products and services, or creating broader markets. Leading users of these technologies have found that the value of the Internet lies not simply in automating existing business processes but in creating new means of interaction between suppliers and consumers of products and services, often with significant implications for industry structure.

Industries differ in many respects, and their degree of success achieved in applying the Internet varies as well,² but the experiences of leading companies in different industries in which Internet use is common suggest a number of general trends. The Internet clearly is transforming the retail marketplace. Online retailers such as Amazon.com and Barnes and Noble have changed the nature of the book industry by creating direct relationships between book buyers and book suppliers that have significantly reduced inventory costs and eliminated many middle layers in the distribution chain—a process called disintermediation.³ Manufacturers of personal computers, such as Dell and Compaq, increasingly use the Internet to market their wares directly to consumers, enabling the consumers to customize their orders and enabling the firms to control inventory at the lowest possible levels. Online auction sites, such as eBay, have pioneered new ways to link buyers and sellers in a virtual marketplace, with some companies expanding on the auction concept to exploit spot markets for last-minute airline tickets, car rentals, hotel rooms, and other services.

Internally, the effective use of information technology (IT), including Internet technologies, can have a profound impact on organizational structure and function. As information is distributed efficiently to those who need it when they need it, lines of control and influence become clearer, and individual units often self-organize in new and more effective ways. The impact may be multifaceted, not only flattening organizational structures but also changing the skill mix of employees. Early evidence suggests that online sales of automobiles reduce the set of skills needed by salespeople (McGarvey, 1999). In contrast, some stock brokerage firms report that online trading requires brokers to have a broader set of skills, although the total number of brokers needed may decline because much of the effort of executing a stock transaction can be passed on to the consumer.

Online interactions boost consumer expectations. Many traditional storefront industries—from retail to manufacturing to news—are now open around the clock, competing in a highly visible and competitive environment. Consumers conduct many transactions at night or on holidays, when many traditional merchants shut their doors, and buyers often compare the prices of many Internet vendors before making purchases. In fact, many Internet companies encourage consumers to discuss topics or items of particular interest. Internet book merchants, for example,

allow readers to contribute reviews and rate the quality of an offering. Internet-based vendors of financial information sometimes support client-generated discussion groups on specific equities or investing techniques. These techniques are intended to assist consumers in making educated decisions and, simultaneously, attract them to particular sites.

Internet technologies also allow merchants to develop a deeper understanding of consumers. By automatically recording consumer choices and preferences, merchants can offer both goods and advertising that have a high likelihood of reaching a desired consumer audience. If applied successfully, these technologies enable merchants to develop a sense of one-on-one personalized service for thousands or even millions of customers—a process sometimes called mass customization. Vendors can also allow consumers to preview, or experiment with, products prior to purchase. The film industry now routinely provides previews of upcoming movies on company sites on the Web. The music industry also distributes promotional material online in the hopes of generating traditional sales. Online mortgage vendors allow consumers to simulate the cash-flow implications of various mortgage packages.

Most importantly, the customer relationships established by successful Internet companies are not static. Rather, the companies evolve as their customers' needs and sophistication evolve. The active empowerment of consumers forces companies to provide highly targeted services and a degree of variety commensurate with the buyers' needs. During this period of rapid Internet evolution, companies are literally reinventing their online demeanor on a weekly basis in response to changing perceptions developed through continued experimentation with the Internet. In essence, merchants and consumers are engaged in a consensual exploration of the means by which this technology can more effectively satisfy perceived mutual needs.

ADVANCING THE STRATEGIC INTERESTS OF HEALTH CARE

Just as it is transforming other industries, the Internet could enable profound changes in the nature and structure of the health care industry and, ultimately, the delivery of health care services. The health care industry is—and will continue to be—diverse, with individual organizations facing different environmental pressures, pursuing different missions, and cultivating different cultures, but the Internet appears capable of supporting at least a handful of common strategic interests. It could, for example, help organizations to do the following:

- Improve the efficiency and effectiveness of processes that customers use to judge organizational performance (e.g., scheduling an

appointment) or processes that form the core of the organization's business (e.g., medical management);

- Develop partnerships with related organizations in an effort to leverage respective strengths (e.g., MCOs partnering with pharmaceutical companies to develop disease management programs or regional alliances of providers partnering to form a continuum of care);
- Reach consumers directly to solidify brand names and eliminate intermediaries (disintermediation);
- Improve, differentiate, and deliver new services to key customers; and
- Improve organizational decision making.

The Internet already is empowering consumers to become more involved in and take greater control of their own health and care. Patients are coming into doctor's offices armed with information downloaded from the Internet and suggestions for diagnoses and treatments. They may soon be able to access information on the quality of care delivered by different health care providers or facilities in their geographic region. The relationship between health care organizations and consumers could change even further when an organization uses networks effectively to expand its customer base beyond a specific geographic region. The assumption that just because a patient lives near a particular hospital he or she will opt to be treated at that hospital is challenged by the ease with which networks afford patients access to clinical specialists worldwide.

The Internet also offers institutions the capacity to separate their business operations from the operations of the physical plant. The effective use of computer networks could change the fundamental nature of a health care organization: from a unified entity providing a specific, fixed set of resources and services to a broker entity that acquires services and resources and offers them on an as-needed, on-demand basis. For example, physician practice groups could use the Internet to acquire necessary expertise on demand and deliver health care from a distance or to gain access to distributed decision-support systems and high-end applications offered by application service providers, contracting for specific medical services from remote practice groups rather than investing in an expensive, but seldom used, on-site resource. The use of network technology to reach customers in their homes, schools, and workplaces would extend the opportunity for the delivery of services and products at locations away from a health care organization's physical site (e.g., clinic, hospital, or campus).

The Internet also allows the integration of clinical data from affiliated organizations (such as two hospitals that contract with the same MCO), enabling the assembly of a medical record that is more complete than

before and ensuring greater continuity and documentation of care. This is a useful feature because the increased specialization of clinical treatment has contributed to a differentiation of health care services. Specialization has led patients to acquire health services from different organizations (e.g., hospitals, physician's offices, dentists, laboratories), with the result that the clinical records of most individuals are scattered across several institutions. Other opportunities exist for hospitals to use the Internet to improve internal operations. They could communicate, in a just-in-time fashion, with distributors of supplies—perhaps diminishing the need for an inventory management department—and with each other to share patient records. They could reduce their reliance on local specialists by centralizing specialty services and offering them remotely via teleconsultations. Care providers and payers could link together to consolidate utilization review activities.

The Internet also could affect the nature of health care research. Network-enabled “collaboratories” could allow individual research operations within a particular health organization (such as a research hospital) to interact with research groups at other organizations, creating a complex set of interacting research groups that share information and results across multiple organizations and even national borders. Private research institutes, either free-standing or organized as research and development arms of pharmaceutical companies, could enhance their scientific collaborations with academic medical centers. The Internet also would enable schools and universities to educate not only the students who arrive on their campuses but also authorized learners anywhere in the world who have access to computers and Internet connections.

The contours of such change are beginning to emerge. Health care organizations are beginning to experiment with Internet-based systems to serve a variety of functions. Many MCOs, for example, are developing Web sites to provide consumers with health-related information, the ability to select physicians or schedule appointments, and the tools to evaluate their immediate health care needs. Other organizations—both MCOs and new Internet start-ups—offer personally tailored health information over the Web, matching profiles of a patient's health status with existing network-based health information resources. They also offer specific programs for monitoring personal health status and can support e-mail exchanges between patients and designated providers. Still other firms focus on the administrative side of health care, using secure communications channels to deliver a broad array of services required in a managed care environment, with an emphasis on linking different types of health care organizations. Offerings include products for enrolling in and managing personal health care plans; information services and portals designed for consumers, providers, and physicians; systems for viewing

laboratory and other clinical data; and methods for referring patients to specialists, ordering medications, and performing other clinical tasks.

The returns to health care organizations on investments in Internet-based applications are not yet clear, but the early evidence is encouraging. Kaiser-Permanente of Northern California, for example, reports that a pilot program to test a consumer-oriented Web site for 100,000 members reduced the number of visits to physicians' offices by 11 percent, reduced the number of calls to nurses by 46 percent, and allowed 14 percent of the patients to treat their illnesses at home. The result was not only significant cost savings for the organization but also improved consumer perceptions of Kaiser-Permanente and better understanding of health concerns. Partners Healthcare System, an integrated delivery organization based in Boston, reports that Internet-based systems reduced the time needed to return a radiology report to a health center from 72 hours to 4 hours. The organization expects to realize a 20 percent reduction in the cost of specialist dermatology by using telemedicine.

More study is required to fully evaluate the benefits of Internet-based applications in health care, but the evidence cited above offers hope for improvements in customer satisfaction and reductions in cost. Given the potential of the Internet and the economic and other pressures facing the health care industry, it would be reasonable to expect significant investment in Internet technologies and applications by more health care organizations.

IMPEDIMENTS TO ADOPTING INTERNET APPLICATIONS

Despite the promise of many Internet-based applications, health care organizations can be expected to encounter many obstacles as they attempt to apply these technologies to realize their strategic visions. They will face barriers to, and constraints on, organizational change, as well as uncertainty about the efficacy and effects of Internet-based applications. A resistance to change might come from denial of the need to change, the inability to manage change, uncertainties about the types of changes needed and how best to make them, mistaken assessments of optimal changes, and failures in executing changes. These issues are not unique to the adoption of the Internet and could arise in many other areas of organizational change, including those driven by other types of information technology.⁴ What makes Internet-driven change different is its magnitude and the high degree of uncertainty. Organizations that have difficulty making the necessary investments in, and managing, information technology in general will have even more difficulty adapting to the Internet.

Barriers to Change

Barriers to change can assume many forms, from characteristics of the marketplace to organizational capabilities. The discussion below classifies barriers into two broad categories: (1) external factors that define the environment in which an organization operates and (2) internal factors that define the ability to implement change. External factors are often difficult for an individual organization to address directly because they demand collective action. Internal barriers are easier to overcome but in many cases still present significant difficulties, even if the need for change is recognized as urgent.

External Barriers

External factors define the environment in which health care organizations operate and shape their ability to capitalize on the Internet. The barriers here assume many forms, including market forces, policies and standards, finances, and technology.

Market Forces. Changes in the health care marketplace—such as an aging population, escalating health care costs, and changes in consumer preferences—can have positive and negative effects on health care organizations and the viability of different Internet-based applications. Often the effects cannot be anticipated. For example, an aging population can be expected to create a growing demand for health services, a trend that could benefit local hospitals but could raise costs for MCOs, which would need to provide more care on a per capita basis. An aging population also could boost demand for Internet-mediated care, but at the same time as this would allow the patients to avoid travel, it would raise challenging interface considerations. To the extent that changes in the marketplace signify to organizations a need to change, they can be considered stimuli for transformation; to the extent that they work against organizations, they can be considered barriers.

Other market forces also affect the ability of health care organizations to adopt Internet technologies. Many health care organizations have seen shrinking operating margins over the last several years, along with a marked increase in regulatory compliance requirements (such as those related to the Health Insurance Portability and Accountability Act of 1996—see Chapter 5) and in competing programmatic needs (Glaser and Hsu, 1999). The lack of coherent business models at work in many of the new Internet-based health companies (drugstore.com, Healtheon/WebMD, drkoop.com, etc.), along with their lack of demonstrated strong

financial returns, does little to persuade established health care organizations to aggressively shift to an Internet-based strategy.

Policies and Standards. State and federal laws and regulations, professional standards, and technical standards, which remain largely outside the exclusive control of the health care industry, are among the policy barriers that can impede health care organizations. These barriers can interfere with business opportunities enabled by the Internet, reducing the incentive to undertake initiatives that would benefit the nation's health care system and citizens' health. For example, policies on professional licensure can impede attempts to deliver health services across state lines using the Internet; regulations on data security and patient privacy can create disincentives for the transmission of information across public communication networks such as the Internet (these issues are addressed in greater detail in Chapter 5). Policies of third-party payers and health insurance companies on payment for medical services can also affect a care provider organization's incentives to develop new Internet applications. If services rendered cannot be reimbursed, the incentives for deploying capabilities are reduced. For HMOs the incentives may be different, because HMOs that reduce the costs of care can improve their profitability regardless of payment policies. As with market forces, changes in policy can serve either to motivate change or to constrain it.

Finances. Financial barriers arise from the complex and sometimes perverse mechanisms for funding health care. Care providers are often rewarded for treating disease rather than preventing it, a situation that can be expected to reduce interest in developing wellness-oriented Internet material. MCOs may exhibit a similar lack of interest in wellness activities if their subscribers are young and transient. Payment is often made for each visit to the care provider's office rather than for overall treatment of a particular health problem, with the result that a provider has little incentive to invest in Internet-based disease management capabilities. Moreover, providers generally cannot receive reimbursement for care provided at a distance using electronic communications technologies (i.e., telemedicine), even if such practice could reduce the cost of care (see Chapter 5). Insured consumers do not pay for the vast majority of their care and are constrained (by insurance companies) in their ability to find and engage their preferred source and form of care—choices the Internet could facilitate. In effect, Internet applications, while benefiting the patient and society, require investments by health care organizations without providing any balancing revenue. Furthermore, the economic anomaly whereby the consumer of care is not the purchaser of that care limits the

consumer's ability to fully pursue the use of the Internet to support personal health.

Technology. Technology itself can be a barrier to change in organizations. Technological change is especially rapid in information technology, a supreme challenge for organizations that try to keep up with the pace of innovation while controlling costs. Significant technological changes can create major dislocations, rendering investments in existing technologies obsolete. Organizations cannot depreciate prior investments fast enough to keep up with the rate of change or shift their technical and human infrastructures rapidly enough without undermining organizational performance. This process of technological change is particularly challenging in a health care context, where there is a heightened need to demonstrate the efficacy of such change—a process that can be expensive, time-consuming, and difficult to do well—and the wide range of potential users involved—patients, physicians, nurses, administrators—all of whom may have different work flows and skills. Managers attempting to learn about the performance, maturity, and potential of new technologies are often confronted with exaggerated claims. The distillation of truth from an overload of information is an error-prone and inefficient process.

External barriers present a significant challenge to organizations. Some barriers, such as demographic trends and the complex and politicized decision-making processes in health care, are often out of the control of organizations, although they can take steps to anticipate and mitigate the effects. Policy barriers, in general, cannot be addressed effectively by individual organizations but require a coordinated and concerted effort by multiple players. Policy barriers can be addressed in a number of ways, as explained in Chapters 5 and 6 of this report.

Internal Barriers

Internal barriers can prevent organizations from recognizing the need to change and properly implementing the required changes. A lack of organizational self-awareness, responsiveness, and competency, and a reluctance to change—characteristics of organizational inertia—all impede attempts to implement necessary change. Inertia is associated with a large size, long history, and complicated internal hierarchies—characteristics of many health care organizations. Resilient organizations that overcome inertia have the capacity to revise their structure and function to effectively manage external forces; organizations characterized by inertia are less likely to do so.

An organization's ability to change is influenced by many factors, including its competence, sophistication, and history of action with other technologies. When basic operating principles must be reinterpreted, organizations require time to promulgate, implement, and assimilate the new standards, policies, and guidelines based on the new principles. The appropriate use of new technology often requires a degree of process and role redefinition not usually encountered in health care settings. It may pose threats to individual roles or positions, challenge the rationale of current business or clinical practices, demand rapid political mobilization, encounter user resistance, and require additional funding. Processes used in implementing new technologies, particularly if reliant on consensus, can slow the effort to adopt them. New forms of communication can require new interpretations of basic principles such as the core nature of an organization's services. Achieving consensus takes time, perhaps more time than is feasible given the rapid pace of technological advances.

Hesitancy to change may result from legitimate organizational concerns as well as organizational inertia. Organizations may be reluctant to adopt a new technology because it is a poor fit with the existing strategy or because they are in a market segment that does not reward innovation. In addition, change may be hindered by conflicts between the ability of management to maintain coherent integration throughout an enterprise and the need for rapid deployment of a new technology in a more localized setting, such as in a particular department or laboratory.

Internal barriers must be changed by individual organizations. At the very least, organizations need to be aware that these barriers exist and that their persistence will hinder the efficient and effective use of the Internet (and a wide range of other technologies and innovations). An organization should also be aware of the special problems associated with its market sector. Barriers actually encountered, and the degree to which a particular barrier is viewed as problematic, will vary among organizations. For example, clinics, particularly small group practices outside academic health science centers, lack the administrative and financial resources of larger organizations. They may face barriers to full Internet implementation as a result of reimbursement policies that prohibit cost recovery. Additionally, they may lack the organizational resilience to deliberate and manage the uncertainty inherent in investing in technologies such as electronic medical records or distributed information systems.

A hospital, on the other hand, may fail to capitalize on emerging technologies because its staff is not competent to understand how such technologies could benefit the hospital. Researchers whose collaborators once worked across the hall now must communicate in an up-to-the-minute fashion with investigators around the world and are hampered from doing so by the lack of common standards for reporting. Knowl-

edge providers (e.g., libraries) face intellectual property and ownership considerations that are rapidly changing and for which existing rules no longer provide adequate guidance. Professionals (e.g., physicians and nurses) are both protected by and limited by state-based licensure requirements that protect their scope of practice but limit their ability to practice across state lines.

Uncertainties Surrounding Internet Strategies

Even if organizations can identify the need to embrace an Internet-based strategy and can overcome some of the internal and external barriers to doing so, their progress can be inhibited by considerable uncertainty about how best to proceed. Uncertainty could cause an organization to hesitate to pursue an Internet-based strategy and could result in sub-optimal, deleterious, or unnecessarily turbulent or inefficient change. Health care organizations face uncertainty in four areas: organizational and industry structures for Internet-based care, internal policies and procedures to guide Internet use, technological capabilities of Internet-based systems, and human resource issues. For each type of uncertainty, examples of related issues and questions are presented below. The examples are not intended to be comprehensive but to illustrate the range and depth of the uncertainty surrounding organizational Internet-based health care strategies.

Organizational and Industry Structure

Organizations are experimenting with the use of the Internet to extend or alter current processes. Patients communicate with providers through e-mail. Telemedicine is offered as a means of increasing referrals from remote clinics. Some organizations offer medical content to providers for a subscription fee. Others attempt to connect a range of providers and payers in a region to improve the efficiency of care. Some pharmaceutical benefits companies may choose to ally themselves with online pharmacies, whereas others may open their own online pharmacies. As these examples demonstrate, the Internet has the potential to fundamentally alter the economics, organizational form, and interorganizational processes that define health care today. These new arrangements and relationships are complex and not well understood. Poor conception of new organizational forms and poor execution of a form place the transforming organization at risk. Faced with such risk, many organizations hesitate to engage in new forms. Others try different relationships between participants and different types of services until one or more successful models emerge.

The effectiveness of many new health care processes and initiatives is

not clear. Will providing patient access to medical information improve the quality of care and patient satisfaction? Does telemedicine increase referrals? Will linking providers with retail pharmacies or pharmacy benefits managers improve medication compliance? Will health care organizations become, to a large degree, virtual? Their limited understanding of the business value of these initiatives causes organizations to be hesitant in carrying out Internet-enabled change. The examples below explore particular uncertainties in greater detail.

Example 1: Changes in Industry Structure. Internet investments and start-ups target industries with multiple distribution channels, large numbers of middlemen, and massive product selections (Fitzgibbons and Lee, 1999). Health care is such an industry; accordingly, a number of Internet-based start-ups, such as Healtheon, CareSoft, and drkoop.com, have entered the marketplace. The intent of these investments and start-ups is to significantly alter and improve the mechanisms by which health care is distributed, in some cases simplifying administrative transfers of information and in others creating new channels for the provision of health care and health information. It is not clear whether extensive use of the Internet will dramatically change the health care industry's structure, altering the relationships among consumers, care providers, managed care organizations, and insurers. At one extreme, the introduction of the Internet could so empower consumers that they force major changes in MCO operations that now tend to limit their choice of care providers and to direct provider-patient interactions. At the other extreme, the Internet could simply reinforce existing structural arrangements, perhaps making them more efficient.

Example 2: Separation of Content from Transaction. The Internet allows separating content acquisition (e.g., conducting research on vacation spots) from the execution of a transaction (e.g., booking a vacation). Consumers often use information gathered in one location to inform decisions made and transactions executed elsewhere. This practice is already occurring in health care as consumers find health-related information on the Internet and later present it to care providers to help guide decision making on diagnoses and needed tests, procedures, and medications. Similarly, discussions with care providers can lead consumers to order medical supplies or nonprescription drugs through an Internet-based service. The disaggregation of transaction and content, and the commoditization of both, may have profound impacts on the value proposition on which care providers have long based their existence. The ramifications are not well understood by the provider community (or most other communities facing similar changes in value propositions). Further-

more, health care providers do not understand how to take advantage of such a change or ensure that patients are not harmed by it.

Example 3: Alteration of Roles. The Internet is altering the traditional roles of and relationships among patient, primary care provider, local specialist, and tertiary hospital specialist. As consumers become better informed about health issues, they may pressure MCOs to relax restrictions on the pursuit of care directly from a specialist. Patients may also seek greater flexibility in selecting specialists who, although they may not be local, offer superior service and expertise. Such specialists would include radiologists, dermatologists, nutritionists, pharmacologists, and others with knowledge of particular medical domains. If this occurs, then existing referral patterns may be disrupted (or strengthened), and the definition of high-quality service and care may change. Patients may come to expect, as they have with many Internet-based services, 24-hour availability of certain service offerings that are customized to their needs. This trend, in turn, may reinforce the already growing roles of physicians' assistants and nurse-practitioners in the provision of care.

Example 4: Management of the Chronically Ill. The Internet may markedly improve the ability to manage and treat the chronically ill. Improvements can result from a range of Internet-based services, such as the translation of medical text into lay language, patient-moderated chat sessions, "ask-a-doctor" chat rooms, remote consultations, and remote monitoring of patient health data. The extent to which such improvements will occur is not clear; also unknown is the relative utility of various types of Internet-based services and content. In both cases, the answers will vary by disease and across cohorts of patients.

Example 5: Spot Market Purchase of Health Care. The Internet supports consumers seeking to make spot market purchases of books, airline tickets, and automobiles by providing information on sources, cost, availability, and quality for the item or service to be purchased. This trend has yet to become important in health care, but one can imagine consumers engaging in spot market purchases of medications, vitamins, and durable medical equipment. Continued experience with this model could encourage consumers to pressure MCOs to allow spot market purchases of radiology procedures, simple surgical procedures, medical education, and routine checkups. Consumers in fee-for-service insurance plans might be able to enter this marketplace more quickly. A thriving spot market could reduce the cost of medical procedures by allowing rapid comparisons of services based on cost and quality. It is not clear at this time which—if any—elements of health care delivery would be amenable to spot market pur-

chases, whether and to what degree consumers would be comfortable making them, or the degree to which such purchases would pose risks or improve care.

Internal Policies and Procedures

Effective management of clinical, administrative, and fiscal processes and activities requires policies and procedures. Internet use often creates a need for additional policies and procedures to guide and manage the technology and the new patterns of communication and processes it enables. For example, e-mail interchanges between patient and provider—which are becoming more common in the health care sector—require policies on the appropriate types of clinical conversation (e.g., e-mail should not be used to convey highly sensitive information unless one can ensure reader authentication).⁵ A provider, serving as a moderator of a disease-specific chat room, may inadvertently create a provider-patient relationship. The ease with which a Web page can be created by an employee may make it appear as if an organization is sponsoring or supporting the content when in fact it actually disagrees with it—if indeed it is even aware of it. As organizations grapple with the need for new management and medical policies and procedures, they will find little precedent for how to deal with these issues. Faced with a policy and procedure vacuum and an absence of models, organizations may hesitate to sanction Internet use because they are not sure they can manage it well enough.

Health care organizations would benefit if they could refer to reference models to assist them in developing policies and procedures. Reference models have been developed to guide other activities, such as general confidentiality and security practices unrelated to the Internet. These models include sample forms for consent and release of information, confidentiality policies, and chief security officer job descriptions. Other reference models have been developed to guide the use of e-mail between patients and care providers (Kane and Sands, 1998). The examples below explore some of the uncertainties surrounding Internet policies and procedures.

Example 1: Monitoring/Conducting Health-Related Chat Sessions, Bulletin Boards, and Forums. A number of health care organizations and consumer health Web sites have established chat rooms and forums for their members. Chat rooms and forums would appear to have clinical, service, and marketing value for a wide variety of users: the chronically ill, family members of patients with debilitating diseases, consumers interested in health care topics, and providers. Hosting such sessions, however, can raise a number of issues that the sponsoring organization and the mod-

erator must address. For example, can the care providers who moderate them diagnose or suggest treatment in such forums? What types of liability does the sponsor assume for the quality of information, the behavior of the forum participants, or the outcome of a participant's efforts to seek care? How do providers and patients prevent the formation of inappropriate provider-patient relationships? What are the characteristics of a high-quality forum service?

Example 2: Assessing Trade-offs Between Security/Confidentiality and Access. Some providers and payers will extend their services to enable patients to access their personal clinical data over the Internet. A number of online health companies are already providing such a service. These services may offer information related to test results and suggestions for pursuing related health care (e.g., a test showing high cholesterol levels might elicit a recommendation to seek dietary counseling). Such access and the information it provides may improve care as well as patient perceptions of service quality. However, such access also poses security and confidentiality risks. How can organizations balance the technical and management mechanisms required to provide this access against the need to maintain acceptable security? Furthermore, as consumers are offered the opportunity to create and maintain their own health and medical records online, new issues arise regarding data availability and integrity. Data for these records will come from claims, electronic medical record systems, and provider questionnaires, and from the patients themselves. A patient's physician may have only a limited opportunity to review this record to assure its integrity and quality.

Example 3: Evaluating Content. A large number of Web sites (20,000 by some estimates) offer health information, algorithms for evaluating health risks and status, and tools for managing chronic diseases. This is a potential boon to consumers because the information is widely available; however, providers and consumers often find it difficult to evaluate the quality of information available on sites that provide such content. Health care organizations need to determine, among other matters, how they can help individuals distinguish reliable information dispensed by reputable sources from less reliable information, and how providers (especially physicians) can deal with the demands for untested treatments sometimes coming from individuals who perform their own Internet research.

Technological Capabilities

All industries are besieged by claims that Internet-based content, services, and applications can solve a wide range of problems that have

plagued prior efforts to implement other information technologies. In health care, these problems include the difficulty of integrating a diverse, heterogeneous set of legacy systems; a lack of data standards; the complexity of medical practice; and complex, politicized health care organizational structures. It is not clear whether the traditional hurdles to implementation of complex clinical information systems, such as the electronic medical record, are significantly lowered by the use of the Internet and related technologies. For example, although the capability to provide access to a set of services, content, and applications through a wide array of workstations and other devices might reduce the support costs for an application, it is not clear that the aggregate cost of the application infrastructure is reduced.

It is highly likely that the use of the Internet and related technologies will ease some of the challenges that have plagued the implementation of information systems in health care over the years. It is also highly likely that some challenges will persist, unaffected by the presence of the Internet. Regardless of these uncertainties, new applications will be based on the Internet and related technologies, and legacy applications will be fitted with interfaces to the Internet. Inappropriate expectations of the Internet's ability to solve certain problems can cause the Internet's true contributions to be disregarded or organizational resources to be squandered. As they assess Internet-enabled applications—particularly those retaining the same labels as non-Internet applications (e.g., electronic medical records)—health care organizations will have limited resources with which to understand the value added by the Internet. The following examples illustrate some of the dilemmas that may arise.

Example 1: Cost of Applications. Many organizations assume that Internet applications will reduce hardware costs by allowing the use of networked computers (or “thin clients”) with less internal capability than traditional desktop computers but with access to shared information and applications stored on centralized Web servers. To assist with this transition, application service providers (ASPs) such as Abaton.com and QCSI are offering online access to a variety of computer programs. Unfortunately, few full cost comparisons have been made between ASP-delivered applications and those delivered with existing architectures. Moreover, there may be hidden costs in the transfer to an Internet-based infrastructure. For example, an Internet-based infrastructure introduces new security vulnerabilities that must be addressed, and the ability of organizations to benefit from the Internet will depend on other IT-related expenditures. Variations across institutions in the implementation of EMRs have a significant effect on an organization's capability to capitalize on the Internet

in support of delivery of care, payment and administration of care, and research. Hospitals and insurance plans need to have networking infrastructure, electronic medical records, and clinical systems in place before they can develop more interesting Web applications. Elements of the infrastructure include operating systems, networks, information, and applications.

Example 2: Design Challenges. Clinical applications of the Internet pose complex design challenges. The applications must support comprehensive interactions in a manner that is intuitive, efficient, dependable, and fast and that clearly indicates the next set of options. Are there design situations and tasks for which the Internet, specifically the Web, is either ill suited on the one hand or well suited on the other? Are additional refinements needed to tailor Internet-based services to different user groups with different capabilities and needs? The enrolled clients of MCOs and HMOs may vary considerably in their socioeconomic status and educational levels, for example, and different capabilities may be needed to get them to use online services. The match between what health care organizations can do to better structure information access, and what clients require and can use, needs to be examined.

Example 3: Ease of Use. Individual human factors constrain Internet use in health care organizations. Applications must be easy to use and well integrated into the daily routines of users. The use of IT by health professionals is growing, but it is not uniform across regions (e.g., rural vs. urban areas), medical specialties, or practice types. Some applications have a long history (e.g., results reporting systems), yet it is not clear what factors facilitate the use of information technology by health care professionals or what types of organizational arrangements or financial incentives facilitate or impede the technology's use. Not enough is known about the barriers faced by care providers and consumers in using the Internet routinely in health-related activities. A substantial gap exists between a relatively small group of providers and consumers who actively use the Internet for any purpose and the much larger group that has not used it or cannot use it. Similarly, not much is known about the training requirements that Internet technologies will impose on provider organizations. Substantial innovation and development will be needed to train and prepare care providers. Training for clinical information systems has always been difficult and labor-intensive. Applications often have hundreds of important features that are difficult to teach to care providers who are impatient and pressed for time. Does the ubiquity of the Internet and Web ease the training burden?

Example 4: Systems Integration. Internet technologies are touted as a new answer to the challenges of integrating large numbers of diverse legacy information systems. However, it is not clear to what degree the technology can overcome difficulties such as the lack of data syntax and semantic standards, technology incompatibilities, and vendor unwillingness to participate actively in projects that integrate their systems with those of competitors.

Human Resource Needs

A significant organizational commitment to Internet-based applications requires information technology workers who are skilled in new technical areas such as HTML development, Internet security, and digital commerce. It may also result in unanticipated needs for expanded computing and information processing skills among the organization's information technology staff and other workers. For example, a patient at home whose workstation experiences problems while accessing an organization's Web site may call that organization's help desk for support

BOX 4.1 **Unanticipated Support Costs**

The emergence of unexpected technical challenges is illustrated by the experiences of Kaiser-Permanente of Northern California with its consumer-oriented Web site. The goals of the Web site are to provide value-added services to members, facilitate the appropriate use of services, and encourage members to take a more active role in their health care. Kaiser members can also use the Web site to obtain health information and service directories, request and receive problem-specific answers to clinical questions, participate in electronic discussion groups with other members, and conduct simple tasks such as scheduling a clinic visit. More than 16,000 members have access to the site.

One of the biggest challenges faced by the Kaiser staff after system launch was the unexpected need to provide technical coaching and services to members. Member calls indicated that the quality of a Web session was highly dependent on the configuration of the member's browser, and Kaiser was placed in the position of providing technical support for the members' home computers. Additional challenges included (1) assuring that clinicians would agree to use the clinical protocols developed to answer questions and (2) developing policies for monitoring—and intervening in—discussion groups. For instance, Kaiser needed to rapidly determine how to respond in a safe and appropriate manner to a message posted by a member who was in a life-threatening situation. In spite of these challenges, Kaiser's efforts have been successful enough to inspire imitation by other managed care organizations, not-for-profit support groups, and for-profit health care sites.

(see Box 4.1). The organization can thus be confronted with potentially enormous and unanticipated support costs. Providers who enable their patients to send e-mail may discover the need to redirect office staff and nursing time to responding to e-mail and printing messages for filing with paper records.

Health care organizations would benefit from information on the staffing levels and skills needed to develop, implement, and support Internet-based applications. Some experience has accumulated during early demonstrations and evaluations of electronic medical records. In addition, real-life experiences or lessons learned, in the form of case studies or conference presentations on new or more complex situations, would assist many organizations in planning for Internet use. The following examples suggest some of the staffing issues that may arise.

Example 1: Information Systems Staff. Internet-based applications are likely to present information systems staff with situations that are new or more complex than those they encountered with other technologies and applications. As more patients begin using its Internet-based offerings, an organization may be asked to provide workstation support to an increasingly broad and less-well-controlled user base. Troubleshooting problems on a network shared by multiple organizations can be difficult. Organizations may need to expand their staffs in traditional information systems categories, such as application developers and analysts, implementation and training, database and network technical issues, and support (e.g., for the help desk). They may also need to retrain existing workers to enhance their understanding of new Internet technologies and the ways in which they can be used in a health care environment.

Example 2: Process Support Staff. Management, medical and allied health personnel, and clerical staff will need to support new clinical, administrative, and financial processes. New processes may include triaging, responding to, and filing patient e-mail and responding to patient requests for appointments; managing Internet content updates and versions; developing Internet-based curricula and serving as teaching assistants to remote students; and managing telemedicine consultation setup and sessions. Additional staff may be needed to carry out these functions, and existing staff may need retraining.

ESTABLISHING ORGANIZATIONAL LEADERSHIP FOR INFORMATION TECHNOLOGY

Health care organizations will not increase their commitment to Internet-based applications without strong leadership. Issues of vision

and leadership are often crucial determinants of successful health care applications of information technology (IT) in general. Limited, narrowly focused applications may be implemented at a grassroots level in an organization and successfully applied, but when applications require complex interactions across—and potentially beyond—the organization, the skills and talents of individual participants must be augmented by strong institutional leadership and a shared vision of what the organization is trying to accomplish.

Until the 1980s, information systems managers for major corporations typically played a technical, service-oriented role. They brought technical computing and communications skills, plus management abilities, to the organization. In recent years, however, the strategic role of information systems has become clearer as corporations plan for the future, identify new business opportunities, and implement new practices for communicating with clients, distributing products, and managing inventories and finances. As a result, the technology managers are increasingly identified as key strategic leaders for the organizations—at least outside the health care industry. Their titles have generally evolved (typically to chief information officer (CIO) or vice president for information systems and technology) to reflect this central role. When their roles were considered technical rather than strategic, they often reported to the chief financial or administrative officer. Today, however, they more typically report to the chief executive officer (CEO) and participate actively in high-level strategic planning, priority setting, and decision making. In fact, Peter Drucker has suggested that information and supporting technology is becoming so central to strategy that the CEO of the future will be the CIO (Drucker, 1999).

Another important change has occurred during the past two decades. Originally the IT leader had little industry-specific expertise (e.g., a drug company CIO typically would lack a medical or pharmacology background). Their responsibilities were largely confined to managing large-scale technical installations and implementations, and the IT activities were viewed primarily as a cost center for the organization. It is now axiomatic in many industries, however, that the CIO has deep industry expertise. Ideally, IT leaders “grow up” in that industry and combine domain training or expertise with education in, or an inclination toward, information systems. Rather than constituting a cost center, the IT functions overseen by the CIO are increasingly viewed as enablers of key business or strategic opportunities. Several studies across a range of industries have demonstrated a relationship between the effectiveness of an organization’s application of IT and the presence of a senior, strategy-oriented information systems executive (CHIME, 1998; Earl and Feeney, 1995; Kilbridge et al., 1998; McKenney et al., 1995; Ross et al., 1996;

Sambamurthy, 1996). Effective use of the Internet is unlikely to be an exception.

In health care, the pre-1980 model lingers, with CIOs who often have limited experience in medical environments and a set of expectations and reporting structures that de-emphasize a strategic role and instead suggest a largely service and infrastructure orientation. This is true at all industry levels, from community-based hospitals to large biomedical research universities. The reporting level and professional background of the information strategy leader in an organization inevitably influence the attention that information management receives and can affect the types of initiatives undertaken. A CIO with a background in clinical medicine might place a higher priority on the establishment of EMRs and decision-support systems than would a CIO with a strictly technical background. This is not to suggest that every information strategy leader in health care needs to be clinically trained as well. What is critical is that the CIO function as a professional peer of the organization's senior leadership. This individual will require a deep understanding of the business and mission of the organization and a place at the table when strategic decisions are made.

Health organizations need to recognize the importance of IT management. Certainly there are important educational issues: health care leaders must be familiarized with the fundamental role that IT needs to play in their organizations, and an expanded cadre of future CIOs must be produced who have a combination of technical and management skills as well as knowledge of the medical environment and its complex cultural constraints. The effective mobilization of the Internet for health care, and its use for highly leveraged applications and demonstration projects, will require an investment and commitment from medical organizations and visionary leadership from CEOs and other key managers who understand the strategic role of the technology and the return on investment that can be expected.

SUMMARY

Although Internet use is embryonic across U.S. industry, its transforming potential can be seen in the existing Internet initiatives and experiences of health care and other organizations. Health care organizations and Internet-based suppliers of products and services are experimenting with different relationships between participants and different types of services, seeking to find business models that work. It is clear, however, that the Internet can advance the strategic interests of the health care industry. Nevertheless, health care organizations will confront many barriers to Internet use. As improved technologies become deployed in the Internet,

organizations will be able to consider new applications, and new barriers are likely to arise.

Research in several areas could give health care organizations the confidence they need to move forward in using the Internet. First, research on Internet-induced changes in health care economics, organizational form, and interorganizational processes would provide guidance for organizations and patients, helping to ensure that the changes are effective and that they do not materially damage the health care system or harm the health of patients and consumers. Second, model policies and procedures for the effective management of Internet-related clinical, administrative, and fiscal processes and activities would help organizations address these issues before they become problems. Third, an assessment of the Internet's capability to resolve (or not resolve) the health care system's persistent difficulties with implementing and managing older information technologies would provide guidance on the value added. Fourth, case studies of the staffing levels and skills needed to develop, implement, and support Internet-based applications could assist planning for Internet use. Poorly conceived or poorly executed change could have serious negative consequences for the nation's health care delivery system, health care organizations, and health care research and education, as well as for its citizens. All of this research is needed for developing information and guidance to promote a positive transformation of the health care industry.

REFERENCES

- College of Healthcare Information Management Executives (CHIME). 1998. *The Healthcare CIO: A Decade of Growth*. College of Healthcare Information Management Executives, Ann Arbor, Mich.
- Drucker, P. 1999. *Management Challenges for the 21st Century*. Harper Business Press, New York.
- Earl, M., and D. Feeney. 1995. "Is Your CIO Adding Value?" *McKinsey Quarterly* 1995(2):144-161.
- Fitzgibbons, Stephen M., and Richard Lee. 1999. *The health.net Industry: The Convergence of Healthcare and the Internet*, Industry Report. Hambrecht and Quist Institutional Research, San Francisco, Calif., January 8.
- Glaser, John P., and Leslie Hsu. 1999. *The Strategic Application of Information Technology in Healthcare Organizations: A Guide to Implementing Integrated Systems*. McGraw-Hill, New York.
- Kane, Beverley, and Daniel Z. Sands. 1998. "Guidelines for the Clinical Use of Electronic Mail with Patients," Report for the AMIA Internet Working Group, Task Force on Guidelines for the Use of Clinic-Patient Electronic Mail, *Journal of the American Medical Informatics Association* 5(1). Available online at <<http://www.amia.org/pubs/pospaper/positio2.htm>>.

- Kilbridge, P., et al. 1998. "Information Systems for IDNs: Best Practices and Key Success Factors," pp. 229-241 in *Proceedings of the 1998 Annual HIMSS Conference, Volume 2*. Healthcare Information and Management Systems Society, Chicago.
- McGarvey, R. 1999. "Online Care Chase," *Upside* 11(12):154-161.
- McKenney, J., D. Copeland, and R. Mason. 1995. *Waves of Change: Business Evolution Through Information Technology*. Harvard Business School Press, Boston, Mass.
- Organization for Economic Cooperation and Development (OECD). 1999. *The Economic and Social Impact of Electronic Commerce*. OECD, Paris.
- Ross, J., C. Beath, and D. Goodhue. 1996. "Develop Long-Term Competitiveness Through IT Assets," *Sloan Management Review* 38(1):31-42.
- Sambamurthy, V., and R. Zmud. 1996. *Information Technology and Innovation: Strategies for Success*. Financial Executives Research Foundation, Morristown, N.J.

NOTES

1. Other types of health organizations that are not directly involved in the provision of health care—such as educational institutions, insurers, and public health agencies—will face similar challenges, but they are not addressed specifically in the chapter.
2. The Organization for Economic Cooperation and Development (1999), for example, estimates that between 2000 and 2005, 60 percent of retail stock trading but only 20 percent of book sales and 7 percent of music sales will migrate to the Internet.
3. Similarly, online travel service companies, such as Travelocity and Expedia, have adapted airline reservation systems in ways that allow consumers to make purchases directly from airlines, eliminating the travel agency as an intermediary.
4. Security, scalability, human resources, and cost are common areas of uncertainty with new information systems.
5. The American Medical Informatics Association has developed a set of guidelines to help providers determine how to use e-mail for patient care (Kane and Sands, 1998).

5

Issues for Public Policy

A variety of cultural, social, and political forces influence the ways in which health care providers, health management organizations, publishers of biomedical knowledge, researchers, consumers, and others use the Internet. By shaping the prevailing policies of the times, these forces influence not only the types of Internet-based applications that are likely to find widespread, routine use in the health sector but also the underlying technical capabilities of the Internet. Policies affecting Internet use, and the debates surrounding them, reflect fundamental beliefs about the way things work—or should work—in the United States. In this sense, these policies are points of potential conflict, where ideas about individual rights, the public good, equal access, free enterprise, and the role of government collide with material, economic, and technical realities and possibilities. The uncertainties surrounding the resolution of policy issues, combined with the organizational dilemmas outlined in Chapter 4, generate considerable doubt in regard to the technological trajectories that are likely to be pursued in the health sector and the capability of the Internet to support health applications. Clearly, the resolution of technical issues will not, by itself, enable greater use of the Internet in the health sector.¹

This chapter examines six policy issues that influence the use of the Internet in support of health objectives: the protection of personal health information, access to information infrastructure, the protection of intellectual property contained in educational and reference materials, regulatory issues associated with the electronic delivery of medical services, federal support for health informatics research, and human resources.

These topics have been addressed in several national reports in the last five years (Box 5.1), a sign of both their significance and the difficulties of resolving them. Many of these issues have implications outside the health sector; nevertheless, their importance in health applications argues for strong leadership by the health community in their resolution. The chapter describes the issues, the uncertainties they introduce to the Internet's deployment in health applications, and ongoing efforts to address them. Consistent with the charge to the committee and the expertise of the committee's members, the chapter does not attempt to offer recommendations for resolving these issues. In many cases, additional study will be required to delineate more fully the trade-offs among possible solutions and gather sufficient information to render reliable guidance. These issues will need to be resolved if the benefits of the Internet are to accrue to the health community.

PROTECTION OF PERSONAL HEALTH INFORMATION

Technology is, to a large extent, both the cause of and the solution to concerns about the protection of personal health information. The capability to connect health information systems to the Internet exposes personal health information to hostile attacks that can alter, delete, or divulge it (see Chapter 3). At the same time, technologies such as passwords, encryption, and firewalls offer reasonably effective means of protecting information systems and the data contained within them. Nevertheless, the effective assurance of data and service protection ultimately depends on the implementation of policies and practices within the organization (see CSTB, 1997). To whom should organizations be allowed to disclose personal health information with and without patient consent? Under what conditions may such disclosures be made? What steps must organizations take to protect personal health information from loss, unauthorized editing, or mischief? What types of security technologies and administrative policies will be considered sufficient protection?

Concerns over patient privacy are not new in the health sector. Patients and consumer advocate groups have long expressed concern about the collection, use, and sharing of personal health information data and the practices used to maintain its confidentiality. These groups view the Internet as further eroding patient privacy by making health information more easily available to a larger number of users (e.g., insurers, direct marketers, and pharmaceutical benefits managers) and more susceptible to security breaches. Health care organizations have been sensitive to the vulnerabilities inherent in the Internet and have long used private networks for data exchange, both because products and services are available to support such networks and because the organizations are confident of

BOX 5.1 Other Reports That Identify Policy Issues Related to Health Informatics

Realizing the Information Future: The Internet and Beyond (CSTB, 1994). This report looked at the powerful potential of the emerging national infrastructure to “enrich people’s economic, social, and political lives . . .” (p. 1). It went on to say, “Today we lack a consistent technical, legal, and business framework for the dissemination of intellectual property over networks” (p. 100). Noting the need to balance public and private interests, the report called for a fuller consideration of competing interests, particularly as they affect “decisions relating to societal equity—including access to networks and the information resources available on them” (p. 211).

Telemedicine: A Guide to Assessing Telecommunications in Health Care (IOM, 1996). Noting that “most clinical applications of telemedicine have not been subjected to systematic comparative studies that assess their effects on the quality, accessibility, or cost of health care,” this report presented a framework for evaluating the practicality, value, and affordability of telemedicine. Concluding a lengthy chapter on the policy context of telemedicine, the committee noted: “The task for this committee was to develop an evaluation framework for clinical telemedicine—not to develop policy recommendations. The committee recognized, however, that policies related to licensure, malpractice, and other matters need to be considered . . . because they may affect the availability, acceptability, effectiveness, and cost of telemedicine services” (p. 115).

For the Record: Protecting Electronic Health Information (CSTB, 1997). This report looked at technical and nontechnical mechanisms and issues relating to the privacy and security of health care applications of the national information infrastructure. It led off with a chapter on the public policy context. Among other things, it said that “better protection of electronic health information will require efforts at the national level. The lack of uniform national standards for the privacy and security of health information creates particular problems for health care organizations that serve constituents in multiple states and creates additional confusion for patients regarding their rights.” It then suggested that “conflicting views of data ownership and a lack of patient understanding of health data flows and of their rights to privacy and confidentiality also need to be addressed . . .” (pp. 49-50).

The Computer-Based Patient Record: An Essential Technology for Health Care (IOM, 1997). Updating and expanding the original report, published in 1991, this revised edition provides a scorecard on the implementation of the original recommendations. The original committee concluded that “computerization can help to improve patient records and that improved patient records and information management of health care data are essential elements of the infrastructure of the nation’s health care system” (p. 46). The authors of the revised edition’s progress report noted: “Security, privacy, and confidentiality concerns have become major barriers to widespread implementation of [computerized patient record] systems and [the] sharing [of] data. There is, as yet, no agreement on what must be done to establish the balance between appropriate use of health care data and the individual patient’s rights to privacy” (p. 14).

Health Data in the Information Age: Use, Disclosure, and Privacy (IOM, 1994). This report examined the potential of health data organizations to improve health

and the performance of the health care system, as well as issues relating to the quality of health information contained in data repositories and protection of the confidentiality of personal health information. As the report noted, “[e]xisting ethical, legal, and other approaches to protecting confidentiality and privacy of personal health data offer some confidentiality safeguards, but major gaps and limitations remain” (p. 15). The report recommended preemptive legislation to fill these gaps.

their ability to manage and protect the networks. Although some organizations have taken steps to use corporate intranets or virtual private networks (see Chapter 3 for a description of these technologies) for sharing personal health information among their affiliates, they have been reluctant to transmit such information over the public Internet because of concerns about privacy and security. Until key decision makers in health organizations are confident of their ability to protect the personal health information for which they are responsible, opportunities to harness the Internet’s full potential in health care will be seriously constrained.²

Policy makers have already entered into the debate over privacy. Numerous bills have been introduced in both houses of the U.S. Congress relating to the use of medical records and personal health information (Box 5.2), but the enactment of legislation has been impeded by factors such as differences of opinion on the restrictions that need to be placed on the ability of different stakeholders (e.g., pharmaceutical companies, direct marketers, and legal authorities) to access health information. The most notable advance was contained in the Health Insurance Portability and Accountability Act of 1996 (HIPAA, P.L. 104-191), which gave the Congress until August 1999 to pass legislation regarding the privacy of individually identifiable health information. The HIPAA also directed the secretary of Health and Human Services to promulgate regulations on the topic by February 2000 if no legislation was passed by the August deadline. President Clinton announced proposed regulations in October 1999. These regulations (1) allow health information to be used and shared easily for the provision of care and payments for care, (2) establish procedures for disclosing health information without the patient’s consent for purposes such as research, public health, and oversight, (3) require written authorization to use and disclose personal health information for other purposes, (4) create a set of fair practices to inform consumers about ways in which their information is used and disclosed, ensure that consumers have access to information about them, and allow patients to propose corrections or amendments to such information, and (5) require organiza-

BOX 5.2 Legislation Related to Health Privacy Issues

By May 1999, more than 70 bills on privacy and confidentiality had been introduced in the 106th Congress. Some of these were related to health information, others to financial information or personal information collected by companies for marketing purposes. Although obviously not exhaustive, the following three Senate and House bills attest to the political importance of confidentiality and privacy issues as they relate to health:

- S. 881 (Bennett): *Medical Information Protection Act of 1999*. The bill requires specified health entities in possession of protected health information to allow the subjects of the information to inspect, copy, and amend it. It directs the secretary of Health and Human Services to develop model notices of confidentiality. It also mandates (1) administrative, technical, and physical safeguards for protected health information, (2) a record of any protected health information disclosures, and (3) identification of disclosed information as protected health information. In addition, the bill prescribes guidelines for the disclosure of protected health information with respect to (1) authorizations for treatment, payment, and health care operations, (2) the individual's next-of-kin and directory information, (3) emergency circumstances, (4) certain oversight agencies, (5) public health authorities, (6) health researchers, (7) civil, judicial, and administrative procedures, (8) certain law enforcement procedures, (9) payment for health care by credit card or electronic means, (10) certain duly authorized representatives acting on behalf of a subject individual (including a deceased subject individual, or a minor), and (11) certain business sales, transfers, or mergers. Finally, it imposes criminal penalties for knowingly and intentionally disclosing protected health information and establishes civil monetary penalties for failure to comply with the act.

- S. 578: *Health Care Personal Information Nondisclosure Act of 1999*. The bill requires specified parties to permit an individual who is the subject of protected health information to inspect, copy, and request amendment of the information (or, if amendment is refused, to permit appending the individual's disagreement statement). It requires specified parties to maintain safeguards to ensure the confidentiality, security, accuracy, and integrity of protected health information. It also requires specified parties to maintain a disclosure record and prohibits them from disclosing protected health information. In addition, the bill requires (1) a single authorization form for each individual for disclosure in connection with treatment, payment, and health care operations and (2) a separate authorization for other purposes, including disclosure with intent to sell, transfer, or use protected health information for commercial advantage. It requires the development of model authorizations for circumstances other than treatment, payment, and health care operations. It also allows and regulates disclosure (1) to next of kin (or another person identified by the information subject), (2) of directory information, (3) regarding a deceased individual, (4) in emergencies, (5) for oversight, public health, or health research, (6) in civil, judicial, and administrative procedures, and (7) for law enforcement. Finally, the bill imposes criminal and civil penalties.

- S. 573/H.R. 1057: *Medical Information Privacy and Security Act*. The bill requires specified parties to permit an individual who is the subject of protected

health information to inspect and copy the information. It sets forth rules concerning (1) supplements to protected information and (2) provision of notice of privacy practices. It requires specified parties to establish safeguards to ensure the confidentiality, security, accuracy, and integrity of protected health information. It also mandates the development of model safeguard guidelines; requires specified parties to establish a record of disclosures not related to payment or treatment; prohibits specified parties from disclosing protected health information, except as authorized under this title; and allows disclosure if authorized by the information subject. In addition, the bill sets forth rules governing authorizations for the disclosure of protected information for purposes other than treatment or payment. It allows any person to disclose protected health information (1) to allay or remedy a threat of imminent physical or mental harm to an information subject and (2) if there is an identifiable threat of serious injury or death to an identifiable individual or group and other requirements are met. It authorizes disclosure to (1) a public health authority, (2) certain protection and advocacy agencies if an individual is vulnerable to abuse or neglect by an agency providing health or social services, (3) a health oversight agency, under specified circumstances, and (4) on court order, a law enforcement authority. The bill also regulates disclosure (1) to next of kin and (2) in directories of individuals admitted to a facility. It directs the secretary of Health and Human Services to report to the Congress whether written informed consent should be required and, if so, under what circumstances, before personally identifiable data can be used for medical research. It establishes the Office of Health Information Privacy, including in its duties the receiving and investigating of violation complaints and providing for the conduct of audits. It imposes criminal and civil sanctions. Finally, it amends the Privacy Act of 1974 (P.L. 93-579) to require an agency that receives protected health information to promulgate rules to exempt a system of records within the agency from all but specified provisions of that act.

SOURCE: Information derived from bill summaries contained on the Library of Congress's THOMAS system, available online at <<http://thomas.loc.gov/home/thomas.html>>.

tions to implement technical and administrative mechanisms to protect electronic health information. As a general rule, the regulations limit organizations to releasing the "minimum amount [of information] necessary to accomplish the relevant purpose."³

These regulations represent a significant step forward in the protection of personal health information and should reduce some of the uncertainty regarding allowable exchanges of health information and minimum levels of security protection. Other issues remain to be addressed, however, because the regulations were limited in several ways by the nature of the HIPAA legislation. First, the regulations apply only to

electronic health information—they do not apply to paper-based records, which constitute the bulk of all medical records held by provider organizations. Second, the regulations apply only to health care providers (that transmit information electronically), health plans, and health care clearinghouses (e.g., third-party administrators who process health care bills); they do not apply to the many other organizations, such as other insurers, pharmacies, and direct marketers—or consumer health Web sites—that routinely handle health information. Furthermore, the regulations leave unanswered the question of whether consumers need an explicit private right of action to enforce their privacy rights. Accordingly, the President called on the Congress to pass comprehensive health privacy legislation that addresses these issues (White House, 1999).

Further uncertainty surrounds the development of unique patient identifiers. The HIPAA requires the secretary of Health and Human Services to develop standard identifiers for care providers, health plans, and patients to streamline the electronic administration of health care benefits and payment of claims. Work on both the provider and plan identifiers has moved forward without great difficulty, but efforts to develop a patient identifier have encountered resistance because of concerns that a unique identifier might facilitate the linking of personal information from different sources, thereby eroding privacy. The Department of Health and Human Services (DHHS) has been reluctant to specify and implement a unique patient identifier before a consensus is reached on the larger issue of protecting personal health information, but it is evaluating several alternatives.⁴ The outcomes of these initiatives will have long-ranging effects on the health care providers, payers, and delivery systems that must implement the recommendations.

International actions may further affect the ways in which personal health information is transmitted over the Internet. The European Union (EU) Data Protection Directive, which went into effect on October 25, 1998, requires EU member states to block outbound transmissions of data to countries that do not have laws providing a level of privacy protection similar to that in the country where the data originated.⁵ The directive affords the people to whom the data refer a host of rights, including the right to be notified of data collection practices, to access information collected about them, and to correct inaccuracies. The directive is meant to facilitate the flow of information among EU member states, but its provisions threaten to cut off data exchanges with the United States, which has no national law governing data protection (BNA, 1998). The Clinton administration favors a safe harbor proposal that would allow firms to self-certify privacy policies and has been working with the EU to find a mutually agreeable solution to the problem. Privacy rights groups see the EU directives as an opportunity to push harder for national legislation on

privacy. The Trans Atlantic Consumer Dialogue adopted a resolution that called on the European Commission to reject the U.S. safe harbor proposal and recommended instead the establishment of an international convention on privacy protection to address public concerns about transborder data flows. Leading U.S. and European consumer organizations agree that neither the industry self-regulation nor the safe harbor proposal would provide adequate privacy protection for consumers.⁶ How this issue will be resolved is unclear.

ACCESS TO INFORMATION INFRASTRUCTURE

The promise of the Internet in health applications is related to its ability to interconnect the diverse members of the health community—care providers, insurers, consumers, researchers, educators, and others—in a dynamic fashion that allows the sharing of information and resources in response to changing needs and affiliations. Health applications of the Internet pose particularly challenging requirements for access. First, they demand that access be widespread—extending to the point of care, whether it be a major hospital, rural physician's office, a patient's home, or a hotel room. Second, to the extent that applications such as home-based patient monitoring and telemedicine to the home become more viable, access links will need to provide high-bandwidth connections to and from the end user. Third, to the extent that the Internet is used for consumer-oriented health initiatives, near-universal access to the information infrastructure will become important to avoid exacerbating existing inequalities in the access of different population groups to health information and health care. As more health transactions move online, there will be strong incentives to ensure broad-based, universal access to the Internet for patients, providers, and administrators—particularly among disadvantaged groups, who are often those most in need of health care. Furthermore, as different technologies for high-speed (or broadband) connections to the Internet are deployed that cost more than a standard telephone line, concerns will arise about unequal access to the Internet by different population segments.

Technology and business trends will go a long way toward improving Internet access. Declines in the price of Internet service and access devices, such as personal computers, promise to make connectivity more affordable. Over the past five years, the average price of an Internet-capable computer has declined noticeably, driving further penetration of computing—and the Internet—into the home. Statistics from the U.S. Department of Commerce indicate that roughly 42 percent of all U.S. households owned a computer in 1998, and 26 percent had Internet access. These figures compare to penetration rates of 24 percent for computers in

1994 and 19 percent for the Internet in 1997 (NTIA, 1999). Several companies have begun discounting the price of computers for customers who purchase a contract for Internet service. In addition, lower-cost alternatives to the computers have already begun to enter the marketplace. WebTV (now owned by Microsoft) offers a device for as little as \$99 that enables users to send e-mail and access the Web through a standard television set, using a remote control or keyboard. The 3Com Corporation sells several models of the Palm Pilot, ranging in price from \$100 to \$700, which offer various degrees of connectivity to the Internet but less functionality than a computer. Less expensive products like these could help overcome some of the existing economic barriers to Internet access.

Despite these trends, the marketplace may not ensure equitable access to information infrastructure across demographic lines. Broadband connections to the Internet are considerably more costly than traditional telephone line access and are not available in many parts of the country—especially in rural areas that might benefit most from delivery of medical services via the Internet. In addition, many offers of discounted Internet access require consumers to accept additional advertising or to allow greater monitoring of online activities, a condition involving privacy that not all users are willing to accept. Furthermore, Department of Commerce statistics indicate a widening gap between those with computers and Internet access and those without, along a number of socioeconomic lines.⁷ In 1998, households at the lowest income levels were nine times less likely to have Internet access than those with incomes exceeding \$75,000. Fewer than 20 percent of Americans with incomes of \$25,000 or less have access to the Internet either inside or outside the home (e.g., at work), compared with almost 60 percent of those with incomes of \$75,000 or more (NTIA, 1999). Caucasians were two to three times more likely than blacks or Hispanics to have Internet access from any location (e.g., home or office), and rural residents lagged their urban counterparts by several percentage points at all income levels, although the differences were most notable at lower incomes. In sum, the Department of Commerce figures reflect increased Internet access at all income levels over the past few years but show growing gaps based on income, education, and race. A similar study conducted in the emergency department of a large urban pediatric teaching hospital found similar results (Mandl et al., 2000).

Considerable public debate centers on the best way to provide universal, low-cost access to the Internet. Many observers argue that the federal government should stay out of this matter, that the marketplace will take care of consumer needs. However, telecommunications companies point out that they have little financial incentive to lay the infrastructure in locations where populations are sparse (i.e., rural areas) or property costs are prohibitive (i.e., urban areas). Consumer groups have

emphasized a need for increased pressure on telecommunications and cable companies to ensure equitable access to broadband Internet service (see CFA, 1999). Given the prevalence of mergers, takeovers, and partnerships that bring together cable, telephone, and wireless companies with Internet service providers (ISPs) to forge single entities, the issue is murky at best.⁸ The resolution of this debate certainly will affect the size and scope of, and the need for, special programs such as those described in the present report. Numerous strategies have been proposed for expanding consumer access to the Internet for health purposes specifically. These ideas range from encouraging health plans to offer their members (or the general public) access to online health resources and to subsidize the cost of Internet access for their members to providing such support via the Medicare and Medicaid programs, which work with populations that might benefit substantially from improved Internet access (Eng et al., 1998).

The federal government has long played a critical role in promoting universal access to basic telecommunications services. The notion of universal service was first promulgated in the Communications Act of 1934 (P.L. 73-416), the goal of which was “to make available, so far as possible, to all the people of the United States a rapid, efficient Nation-wide and world-wide wire and radio communication service with adequate facilities at reasonable charges.” These words continue to provide the ideological and legal basis for numerous telecommunications practices and programs aimed at ensuring ubiquitous, affordable access. Prior to 1983, attempts to ensure universal service were pursued through AT&T’s internal rate structure, which effectively subsidized the extension of telecommunications services to rural areas. With the breakup of AT&T, the Universal Service Fund (USF) was established to keep telephone service affordable in a competitive telecommunications market. Telecommunications companies operating in the United States (including local and long-distance phone companies, wireless and paging companies, and pay phone providers) contribute to the USF, and companies can draw money directly out of it to defray the cost of delivering discounted service to both low-income communities and rural areas, where the cost of providing service is high. The Telecommunications Act of 1996 (P.L. 105-125) also mandated that the USF provide support for schools, libraries, and rural health care providers.⁹ The Universal Service Administration Company (USAC), under the direction of the Federal Communications Commission, now operates three programs that can support Internet connections: the High Cost and Low Income Program, the Rural Health Care Program, and the Schools and Libraries Program. Each of these programs provides affordable access to modern telecommunications services for schools, libraries,

rural health care facilities, and consumers, regardless of geographic location or socioeconomic status.

USAC's Rural Health Care Program is only one of several programs sponsored by the federal government that subsidize telecommunications services for health care applications. Other programs are operated by the Department of Commerce's National Telecommunications and Information Administration (NTIA) and the U.S. Department of Agriculture (Box 5.3). Although these programs make significant strides toward improving the access of health care organizations and consumers to the information infrastructure, they do not necessarily provide sufficient incentives to ensure connectivity from all home and health care settings or adequately target the full range of participants in health care. For example, the NTIA's Technology Opportunities Program (TOP) provides project funding for a limited time, after which the project is expected to become self-sustaining. A 1998 survey found that 70 percent of the demonstration and access projects initially funded in 1995 were still in full operation or serving an altered or expanded function (Westat, 1999). The remaining 30 percent had either been scaled back or terminated. Lack of maintenance funding was cited as the primary threat to the sustainability of these projects. Projects funded for 21 months or longer were more likely to have expanded to serve additional end users than were short-term projects, evidence of the need for longer-term support.

Ironically, the DHHS does not, itself, have ongoing programs to ensure Internet access for care providers, consumers, or other members of the health community. The National Library of Medicine (NLM) does provide small (\$30,000) grants for Internet connections to public and private nonprofit institutions (or consortia) engaged in health sciences administration, education, research, and/or clinical care and to consortia of health-related institutions. However, these grants do not subsidize the subsequent operational costs.¹⁰ In December 1999, the Clinton administration directed federal departments and agencies to take a number of steps aimed at improving access to the Internet and narrowing the gaps in access across demographic groups. The DHHS was charged, along with the Education, Labor, and Housing and Urban Development Departments, with expanding the nation's network of community technology centers to provide access to technology for low-income Americans and encourage the development of information technology applications that would enable those populations to start and manage their own small businesses (Clinton, 1999). The provision of Internet access at public sites such as libraries, schools, and community centers may suffice for some federal missions (e.g., education), but to serve health purposes adequately, connectivity must extend to all likely points of care, including homes. Many health-related inquiries are too personal to be made in a public

BOX 5.3
**Federal Programs That Support Health-Related
Access to the Internet**

Rural Health Care Program

The Universal Service Administration Company (USAC) supports Internet access for health care providers through the Rural Health Care Division (RHCD, formerly Rural Health Care Corporation), which was formed to ensure that health care providers in rural areas obtain the benefits of current telecommunications technology. The Universal Service Support Program established a fund of up to \$400 million annually to ensure that rural health care providers pay no more than their urban counterparts do for telecommunication services, including Internet access. In particular, the RHCD aims to provide support to rural health care providers for services related to the delivery of telemedicine. Initially, the RHCD expected thousands of physicians to apply for this program. For a variety of reasons, however, the program has not been as successful as counterpart programs that support schools and libraries. In 1999, the RHCD approved and funded the first applications for physician access to the Internet through this program. On May 2, 1999, the Federal Communications Commission established \$12 million as the collection level for the second funding year of the RHCD support program (July 1, 1999, to June 30, 2000). As of October 1999, the RHCD had committed total funding of only \$1.2 million to 223 rural health care providers (USAC, 1999).¹ Although the FCC modified the collection level, it did not revise the \$400 million annual cap for the rural health care support mechanism. Among their drawbacks, the RHCD programs support Internet access by physicians but not mid-level practitioners or allied health professionals. Furthermore, the application process is burdensome, certain telecommunications companies are excluded from participation, and the program does not underwrite bandwidth above T-1 capacity. Without changes, the RHCD program cannot possibly meet its intended goals.²

Technology Opportunities Program³

The Technology Opportunities Program (TOP) is a competitive, merit-based grant program run by the U.S. Department of Commerce's National Telecommunications and Information Agency. The TOP provides matching grants to nonprofit organizations such as schools, libraries, hospitals, public safety entities, and state and local governments. The grants are intended to fund projects that improve the quality of, and the public's access to, health care, education, public safety, and other community-based services. The funds can be used to purchase networking equipment, including computers, videoconferencing systems, network routers, and telephones; buy software for organizing and processing all types of information, including computer graphics and databases; train staff, users, and others in the use of equipment and software; purchase communications services, such as Internet access; evaluate the projects; and disseminate the project's findings.

continued

BOX 5.3 Continued

From its inception in 1994 through October 1999, the TOP had awarded 421 grants totaling \$135.8 million and leveraging \$203 million in local matching funds. Many grants supported projects in rural areas. Sixty-six of these grants, totaling \$25 million in funding, supported health-related projects, ranging from efforts to develop a digital, wireless home health care service network, to the coordination of responses from health and emergency services to high-risk patients in the Pine Ridge Reservation, to the creation of a network linking 11 county health departments in California, through videoconferencing and data communications. Nevertheless, the TOP has been able to fund only a small percentage of the proposed projects. More than 5,300 grant proposals were submitted between 1994 and 1998, requesting \$2.1 billion in funds; only 378 of these proposals were funded, using \$188 million in federal funds.

Distance Learning and Telemedicine Grant and Loan Program

The U.S. Department of Agriculture's Distance Learning and Telemedicine Grant and Loan Program (DLT) was established to encourage, improve, and make affordable the use of telecommunications, computer networks, and related technologies for rural communities to improve their access to educational and/or medical services. The DLT helps rural schools and health care providers invest in telecommunications facilities and equipment to bring in educational and medical resources that might not otherwise be available in rural areas. Demand for the DLT has been high. Approximately 500 rural medical facilities will access improved medical care through linkage with other rural hospitals and urban medical centers for clinical interactive video consultation, distance training of rural health care providers, management and transport of patient information, and access to medical expertise or library resources.

The program is intended to fund projects that deliver critically needed educational and medical services in rural areas through structured, interactive educational training and/or medical professional presence over distance. It facilitates the networking of multiple sites dispersed over a large area rather than single, stand-alone entities. The DLT covers capital costs of acquiring and installing telecommunications hardware at schools, hospitals, and other eligible sites. It also covers other nonrecurring capital costs of establishing a distance learning and telemedicine system; software, training, and technical assistance are among the items that may be purchased. Funded projects are required to become self-sustaining through mechanisms such as user fees, tax assessments, or school budgets. Fifty-two grants totaling \$13 million were awarded in 1999.

¹Detailed information on the awards is available online at <www.rhc.universalservice.org>.

²Presentation by William England, Director of Operations & Systems, Rural Health Care Program, Universal Service Administrative Company, at the Emerging Health Information Infrastructure Conference (HII99), April 27, 1999, Washington, D.C.

³The Telecommunications and Information Infrastructure Assistance program was renamed the Technology Opportunities Program in January 2000 to reflect the opportunities new technologies provide for economic advancement.

arena (e.g., a library or community center), and many health care emergencies would require rapid access to the Internet outside the normal operating hours of public facilities.

It may be possible to encourage organizations to enable this type of access or find mechanisms to pay for Internet access for health applications. For example, incentives could be designed to entice health plans (public or private) to support some of the costs of Internet access for their patients. ISPs and cable companies could be given access to USF funds (assuming they also contribute to the fund) to encourage them to expand their service markets. Programs similar to those of the Department of Agriculture and the Rural Health Care Program (see Box 5.3) that target remote and rural areas could be created for health care providers and delivery sites in underserved and impoverished urban areas. The feasibility of the USF's High Cost and Low Income Program, providing a basic level of Internet service to individuals, could also be examined. Could the same economic principle that makes telemedicine an affordable solution for prison health care also apply to home-based telemedicine services? Greater involvement by DHHS—and others—in these initiatives may be necessary to ensure that Internet access programs meet the needs of all parties in the health community, including consumers.

INTELLECTUAL PROPERTY PROTECTION

The laws and practices that cover the creation, storage, dissemination, and use of intellectual property shape the way in which digital knowledge resources can be used in health professional education, biomedical research, health care, and the information systems that support health care. Considerable experimentation is under way to develop new models for distributing electronic media over the Internet as the technologies for disseminating and protecting health-related information mature. These trends have implications for both the publication of such information and the practice of distance education in the health sector.

Electronic Publishing

The health community has long been among the most active scientific professions in developing online resources providing access to current information (e.g., the NLM's MEDLINE, which predated similar bibliographic resources for other scientific communities), in part because of the rapid growth in biomedical knowledge. Thus, the health community (biomedical researchers in particular) has a special interest in the evolution of electronic publishing and the mechanisms by which scientists and practitioners, as well as the public, will be able to gain unfettered access to

information on research, health care, and disease. Publishers also have a special interest in the evolution of electronic publishing, which is seen as both a blessing and a curse. For example, once an electronic book is produced, the cost of distributing copies is small compared to the cost of printing and distributing the print equivalent (see Shapiro and Varian, 1999). But this very quality of electronic publishing is also its burden—publishers fear that their revenues will shrink as readers simply “cut and paste” the information they want, rather than purchasing the rights to use it.

The remarkable expansion of Web-based publishing and the proliferation of Internet information centers, such as America Online and MEDSCAPE, have altered the expectations of consumers, health care providers, and biomedical researchers concerning their rights to use and retain information available on the Internet. These expectations sometimes conflict with publishers’ controls. Technology offers a partial solution by providing the means to protect digital content, but its use can challenge fundamental notions of concepts such as fair use. For example, simple access controls can enable distributors to implement licensing agreements that restrict the use of Internet-based health information resources to a set of users or a specific IP domain. Such agreements would not permit physicians and other health care providers to deliver information or virtual library services to their patients in the course of delivering health care. Digital rights management technologies, such as those offered by Intertrust Corp., IBM, and Xerox, also enable producers to specify rules for the use of information products, but they do not necessarily ensure fair use.¹¹

As health professions, schools, and research centers build Internet-based curriculum materials and databases, they become publishers themselves, subject to many of the concerns that commercial publishers face, and they become involved in a competitive marketplace that strains traditional agreements about the free exchange of scientific information among peers. Recently, the National Institutes of Health (NIH) proposed the establishment of PubMed Central, a repository for the barrier-free electronic distribution of life sciences research reports.¹² If implemented as proposed, the site would alter the way scientific information is exchanged among colleagues, bring biomedical research reports within easy reach of consumers, and raise additional questions about the best way for universities to manage and market intellectual property. Electronic publishing—whether it be article preprints posted on PubMed Central, laboratory data deposited in international scientific databases, or a research set of health outcomes information (scrubbed of personal identifiers) distributed to colleagues from one’s personal desktop workstation—challenges many existing practices. If the full potential of the Internet is to be real-

ized, public and private agencies will need to rethink the principles of ownership, authorship, and priority that have guided scientific communication for centuries (CSTB, 2000).

Distance Education

Two types of intellectual property issues arise out of attempts to use the Internet to deliver distance education, which is becoming increasingly popular in the health sector. For many years, a large group of stakeholders participated in the Conference on Fair Use (CONFU), which sponsored discussions on topics such as intellectual property protection and copyright as applied to digital images, distance learning, educational multimedia, electronic reserves, interlibrary loans and document delivery, and in-library use of computer software. The group's final report, issued in November 1998, contained voluntary guidelines covering several issues but no consensus that led to action.

The issue has been brought to the fore again by passage of the Digital Millennium Copyright Act (DMCA, P.L. 105-304), which became law in September 1998. The DMCA directed the registrar of copyright to consult with copyright owners, nonprofit educational institutions, and nonprofit libraries and archives and to submit recommendations to the U.S. Congress on how to promote distance education through digital technologies, including interactive digital networks.¹³ The recommendations must maintain "an appropriate balance between the rights of copyright owners and the interests of users" and may include legislative changes. The U.S. Copyright Office issued its report in May 1999.¹⁴ The fair use issue needs to be resolved in a way that enhances the ability of health professional educators to take full and fair advantage of the Internet.

The second problem area concerns the rights of teaching faculty to retain the products of their intellectual work. Traditionally, universities have allowed faculty to retain intellectual property rights (and royalties) for the textbooks and monographs they write as part of their scholarly work. But patents that result from laboratory research have been handled differently; faculty members have been allowed to retain only a portion of revenues that result from the sale and marketing of their inventions. The increasing market for computer-based learning tools and courses offered across the Internet has sparked debate over whether electronic curriculum materials should be treated as textbooks or patents. Because the production of multimedia resources and Web-based course materials often involves many staff members, some universities have declared these materials to be work for hire, such that revenues are to flow to the institution rather than to any individual faculty member. A shared-ownership approach for multimedia works, based on the patent model, has been

implemented at some universities as an alternative to the work-for-hire approach.

Several options could be pursued to resolve these issues. First, consideration could be given to the development of fair use provisions within the health domain. A set of practical guidelines and licensing exceptions could encourage and facilitate the use of information for health care while protecting the rights of the owners of that information. This approach has worked well in the educational sector, where agreement was reached on the degree to which students and researchers can reproduce copyrighted material for educational purposes.

Second, a distance education framework could be put in place that removes barriers to the delivery of instruction over the Internet. Without a set of clear, workable guidelines, schools that educate health professionals also are unlikely to develop electronic courses and learning resources that come up to the standards of the courses and resources available to students in the traditional classroom. The same can be said for the societies and associations that provide continuing medical and professional education to their members and health care organizations that wish to provide patient education programs over the Internet. Demonstrations of the practical application of the CONFU guidelines in several health settings could help assess the adequacy of the guidelines for health professions education.

Third, universities could develop a standard approach to the ownership of electronic curriculum materials that balances the rights of the institution with those of the creators of the materials and is compatible with other intellectual property policies. If the proposed PubMed Central makes access to scientific knowledge truly barrier-free for both physicians and their patients, then it could fundamentally reshape the way scientific information is brought to bear on health. However, the mechanisms for assuring the quality and scope of materials published in PubMed Central have not yet been developed. The health care community needs to take an active role in devising these mechanisms.

Congressional activity may affect the use of distance education materials. For instance, a bill in the House of Representatives, the Collections of Information Antipiracy Act (H.R. 354, Coble), proposes copyright protection for owners of collections of electronic information. The bill is modeled on EU law (the data directive) that protects those who make "substantial investment" in developing databases, giving them protection against extraction or reuse (without permission) for a 15-year period. To obtain protection under the EU law, the United States must pass equivalent protections. Opponents of the bill see the proposed legislation as an infringement of existing fair use rights and believe it will have a chilling effect on scientific research.¹⁵

REGULATIONS AFFECTING ELECTRONIC DELIVERY OF HEALTH SERVICES

Various communications technologies (e.g., telephones, satellite links, and private networks) have long been used to deliver health services at a distance—that is, telemedicine. But a number of regulatory issues have impeded the expanded use of information technology for telemedicine. Ongoing issues include payment policies, especially those used in federal programs (Medicare and Medicaid); professional licensure requirements; and standards for malpractice liability, which have impeded attempts to deliver services across state lines. The difficulties that have arisen in the regulatory environment are indicative of the broad range of potential obstacles that could arise as the Internet becomes more widely used for the delivery of health services.

Payment Policies

The lack of suitable mechanisms of payment for telemedicine sessions is a significant impediment to the use of information technology within the fee-for-service environment that still dominates much of health care. Traditional fee-for-service insurers pay for health care in accordance with specialized rules outlining the particular services that will—and will not—be reimbursed. This model stands in contrast to capitated care systems, in which providers are paid based on formulas such as the number of patients under their care rather than the individual services rendered. Providers in capitated systems have incentives to use efficacious and cost-effective approaches, without regard to specific reimbursement policies. They can be expected to view Internet-based service delivery not as competing with face-to-face service delivery but as a tool to reduce the customer service burden on their facilities and confer a competitive advantage. Hence, payment policies are not a significant obstacle to telemedicine in capitated care systems. However, private insurers that rely on a fee-for-service model will need to revisit their payment policies.

Existing payment policies within the fee-for-service framework have only begun to expand coverage into the telemedicine arena.¹⁶ Prior to the Balanced Budget Act of 1997 (P.L. 105-33), the Health Care Financing Administration (HCFA), which pays for Medicare and Medicaid services and thus has a significant effect on payment policies throughout the industry, reimbursed only consultations that occurred in real time and involved face-to-face encounters between physicians and patients (although diagnostic teleradiology and telepathology were covered). The act required HCFA to expand its coverage of telemedicine services under the Medicare program beginning in January 1999, but critics contend that

because the new rules fail to recognize telemedicine as just another way of delivering existing health care services they make artificial distinctions between telemedicine sessions and traditional face-to-face encounters.¹⁷ Under its new rules, which are still being evaluated, HCFA will reimburse telemedicine services according to the following criteria:¹⁸

- The Medicare beneficiary (i.e., patient) must reside in—or be presented from—a designated health professional shortage area (HPSA), typically rural or inner city locations that have difficulty attracting and retaining physicians and other health professionals. HPSAs are designated based on the total number of physicians practicing in a geographic area. The rules do not distinguish between a patient's access to primary care physicians vs. specialists and may not allow for reimbursement of remote consultations by specialists in some areas.

- Services delivered through telemedicine are limited to initial, follow-up, or confirming consultations in hospitals, outpatient facilities, or medical offices. Reimbursement is not available for all the services that are reimbursable in a face-to-face encounter. Moreover, the rules do not specifically address telemedicine delivered to the home.

- Services must be delivered in real time through interactive audio and video telecommunications systems. With encouragement from the Congress, HCFA is expected to consider the use of store-and-forward technologies in the future, but services delivered in this form are not currently covered. Store-and-forward may be especially valuable in locations that have limited access to affordable broadband technologies.

- The professional providing the teleconsultations must be a physician, physician's assistant, nurse-practitioner, clinical nurse specialist, or nurse-midwife. Clinical psychologists, physical therapists, occupational therapists, and speech therapists—who may be reimbursed in face-to-face encounters—may not be reimbursed for telemedicine.

- The referring professional must be a physician, physician's assistant, nurse practitioner, clinical nurse specialist, nurse-midwife, clinical psychologist, or clinical social worker. Another professional may act in place of a referring professional to present the patient to the consultant, provided that the professional is employed by the referring professional and is a practitioner as defined in the regulation. Registered nurses, licensed practical nurses, and others (including family members for home care) are not allowed to present cases to consulting practitioners.

- As specified by the Balanced Budget Act, the regulations require fees for telemedicine encounters to be shared by the referring and consulting practitioners. HCFA will make a single payment to the consulting professional, 25 percent of which is to be paid to the referring professional.

These rules have allowed greater experimentation with telemedicine reimbursement programs. As of July 1999, Medicare teleconsultation demonstration projects in Georgia, Iowa, North Carolina, and West Virginia were allowing more than 100 hospitals to receive reimbursement for outpatient services involving specialist telemedicine. Reimbursement of telemedicine providers under the Medicaid program was ongoing in at least 11 states, with several others considering such programs. But, as outlined above, the rules do not yet address payment for all the diverse forms of telemedicine that could arise in an Internet-mediated environment. For example, they do not cover store-and-forward systems for medical images (either still or video), videoconferencing to the home, or provision of services outside designated shortage areas. Payment policies may have to be modified further to enable these applications as other technical and organizational impediments are overcome and the applications themselves are demonstrated to be efficacious and cost-effective.

A primary requirement for payment—and a current obstacle to it—is firm evidence of the efficacy of telemedicine services. A number of pilot projects have demonstrated the technical feasibility of telemedicine, but rigorous evaluations of the outcomes of care remain sparse.¹⁹ Without a better evidentiary basis and more supportive policies, institutions and providers are understandably reluctant to invest in the equipment and training needed to provide high-quality telemedicine services. Further studies are under way to scientifically evaluate telemedicine applications. The NLM is funding dozens of demonstration projects and has explicitly called for evaluations of efficacy.²⁰ Other demonstration projects funded by federal agencies such as DHHS, the Veterans Administration, the Department of Defense, and the National Aeronautics and Space Administration may help create a stronger scientific basis on which to determine the efficacy of various telemedicine services, leading eventually to more liberal reimbursement policies.

Liability and Licensure

Use of the Internet creates a host of new issues regarding licensure and liability. For instance, it is not clear what rules and regulations control the online prescription of medications or confer credentials on remote interpreters of radiographs. Some people fear that the Internet could allow quacks to masquerade as credentialed physicians. At the same time, existing licensure and liability rules do not adequately cover current and potential Internet applications in the health sector. State-based liability for malpractice poses an additional barrier to widespread use of the Internet for telemedicine services across state lines. Under the principle of long-arm jurisdiction, any state court can claim authority over trans-

actions in another state if it can prove that it has jurisdiction over that transaction (Alberts et al., 1998). The Internet confounds the matter of which state has jurisdiction in the delivery of health services and information because the transaction is distributed among two or more locations. For example, if a patient undergoes a remotely managed surgical procedure at a medical center in a neighboring state, which state has jurisdiction if there is a malpractice claim—the patient's home state, the state in which the medical center is located, or the surgeon's home state? Must the surgeon be licensed in all three states? A mechanism is needed for crafting creative solutions to situations such as these.

Similarly, under the present rules, care providers—whether physicians, nurses, or physical therapists—are licensed to practice in individual states. With telemedicine or Internet-based health care more generally, it becomes practical for providers to practice across state lines. But this would be illegal, and telemedicine programs that do cross state lines must ensure that a licensed physician is located at both ends of the transaction. This practice not only raises the cost of telemedicine consultations or procedures (and all of the cost may not be reimbursed, as noted above) but also makes it difficult to implement telemedicine services in small towns that have other health professionals but lack physicians. These are often the very areas that have the greatest need for telemedicine services. Indeed, one of the reasons the telemedicine program at East Carolina University has been so successful is because it operates only within the state of North Carolina, thereby precluding the need to engage two physicians simultaneously. The university employs specially trained nurses and physician assistants in rural clinics to present telemedicine cases, helping to reduce costs and extend the reach of the telemedicine services. At the same time, local jurisdictions face strong political pressures to maintain strict licensure provisions. Allowing the interstate practice of medicine could make it easier for larger, more prestigious health care institutions to win business from local institutions, hurting the local economy.

A number of mechanisms have been proposed to address licensing concerns. The Western Governors' Association, for example, proposed several options that states could pursue voluntarily, including explicit exemptions for remote consultations under specific conditions; participation in a uniform regional licensure program that would provide a limited license valid in any of the participating states; or institutional or network licensure for all network physicians (Gilbert, 1995; WGA, 1995). An alternative that has been proposed is the development of model legislation for encouraging uniformity in state licensing rules (IOM, 1996). A more recent report from the Southern Governors' Association called for continued examination of the licensure issue to develop a strategy for

evaluating interstate licensure (Southern Governors' Association, 1999). Clearly, continued effort is needed in this area.

The Congress and executive branch agencies continue to express interest in telemedicine. The Information Infrastructure Task Force and the Joint Working Group on Telemedicine are chartered to consider telemedicine applications. Accordingly, the Congress has begun to consider corrective legislation in some of the pertinent areas (see Box 5.4). Other groups, too, have begun to address regulatory issues associated with telemedicine. The Southern Governors' Association, for example, identified payment and licensure as two of the primary impediments to telemedicine and called for the establishment of a national licensing system, although not necessarily one run by the federal government (Southern Governors' Association, 1999). Regardless of the particular mechanisms used, it is clear that regulatory barriers such as liability and licensure need to be remedied before the Internet can be exploited fully to expand the availability of health care.

FEDERAL SUPPORT FOR HEALTH-RELATED INFORMATION TECHNOLOGY RESEARCH

The combined efforts of both private and public organizations will be required to ensure the robust development of health care applications of the Internet. To date, federal health agencies have exhibited uneven interest in pursuing Internet-based applications or in funding information technology research that could support health applications. By comparison, the Department of Defense and the Veterans Administration—both of which operate large health care delivery systems but are not federal health agencies—have invested considerable resources in the development of telemedicine programs and infrastructure for sharing patient health records.²¹ Some entities within the DHHS, most notably the NLM but also other elements of the NIH and the Agency for Healthcare Research and Quality, have invested in research on health-related applications of computing and communications technologies, but DHHS itself has not made use of the Internet in health-related activities a priority. As one DHHS official noted, the department has considerable expertise in information technologies, but that expertise is dispersed throughout its many agencies, with no formal mechanisms for cross-fertilization and exchange of ideas. No clear advocate exists within the DHHS for Internet-based applications.²² As a result DHHS and its constituent agencies sometimes lag behind the institutions they serve (e.g., they do not accept Medicare claims via the Internet) and sometimes mirror the same conservative approach that is characteristic of many health care organizations. For example, HCFA's strategic plan for information systems, dated July 1998,

BOX 5.4 Legislation Related to Telemedicine

A number of bills related to telemedicine were introduced in the 106th Congress. Among those most relevant to the present report are the following:

- S. 770: *Comprehensive Telehealth Act of 1999*. This bill would provide payment under the Medicare program for telehealth services and for other purposes. It amends the Balanced Budget Act of 1997 with respect to Medicare reimbursement of telehealth to (1) include reimbursement for store-and-forward technologies; clinical psychologist and physical, occupational, and speech therapist practitioner services; and items and services covered under Medicare part B (Supplementary Medical Insurance) that are provided via telecommunications systems, (2) extend telehealth coverage to all rural areas (currently, only those designated as health professional shortage areas under the Public Health Service Act are covered), (3) allow any health care practitioner acting on instructions from the referring physician or practitioner to present the Medicare beneficiary to the consulting physician or practitioner for the provision of items and services, (4) prohibit the referring physician and the practitioner from receiving any reimbursement for such presentation other than the payment that the referring physician shares with the consulting physician, and (5) provide that payment for items and services shall include payment for all current procedural terminology billing codes covered under Medicare. It also directs the secretary of Health and Human Services to study, and report to the U.S. Congress on, issues associated with cross-state licensure of health care providers (e.g., numbers of cross-licensed practitioners, status of reciprocal agreements, and state-led efforts to ease licensure burdens) and to provide specified financial assistance to eligible telehealth networks to expand access to health care services for individuals in rural and frontier areas. The bill was referred to the Senate Finance Committee for consideration in April 1999. No further major action had been taken as of February 2000.

- H.R. 1344: *Triple-A Rural Health Improvement Act of 1999*. This bill would promote and improve access to health care services in rural areas. Title V amends the Balanced Budget Act of 1997 with regard to telehealth services, among other changes, to (1) extend Medicare reimbursement for such services to all rural Medicare services, including services by physical, occupational, and speech therapists, (2) revise related payment methodology, and (3) add congressional reporting requirements pertaining to the telehealth services program. It also changes the name of the Joint Working Group on Telemedicine to the Joint Working Group on Telehealth and establishes the mission of the working group, among other things, as identifying, monitoring, and coordinating federal telehealth projects and programs. The bill directs the secretary to provide specified financial assistance for expanding access to health care services for individuals in rural frontier areas through the use of telehealth. The bill was referred to the Committee on Ways and Means and the Committee on Commerce in March 1999. As of February 2000, no further action had been taken.

mentions the Internet only twice in passing—despite the report's focus on enhancing timely access to information for a variety of users (HCFA, 1998).

Some signs of change are visible. Since issuing its strategic plan for information technology, HCFA has revised its security policy to allow Internet-based exchanges of information among its own entities after the information has been received from care providers and health plans. HCFA also is sponsoring an ongoing pilot program to compare several different alternatives for Internet-based submission of claims. Depending on the findings of the pilot study and the results of a cost-benefit analysis of claims transmission, HCFA will decide whether to allow Internet transmission on an operational basis.²³ The Centers for Disease Control and Prevention (CDC) recently initiated a project to use the Internet as part of a Health Alert Network to detect and respond to bioterrorist attacks, although some staff contend that this project is significantly underfunded.²⁴

Federal health agencies have played only a limited role in supporting research and development (R&D) for new Internet-based capabilities. Many R&D and demonstration projects have been funded by the NLM and other NIH institutes, but investigations of new computing and communications technologies that are motivated by health-related goals have been only a minor element in the portfolio of the U.S. biomedical research community. Indeed, the biomedical research community in general, and the NIH in particular, have been accused of failing to provide their fair share of support for fundamental research in information technology, especially in comparison with other mission-oriented federal agencies (PITAC, 1999). Of the \$2.3 billion spent by the federal government on computer science research in 1999, just \$120 million came from DHHS, and none came from the Department of Veterans Affairs (Table 5.1).²⁵ The Commerce, Defense, and Energy Departments and the NSF each funded more computing research, despite the fact that DHHS's overall research budget of \$13 billion exceeds the research budgets of the four other agencies combined.²⁶ This funding imbalance provides further evidence that federal health agencies have been far less active than their counterparts in the education, defense, scientific, and library communities in pursuing initiatives related to the Internet.²⁷

The reasons for the comparative lack of support for computer science research within the DHHS are manifold, but perhaps the underlying problem is that the department as a whole has not embraced information technology (IT) as fundamental to its mission. Federal agencies that have allocated major portions of their R&D budgets to fundamental IT research tend to view IT as an integral part of their particular missions, whether national security, scientific research, or industrial competitiveness. Many

TABLE 5.1 Estimated Federal Funding for Research by Selected Agencies, 1999

Department/Agency	All Fields (millions of \$, % of total)	Computer Science (millions of \$, % of total)	Computer Science Research as % of Agency's Research in All Fields
Department of Health and Human Services	12,982 (39.3)	120 (5.3)	0.9
Department of Defense	4,089 (12.4)	864 (38.3)	21.1
Department of Energy	4,128 (12.5)	656 (29.1)	15.9
National Science Foundation	2,655 (8.0)	395 (17.5)	14.9
Department of Commerce	850 (2.6)	85 (3.8)	10.0
All federal agencies	32,992 (100)	2,255 (100)	6.8

SOURCE: National Science Foundation (2000), Table C-22.

discoveries funded by DOD, DOE, and NSF have found their way into systems applied in the health disciplines. Almost all of DHHS's research funding is obligated to biological and medical sciences; IT is seen primarily as infrastructure to support these other activities. A recent NIH report on biomedical information science and technology recognizes the importance of computing in biomedical research (NIH, 1999), but it assumes that computer science is a source of tools for enabling biomedical research rather than a field of intellectual inquiry worthy of further investigation by biomedical researchers. Until government and industrial research organizations see the health disciplines as a potential source of innovation for new information technologies, fundamental informatics research is not likely to emerge from the traditional biomedical sciences; instead, health informatics research probably will continue to emphasize demonstration projects and applications.

There are several arguments for involving health agencies more closely in information technology research that is motivated by health needs. As noted throughout this report, IT is increasingly critical to the continued success of the nation's health enterprise. Achieving the goal of promoting a healthy citizenry—as well as assuring equitable access to affordable, high-quality health care—is likely to depend on the effective use of new and emerging information technologies, including the Internet. Although the capabilities required of the Internet for health applications do not differ significantly from those for applications in other areas where the government has a mission, research programs motivated by health needs probably would emphasize characteristics and capabilities that would not otherwise receive sufficient attention but that could be widely applicable, just as earlier work in expert systems that was motivated by

health needs resulted in important advances that have proven widely applicable (Box 5.5).

Efforts to involve the DHHS and its constituent agencies more closely in Internet-related research will require careful thought. Structuring mission agency research programs in information technology is vexing, especially given the history of successful research support by the DOD and NSF.²⁸ The DHHS could initiate its own network research program, but it would most likely need to coordinate its activities with other agencies to avoid unnecessary duplication. Lack of experience could be a barrier to this approach; the department has not been active in basic network (or computer science) research for some time, so it may lack the ability to identify relevant and challenging research areas, to select promising proposals, and to manage the research. An alternative approach would be for DHHS to encourage and foster collaborative R&D between health informatics investigators and technical experts in communication and information technologies who are actively engaged in networking R&D. This approach is more likely to lead to solutions that advance the information technology infrastructure in general and less likely to strand health informatics in an isolated technological corner.

WORKFORCE ISSUES

If governments, health care organizations, academic institutions, and professional groups need to attract more effective leaders in the strategic uses of information technology, where will they find them? The number of individuals who understand both the biomedical milieu and the technologies relevant to computing and the Internet is remarkably small. The lack of trained individuals at the interface may also help to explain why the biomedical community has lagged other fields in understanding and adopting computing and communications technology, especially when direct use by health professionals is required. Although the NLM has funded graduate training programs in informatics for more than two decades and some computer science departments and library schools have set up programs to train health technology specialists, the supply of medical information scientists and professionals from these programs does not meet the demand. Directors of the current medical informatics training programs find that their graduates are highly sought after by industry (e.g., the pharmaceutical and health information system industries), health care organizations (e.g., managed care groups, hospitals, and multi-specialty practices), and academic informatics groups with training or research missions.

The problem is not necessarily unique to health care computing. Companies in all sectors of the economy report difficulties in attracting and

BOX 5.5

Medically Motivated Research on Expert Systems

The health community played an important role in the rise of expert systems, which are software environments that analyze data and give advice based on the internal encoding of human expertise (Duda and Shortliffe, 1983). Although early work involved applications in a large number of fields, work on medical expert systems starting in the 1970s had a particularly important impact. This work not only developed technologies and systems that were beneficial to health care and biomedical research but also pushed the frontiers of computer science, developing new knowledge and technologies that could be incorporated into systems for a range of other applications. The medical focus of this work uncovered real-world problems that could be solved only through fundamental computing research.

The real-world issues that arise in medicine pressed the limits of existing software technologies in challenging ways (Clancey and Shortliffe, 1984). For example, medical applications required new approaches to managing and modeling uncertainty—specifically, the uncertainties associated with causality, diagnosis, and treatment. Biological systems were not as well understood as the engineering systems being analyzed by expert systems at the time. Furthermore, because of the diverse types of data that must be considered in diagnosis and the complexity of patient records, medical expert systems could not rely on a simple question-and-response approach to data gathering in support of decision making (an approach that suited many simpler systems). Hence, medical informatics researchers pioneered the integration of expert systems with large databases, such as computer-based medical records. The complexity of medical decision making and the difficulties involved in embedding decision-support tools into the work flows of practicing medical professionals drove further advances in the representation of large and complex knowledge bases and acceptable modes of interaction with users. The work simultaneously forced the community to take a fresh look at the psychological underpinnings of human problem solving in medicine, leading to productive collaborations among physicians, cognitive scientists, and computer scientists.

Ultimately, many of the techniques developed by the medical computing research community, generally with funding from the National Institutes of Health (NIH) or the predecessor to the Agency for Healthcare Research and Quality (AHRQ), were adapted for use in other application areas. In fact, during the 1980s, interest in expert systems soared, largely in areas outside of medicine (Computer Science and Telecommunications Board, 1999a). It is clear that the results of early NIH- and AHRQ-supported research in this area had a major impact on corporate America and the economy in the decades that followed (Feigenbaum et al., 1988).

Ironically, the adoption of expert systems in health care has been slower. Only now is an infrastructure, culture, and health-financing climate emerging that is likely to allow expert systems to have a substantial impact in the clinical setting. The recent emphasis on the implementation of clinical guidelines, reduction in practice variations, and adoption of evidence-based practice has made expert systems seem more attractive. In addition, the introduction of electronic medical records, distributed computing, and networking within institutions has made it possible to integrate decision-support tools in ways that do not require the stand-alone consul-

tation model typical of the earlier medical expert systems. Thus, the true payoff, in terms of the mission and goals of NIH and AHRQ, will have occurred 25 to 30 years after they began to support research in this area, although there have been interim spin-offs in other areas. If the biomedical research agencies of the 1970s had demanded a 2-year turnaround time for their investments in information technology research, the work that has produced so many benefits today never would have been initiated.

retaining sufficient numbers of information technology workers. A recent report by the President's Information Technology Advisory Committee (1999) contends that the supply of information technology workers does not meet the demand in almost all segments of society.²⁹ The Computer Science and Telecommunications Board of the National Research Council has been charged by the Congress to more fully evaluate the issue.³⁰ As the advisory committee report argues, "It is crucial that we produce a continuous supply of well-trained, high-quality professionals in engineering and computer and information science, not merely skilled users, but researchers, creators, and designers of advanced technology" (PITAC, 1999). As with information technology fields in general, niche areas requiring expertise in health care, biomedicine, computing, and communication also appear to be undersupplied. There are important scientific issues to address, practical systems to be built, and strategies to be put in place. Health care will be constrained until there are many more individuals with these combinations of talents.

One option for addressing this concern is to expand educational opportunities in health-care computing and communications. Existing training programs in health informatics could be expanded, and new emphases or tracks could be developed in related disciplines. Alternative training tracks in areas of application ranging from clinical medicine and nursing to bioinformatics and health system management also could be developed. At the same time, educational opportunities could be created within schools of the health professions and computer science departments. Beyond such training programs, which would be designed to develop a cadre of researchers who operate at the nexus of health and computer sciences, initiatives could be taken to infuse some degree of informatics training throughout the health sciences curricula. Health professionals need to become part of the information culture that will define the century ahead, and their educational opportunities, in both professional school and continuing education programs, must be responsive to that need. Such training needs to emphasize areas of overlap among

health care, biomedicine, computing, and communications—not simply computer literacy as traditionally, and narrowly, construed.³¹

CONCLUSION

As this chapter illustrates, the availability of useful information technologies is not enough to ensure their effective application in health care. Myriad policy issues will most likely need to be resolved concerning patient privacy, access to health care and Internet services, intellectual property protection, payment mechanisms for telemedicine, and human resources. The effective implementation of new information technologies in complex environments such as the U.S. health care system will require vision, commitment, and leadership at the highest levels; a well-funded research agenda; effective policies for developing the necessary workforce; and a grassroots community of capable participants. To make decisions related to Internet technologies in health and health care, public and private policy makers need sophisticated analyses and information that goes beyond traditional reporting on a narrow set of facts. The issues are emotionally and politically charged. Their resolution will require the concerted efforts of many public and private sector organizations, including government agencies, companies, and professional associations. Without deliberate, sustained action, the fundamental conflicts seen in these policy debates will keep the Internet from fulfilling its promise in health care.

REFERENCES

- Alberts, R.J., A.M. Townsend, and M.E. Whitman. 1998. "The Threat of Long-Arm Jurisdiction to Electronic Commerce," *Communications of the ACM* 41(12):15-20.
- Appavu, S.I. 1997. *Analysis of Unique Patient Identifier Options*. Report prepared for the Department of Health and Human Services, November 24. Available online at <<http://www.ncvhs.hhs.gov/app0.htm>>.
- Bureau of National Affairs, Inc. (BNA). 1998. "EU Privacy Directive Will Take Effect Even Without Implementing Legislation," *Electronic Commerce and Law Report*, November 17. Available online at <<http://zeus.bna.com/e-law/articles/top0335.html>>.
- Clancey, W.J., and E.H. Shortliffe, eds. 1984. *Readings in Medical Artificial Intelligence*. Addison-Wesley, Reading, Mass.
- Clinton, W.J. 1999. *Narrowing the Digital Divide: Creating Opportunities for All Americans in the Information Age*. Memorandum to executive departments and agencies, December 9.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1994. *Realizing the Information Future: The Internet and Beyond*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1997. *For the Record: Protecting Electronic Health Information*. National Academy Press, Washington, D.C.

- Computer Science and Telecommunications Board (CSTB), National Research Council. 1999a. *Funding a Revolution: Government Support for Computing Research*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1999b. *Being Fluent with Information Technology*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 2000. *The Digital Dilemma: Intellectual Property in the Information Age*. National Academy Press, Washington, D.C.
- Consumer Federation of America (CFA). 1999. *Transforming the Information Superhighway into a Private Toll Road: The Case Against Closed Access Broadband Internet Systems*. CFA, Washington, D.C., September. Available online at <<http://www.consumerfed.org/broadbandaccess.pdf>>.
- Duda, R.O., and E.H. Shortliffe. 1983. "Expert Systems Research," *Science* 220:261-268.
- Eng, Thomas R., Andrew Maxfield, Kevin Patrick, Mary Jo Deering, Scott C. Ratzan, and David H. Gustafson. 1998. "Access to Health Information and Support: A Public Highway or a Private Road?" *Journal of the American Medical Association* 280(15):1371-1375.
- Feigenbaum, E.A., P. McCorduck, and H.P. Nii. 1988. *The Rise of the Expert Company: How Visionary Businesses Are Using Intelligent Computers to Achieve Higher Productivity and Profits*. Times Books, New York.
- Gilbert, F. 1995. "Licensure and Credentialing Barriers to the Practice of Telemedicine," pp. 27-35 in *Telemedicine Action Report: Background Papers*. Western Governors' Association, Denver.
- Goldberg, A.S. 1999. "Taking Healthcare to the Patient—Telemedicine Delivers," *Health Law Digest* 27(7):3-10.
- Gotcher, R. 1999. "End-to-End E-Publishing Service Announced," *InfoWorld Electric*, November 10. Available online at <www.infoworld.com/dgi/bin/displayStory.pl?991110.icpublish.htm>.
- Health Care Financing Administration (HCFA). 1999. *Fact Sheet: Medicare Payment for Teleconsultation in Rural Health Professional Shortage Areas*. HCFA, Baltimore, Md., May.
- Health Care Financing Administration (HCFA). 1998. *HCFA's Information Technology Vision*. HCFA, Baltimore, Md., July.
- Institute of Medicine (IOM). 1997. *The Computer-Based Patient Record: An Essential Technology for Health Care*, rev. ed. R.S. Dick, E.B. Steen, and D.E. Detmer, eds. National Academy Press, Washington, D.C.
- Institute of Medicine (IOM). 1996. *Telemedicine: A Guide to Assessing Telecommunications in Health Care*. National Academy Press, Washington, D.C.
- Institute of Medicine (IOM). 1994. *Health Data in the Information Age: Use, Disclosure, and Privacy*. Molla S. Donaldson and Kathleen N. Lohr, eds. National Academy Press, Washington, D.C.
- Mandl, K.D., S. Feit, B.M.G. Pena, and I.S. Kohane. 2000. "Growth and Determinants of Access in Patient E-mail and Internet Use," *Archives of Pediatrics and Adolescent Medicine*, in press.
- Mueller, Milton L., Jr. 1997. *Universal Service: Competition, Interconnection, and Monopoly in the Making of the American Telephone System*. AEI Press, Washington, D.C.
- National Committee on Vital and Health Statistics. 1999. *Second Annual Report to Congress on Implementation of Administrative Simplification*. July 22. Available online at <www.ncvhs.hhs.gov/yr2-rpt.htm#progress>.
- National Institutes of Health (NIH). 1999. *The Biomedical Information Science and Technology Initiative: Report of the Working Group on Biomedical Computing*. June 3. Available online at <www.nih.gov/welcome/director/060399.htm>.

- National Research Council (NRC). 1999. *A Question of Balance: Private Rights and the Public Interest in Scientific and Technical Databases*. National Academy Press, Washington, D.C.
- National Science Foundation (NSF). 2000. *Federal Funding for Research and Development: Fiscal Years 1998, 1999, and 2000*. Early release tables, National Science Foundation, Arlington, Va.
- National Telecommunications and Information Administration (NTIA), U.S. Department of Commerce. 1999. *Falling Through the Net: Defining the Digital Divide*. Available online at <<http://www.ntia.doc.gov/ntiahome/fttn99/>>.
- President's Information Technology Advisory Committee (PITAC). 1999. *Report to the President, Information Technology Research: Investing in Our Future*, February. Available online at <<http://www.ccic.gov/>>.
- Shapiro, C., and H.R. Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press, Boston, Mass.
- Southern Governors' Association, Task Force on Medical Technology. 1999. *From Promise to Practice: Improving Life in the South Through Telemedicine*, Final Report, Washington, D.C., September.
- Strode, S.W., S. Gustke, and A. Allen. 1999. "Technical and Clinical Progress in Telemedicine," *Journal of the American Medical Association* 281(12):1066-1068.
- Tracy, Joe, Thelma McClosky-Armstrong, Rob Sprang, Sam Burgiss, Jim Reid, and Donna Hammack. 1999. "Medical Reimbursement for Telehealth Encounters." Position paper available from the University of Missouri Health Sciences Center, Columbia, Mo., October 11.
- Universal Service Administrative Company (USAC). 1999. *Report to the FCC: Evaluation of the Rural Health Care Program*. Universal Service Administrative Company, Madison, Wisc., March 5.
- U.S. Copyright Office. 1999. *Report on Copyright and Distance Digital Education*, May. Available online at <http://lcweb.loc.gov/copyright/cypub/de_rprt.pdf>.
- Westat. 1999. *Evaluation of the Telecommunications and Information Infrastructure Assistance Program for the 1994 and 1995 Grant Years*. Report prepared for the U.S. Department of Commerce, National Telecommunications and Information Agency, February. Available online at <<http://www.ntia.doc.gov/otiahome/tiiap/index.html>>.
- Western Governors' Association (WGA). 1995. *Telemedicine Action Report*. Western Governors' Association, Denver, Colo.
- White House, Office of the Press Secretary. 1999. "Remarks by the President on Medical Privacy," Press release, October 29.

NOTES

1. As a reviewer of an early draft of this report noted, issues other than technology, organizational uncertainty, and public policy will determine the degree to which the Internet finds application in the health sector. The Internet must be applied properly to create solutions that work in health applications before wide adoption can be expected. The Internet has been accepted by the public as a tool for gaining insight into illness and health issues because it delivers value. Clinicians use the World Wide Web for MEDLINE searches because they obtain some value from it. On the other hand, clinicians do not use the Internet to deliver care because valuable, usable, affordable, and practical Internet-based solutions have yet to be built. In other words, public policy is a very important factor in the acceptance of Internet technology in biomedicine and health care, but it may not be the first hurdle or the highest.

2. One of the primary benefits that could be lost because of concerns over privacy and security is the capability to compile comprehensive health records of individual patients from a number of different sites at which they were treated. At present, personal health information usually resides at the health care provider organization. For many people, there is no comprehensive health record that merges information from numerous care sites, a gap that can compromise the quality of care received subsequently. Such longitudinal records also would have value in public health and health services research.

3. The text of the proposed regulations, as well as a summary of their contents, is available online at <http://aspe.hhs.gov/admsimp/>.

4. For additional information on the privacy implications of universal health identifiers and standards-setting activities in this area, see Appavu (1997), CSTB (1997), and National Committee on Vital and Health Statistics (1999).

5. The full title of the directive is Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the Legal Protection of Databases, 1996 O.J. (L77) 20. The full text of the directive is available in NRC (1999), Appendix D.

6. See <http://www.epic.org>; main news feature (accessed May 5, 1999). For information on the safe harbor proposal, see <http://www.ita.doc.gov/media/419data.htm>.

7. This gap has been characterized as a "digital divide" between "information haves" and "information have-nots." See NTIA (1999).

8. The Computer Science and Telecommunications Board (CSTB) of the National Research Council has a project under way to examine technology, business, and policy issues affecting the deployment of broadband technologies for the so-called "last mile" to the home. Additional project information is available online at www.cstb.org and at www.nationalacademies.org, under the heading "Current Projects."

9. The 1996 act was the first to codify the notion of universal service (Mueller, 1997).

10. The NLM grants cover gateway and associated connection hardware; internal access equipment, such as personal computers and local area network costs, are expected to be provided by the institution(s). In 1999, seven awards were made for \$232,000. The NLM also has awarded approximately \$6.7 million to nonprofit health centers since 1996 for telemedicine and Next Generation Internet projects. Additional information on NLM's infrastructure programs is available online at <http://www.nlm.nih.gov/ep/connect.html>, <http://www.nlm.nih.gov/research/telemedinit.html>, and <http://www.nlm.nih.gov/research/ngiinit.html>.

11. For information on Xerox's rights management technology, see Gotcher (1999). Information on Intertrust's products is available at www.intertrust.com.

12. Information on this announcement is available online at <http://www.nih.gov/welcome/director/pubmedcentral/pubmedcentral.htm>.

13. The DMCA brings the United States into closer compliance with the World Intellectual Property Organization (WIPO) treaty. It has numerous provisions. First, it assigns liability to online service providers for acts of subscribers who infringe on the intellectual property rights of others. Universities, some of which have medical schools, are determining how best to comply with the act in a way that enables them to qualify for liability limitations. Second, the DMCA outlaws the circumvention of technical protection systems. Exceptions to the act are permitted for educational fair use, reverse engineering to support interoperability, protection of personal privacy, and security testing. Third, the DMCA required the U.S. Copyright Office to determine whether any adverse effects on fair use had been observed after a 2-year moratorium.

14. See U.S. Copyright Office (1999), available online at http://lcweb.loc.gov/copyright/cpy/pub/de_rprrt.pdf.

15. Concerns regarding the possible effect of the EU Directive on scientific research are discussed in a 1999 report from the National Research Council. See NRC (1999).

16. Blue Cross-Blue Shield plans in Iowa, Kansas, and Montana have already developed payment policies for some forms of telemedicine, and North Dakota has announced payment plans that include telemedicine, but even these policies may need to be revised as new applications emerge (Goldberg, 1999).

17. For further elaboration on these points, see Tracy et al. (1999).

18. This summary of coverage rules is derived from Goldberg (1999) and from HCFA (1999).

19. For a discussion of the history and challenges of evaluating telemedicine applications, see Institute of Medicine (1996) and Strode et al. (1999).

20. Additional information on the NLM's telemedicine programs, including descriptions of funded projects, is available online at <<http://www.nlm.nih.gov/research/teledinit.html>>.

21. The Department of Defense has investigated technologies for remote treatment of soldiers and developed telemedicine systems for military personnel and their families, and the Veterans Administration has developed technologies for sharing medical records among its health care facilities. In addition, NASA has invested in technologies for remote consultation and treatment of astronauts.

22. Margarate Hamburg, DHHS, presentation to the committee on March 1, 1999, Washington, D.C.

23. Further information on the pilot can be obtained from either <www.wedi.org> or <www.afecht.org>.

24. William Yasnoff, associate director for science and acting director, Information Network for Public Health Officials in the CDC Public Health Practice Program Office, in his presentation to the committee on March 1, 1999, estimated that \$200 million would be needed, in addition to the \$28 million already appropriated, to get the Health Alert Network up and running.

25. Of the \$120 million in computer science research funding from DHHS, \$110 million came from the NIH and \$10 million came from the AHRQ.

26. In 1999, research funding for the Departments of Defense, Energy, and Commerce and for the NSF totaled \$4.1 billion, \$4.1 billion, \$808 million, and \$2.7 billion, respectively. These figures include funding for basic and applied research but not development (NSF, 2000, Table C-22).

27. Thomas Kalil, National Economic Council, presentation to the committee on March 1, 1999, Washington, D.C.

28. Two forthcoming reports from the CSTB will address issues of structuring federal and nonfederal support for mission-driven information technology research: *Meeting Society's Needs: Expanding the Scale and Scope of Information Technology Research* and the final report of the CSTB Committee on Computing and Communications Research to Enable Better Use of Information Technology in Government. For additional information on these projects, see <www.cstb.org>.

29. Gauging the supply of and demand for information technology workers is a difficult task, and several analysts have noted faults in data given out by industry groups and the Department of Commerce. The CSTB has initiated a study of information technology workers that is expected to shed more light on this subject.

30. Information on the CSTB project workforce needs in information technology is available online at <www.cstb.org>.

31. Insight into the scope of an expanded definition of computer literacy can be gleaned from a report by the CSTB (1999b).

6

Conclusions and Recommendations

Ensuring that the Internet becomes a suitable, ubiquitous medium for supporting health applications is a challenging task. Not only must the Internet provide connectivity among the participants in health-related information transactions, but it must also ensure that such transactions can occur predictably, efficiently, and without endangering patient safety. Consumers must be able to determine the quality and provenance of the information they retrieve from the Internet. Care providers who access patient records remotely must be assured that the network will be available when and where needed. Administrators must be sure that bill payment and enrollment information is not corrupted as it crosses the Internet. Without proper security protections, use of the Internet to transmit medical records could make personal health information more susceptible to breaches of confidentiality and loss of integrity. Without adequate assurances of network reliability and quality of service (QOS), use of the Internet for remote monitoring of patients, controlling remote medical equipment, or conducting remote medical consultations could impair rather than facilitate the delivery of quality health care. Addressing these concerns demands efforts in many areas, both technical and nontechnical.

This chapter summarizes the committee's main conclusions and recommendations for making the Internet capable of supporting a wide range of health applications. Drawing upon the material presented in Chapters 2 and 3 of this report, it identifies the technical capabilities the Internet must possess in order to provide the security, reliability, and

quality of service that health care applications demand. But it does not stop with recommendations on technical requirements. It also discusses the policy and organizational issues that must be resolved to make the health community more capable of adopting Internet applications in both the short and long term. The capabilities the Internet offers to consumers, care providers, public health officials, health care administrators, and researchers promise to reshape the landscape of the health sector. Accommodating these changes will require actions within individual organizations and across them, enlisting the support of technologists, practitioners, legislators, and the general public. This chapter, accordingly, makes recommendations in areas ranging from identification of the needed technical capabilities of the Internet to specification of the organizational and policy issues that constrain its use in health applications. The recommendations are targeted at policy makers, the networking research community, researchers in health-related fields, health care administrators, and managers of health-related organizations. Taken together, the recommendations aim to provide guidance both on short-term measures that can set the process in motion and on long-term and continuing needs in communications, information technology, and health care.

CONCLUSIONS

***Conclusion 1.* The Internet can support a wide range of applications in consumer health, clinical care, health care financial and administrative transactions, public health, professional education, and biomedical research. The networking capabilities needed to support these applications are not unique, but they do reflect distinctive characteristics of the health environment.**

In each of the domains examined by the committee, the Internet could be used to facilitate communications among parties in ways that can improve quality and efficiency. For example, in the clinical care domain, care providers already use the Internet to search the professional literature for information on particular diseases or to examine evidence-based practice guidelines for managing a particular disorder. As ongoing projects demonstrate, the continued research, development, and deployment of Internet applications will allow care providers to more routinely access electronic medical records held by an affiliated health care organization or to interpret medical images (such as mammograms) sent to them from a remote mammography center. They will be able to offer remote medical consultations to patients in rural areas or to adjust settings on remote dosimetry equipment or pacemakers without establishing fixed, dedicated connections between sites. They will increasingly participate in

online discussions with other care providers to consult on particular cases, sharing medical records and images as needed.

The success of any of these applications depends on a variety of factors, including their cost-effectiveness, ease of use, and ability to improve on existing processes. While some applications are already being used in operational environments across networks other than the Internet, many represent new capabilities that have no parallel on other networks or that have not been fully implemented on a large scale, such as remote medical consultations. As a result, not enough information has been gathered to allow evaluation and comparison, and continued experimentation will be needed to explore and evaluate their true potential, their technical needs, and their real-world operational requirements. A preliminary assessment (Table 6.1) shows a diversity of technical needs, with some commonality, at least within a particular domain (e.g., clinical care, public health) or class of application (e.g., real-time video, file transfers, collaboration). Consumer applications, for example, tend to demand high levels of security to protect confidentiality; clinical applications require a combination of security (to protect confidentiality and data integrity), reliability, and QOS. Virtually all collaborative applications—regardless of whether they are in clinical care, public health, biomedical research, or other domains—demand high levels of QOS, and file transfers in any health application tend to strain technologies for authenticating the identity of communicating parties. Determining which technical capabilities health applications will demand must be viewed as an ongoing process as workers envision, develop, and evaluate new applications.

The technical capabilities demanded by a number of health applications of the Internet exceed those provided by the current Internet, but they are not necessarily unique. Applications in other sectors (e.g., defense, entertainment, financial services) also require better security, reliability, and quality of service. However, when these technical characteristics are combined with factors such as the distributed nature and economic structure of the health industry and the constraints of operating in a health care environment, it can be seen that health does occupy a distinct, if not unique, position. Solutions to problems of authentication and QOS, for example, must scale sufficiently to support the activities of numerous independent health organizations and hundreds of millions of potential users. This argues for full participation by the health care community in defining the research agenda and contributing to its resolution, as the Internet moves forward with new architectures and technical capabilities.

Conclusion 2. Security and availability are critical technical needs for health applications of the Internet and are not adequately met by today's Internet.

TABLE 6.1 Primary Technical Challenges and Limiting Technical Factors in Selected Health Applications of the Internet

	Class of Application		
	Real-Time Video Transmission	Static File Transfer	Remote Control
Application Domain			
Consumer health	Remote medical consultations to the home, office, or wherever the patient is located.	Accessing personal health records online. Downloading educational videos. Sending periodic reports on health conditions to a care provider.	Remote control of patient monitoring equipment.
Clinical care	Remote medical consultations between clinician and patient or between two clinicians.	Transfer of medical records and images (e.g., X rays, MRI, CT scans).	Remote and virtual surgery (a long-term possibility being examined by the defense and space communities).
Administrative and financial transactions	Videoconferencing with real-time sharing of documents.	Payment of services, enrollment of patients, quality reviews, etc. Large medical records and images may be transmitted in support of some claims.	N/A
Public health	Videoconferencing among public health officials during emergency situations, such as chemical or biological attacks by terrorists.	Incident reporting. Collection of information from local public health departments and laboratories. Surveillance for emerging diseases or epidemics. Transfer of epidemiology maps or other image files for monitoring the spread of a disease.	N/A

Information Search and Retrieval	Real-Time Collaboration	Primary Technical Challenges
Online searching for health information or self-assessment guides. Looking for a doctor or hospital.	Collaboration with care providers. Participation in chat groups and support groups.	Protection of sensitive patient information from breaches of confidentiality and from corruption. Ubiquity of access so that all health care consumers can be reached at the location at which care is needed. Tools and policies for validating the quality of online information.
Practice guidelines. Searches of professional medical literature.	Consultation among care providers, such as for surgical planning, which may involve manipulation of digital images.	Access to sustained bandwidth and low latency for remote consultations and collaboration. Security of clinical records. Network reliability. Ubiquity of access for care providers.
Consumer access to information about health plans, participating practitioners, eligibility for procedures, covered drugs in formulary.	N/A	Security to ensure confidentiality and integrity of records. Network reliability sufficient to support regular use for business transactions. Standards for data exchange and definitions of data elements.
Access to published literature and research results as well as epidemiological data. Delivery of alerts and other information to practitioners or other health workers.	Videoconferencing among public health officials during emergency situations, such as chemical or biological attacks by terrorists.	Security to ensure confidentiality and integrity of laboratory reports and other public health information that may contain personal identifying information. Network reliability. Security from information warfare or attacks on the network's physical infrastructure.

continued

TABLE 6.1 Continued

	Class of Application		
	Real-Time Video Transmission	Static File Transfer	Remote Control
Professional education	Distance education: either real-time transmission of lectures or on-demand streaming video with integrated graphics. Real-time consultations with experts about difficult cases.	Accessing electronic medical records from remote clinics. Downloading sets of reference images or prerecorded videos of lectures.	Simulations of surgical procedures. Virtual environments for exploration of three-dimensional environments.
Biomedical research	Visual feedback from remote instrumentation. Online conferences. Collaboration among distant researchers.	Transferring large data sets between computers for high-speed computation and comparisons. Reviewing results of remote experiments. Searching archives of three-dimensional medical images.	Controlling experimental equipment, such as electron microscopes.
Limiting Technical Factors	Availability of sustained, predictable, high-bandwidth connections to many locations, including rural health clinics and patients' homes (to support remote consultations).	Authentication of source and recipient of information. Security of personally identifiable information in transit across the network and in storage at either end of the network. Availability of sustained high-bandwidth connections for transfer of large, time-critical files.	Network latency and bandwidth. Ability to obtain guaranteed bandwidth for predictable periods of time.

Information Search and Retrieval	Real-Time Collaboration	Primary Technical Challenges
Accessing reference materials and course materials.	Virtual classrooms. Distributed collaborative projects. Distributed discussions.	Sufficient bandwidth to accommodate large numbers of transactions from a single educational institution or to support access to remote scientific and clinical simulations. Ubiquity of access for students in remote clinical rotations and to support educational applications in the home.
Searching remote databases and professional literature.	Collaboration among researchers. Peer review. Interactive virtual conferences.	Sufficient bandwidth to support rapid transfers of large sets of data for distributed simulations. Low latency to accommodate remote control of equipment.
Tools for locating information of interest and for determining the quality of retrieved information. Means of allowing anonymous searches.	Sustained access to high-bandwidth, low-latency networks for collaborations involving real-time video or manipulation of images. Multicast protocols to make more efficient use of networking resources.	

All applications that involve the transmission of personal health information (such as data contained in electronic medical records, claims for payment, prescriptions, or public health reports from testing laboratories) demand that information be kept confidential. Furthermore, virtually all applications in consumer health, clinical care, health care financial and administrative transactions, public health, and biomedical research require that the integrity of information be assured and maintained both during and after transmission. Meeting these requirements demands a variety of technical supports (as well as policies governing the disclosure of information), including suitable encryption of information during transit and rigorous authentication of both the source and the recipient of information. Access controls are also required to ensure that users can view only the information they are authorized to see; auditing technologies are needed to ensure that successful attempts to circumvent access restrictions are identified so that violators can be punished.

The need for data protection and access control is acute in health applications because some personal health information is extremely sensitive. Not only can loss of confidentiality cause embarrassment and social stigmatization, but personal health information can affect an individual's employment and insurance coverage, especially for people with private insurance or who work for self-insured organizations. Moreover, once health information is divulged, its confidentiality cannot be regained; there is clearly a difference between the prospect of losing \$50 when one's credit card number is stolen and losing privacy when one's HIV status is revealed to friends and co-workers. Confidentiality problems are compounded in health care because many people have a legitimate need to see sensitive patient information. These include workers involved, for example, in the provision of care, payment for services, and filling of prescriptions. In addition, legitimate access may be needed by someone with whom the patient has had no previous relationship—perhaps a physician at an institution that the patient has not visited before (e.g., in an emergency room situation).

Although technologies have been widely deployed for encrypting information transmitted across the network (e.g., Secure Socket Layer encryption), technologies for authenticating the identity of users at both ends of a transaction are not in widespread use, especially among consumers. This approach is effective in electronic commerce applications because most vendors can obtain certificates to authenticate themselves to consumer Web browsers, but despite some early efforts no effective mechanism yet exists for providing authentication devices in large numbers to consumers, including patients. This is not an impediment to electronic commerce because most merchants are willing to authorize a transaction once a valid credit number is presented, but an artifact such as

a credit card may not be suitable to allow access to an online health record. In telemedicine applications or the retrieval of medical records, a compelling need exists to identify the end user reliably. There are comparable circumstances in which an authenticated third party needs to gain access to information via the Internet, such as an emergency room physician accessing a patient's health record that is stored at another institution connected to the Internet.

In addition to security, high levels of network availability are needed to support many health-related Internet applications, particularly in clinical care, in finance and administration, and in public health. Health care organizations must be assured that the network will be available almost around-the-clock if they become dependent on using the Internet for accessing electronic medical records, for remote monitoring of patients, or for clinical decision support. Payers and administrators will also demand high levels of availability if they are to use the Internet instead of private networks for important transactions. The network must be made robust against failure and against hostile attacks, whether directed at its physical infrastructure or at denying its services to end users by flooding its capacity (denial-of-service attacks).

The need for security and availability is compounded by the fact that in many clinical applications of the Internet, human life and health may be at risk. Errant decisions based on incorrect information—whether in diagnosing a condition or filling a prescription—can be harmful or fatal. Hence, ensuring data integrity and properly authenticating individuals are even more important than in many other spheres of application. Inability to access patient information (such as from an electronic medical record), to complete a distant consultation, or to control remote monitoring and dosimetry equipment can also undermine the quality of care and, thus, the health of the populace. Internet-based applications, including the physical networks on which they run, must be robust in the face of failures. A broken fiber-optic cable should not prevent an application from running when an alternative path exists. Remotely controlled instruments should not perform any damaging action if they lose contact with the controller as the result of a network interruption. Guarantees of network performance must be extremely robust in order to prevent statistically unlikely events from having serious consequences.

Conclusion 3. The quality of service needed by a number of high-end health applications will not necessarily be deployed soon across the Internet in a form that meets the needs of the health industry.

A number of potential health applications of the Internet demand guarantees on the quality of service they get across the Internet. The need for QOS derives from the frequent need for smooth and responsive interactivity. For example, care providers engaging in remote video consultations with patients or other care providers need sustained access to high-bandwidth connectivity (roughly 384 kbps for simple interactions and 768 kbps for higher quality video) for the duration of the consultation. So do molecular biologists who wish to control visually an electron microscope located at a remote facility, or medical students who wish to practice a surgical technique using multimedia simulations that are available on remote servers, or groups of surgeons from different parts of the country who wish to collaborate in planning a difficult procedure. QOS is also needed for making practical the real-time exchange of large images, whether medical images such as X rays and mammograms or anatomical images for educational purposes. QOS would not only need to ensure that adequate bandwidth is available to provide timely delivery of images but would also need to offer real-time interaction to allow a primary care provider and a consulting specialist to point out specific items of interest in the image.

Whether the Internet will provide the needed capabilities in the near future is uncertain. The protocols currently deployed across the Internet for routing packets do not contain mechanisms to support guaranteed QOS; rather they provide best-effort service, in which packets are delivered as best the network's resources and traffic levels will allow. Internet service providers (ISPs) are attempting to improve service quality across their networks by deploying additional bandwidth, but this approach does not allow explicit guarantees to be made on bandwidth, latency, and jitter. Protocols have been developed to support different forms of QOS across the Internet (e.g., the differentiated service and integrated service models described in Chapter 3), but they have not yet been deployed, and even if deployed, they may not fully support health applications. For example, the differentiated services (diff-serv) standard does not include mechanisms for providing QOS guarantees for packets that must traverse the networks of different ISPs. Hence, individual ISPs may be able to offer improved QOS to customers attached to their networks, but they cannot provide guarantees related to traffic flows among organizations connected to different ISPs. Because the health industry is highly decentralized and individual care providers' offices may need to interact with a number of different managed care organizations, insurers, and other care providers, inter-ISP mechanisms will be important for health applications. The challenge of providing QOS in a health environment is further complicated by the extremely variable QOS needs of individual health organizations over time. The kinds of information exchanges in which an

organization engages typically vary considerably in the course of a day, from simple exchanges of information regarding a patient's coverage by a health plan, through transfers of medical records with affiliated organizations, to the exchange of large medical images for interpretation and diagnosis. The bandwidth needs of a small medical clinic could, accordingly, vary enormously during the course of a day, ranging from near nothing one minute to several megabits per second the next. Finding ways to satisfy such variable demand for bandwidth economically represents a significant challenge. The integrated services (int-serv) approach to QOS can support variable bandwidth needs through protocols for reserving capacity, but such protocols may not be sufficiently scalable to support widespread deployment across the Internet, as health applications could demand. Nor can it be assumed that QOS mechanisms that are optimized for content distribution (i.e., information flows that are predominantly one-way) will be effective for the more symmetric transactions typical of many health care applications.

Conclusion 4. Ensuring widespread access to the Internet is essential to achieving its promise in health applications.

One of the most dramatic effects of the Internet is its ability to engage patients and consumers more actively in health issues. As the amount of health-related information on the Web increases, patients become more actively involved in maintaining their health: seeking information related to specific ailments or topics of interest, discussing medical problems with peers in online chat groups, e-mailing care providers with questions regarding symptoms or treatments, assessing their health, scheduling appointments with care providers, and maintaining their own health records online. These actions are reinforced by a number of fundamental factors in the nation's health care system, including greater consumer choice in selecting among alternative health plans, concerns about the quality of care provided by health care organizations in an increasingly competitive environment, and pressures to shift the site of care away from the provider location and to the consumer's location and to shift from unilateral to shared approaches that bring patients into the decision-making process. All of these factors encourage—and in fact require—consumers to become more educated about their health and health care.

As these trends continue, the Internet will probably play a more central role in supporting the processes of health care, reinforcing the calls for expanding consumer access to Internet resources. Persistent inequalities of access to the Internet could exacerbate existing inequalities of access to health care. Strong social pressures exist to ensure some degree of equity in access to health resources. Recent statistics show that those

most in need of health care and who could benefit most from Internet-mediated care—such as those in rural and inner city areas—are also those with the least access to the Internet. Furthermore, many of the local strategies used to expand access, such as the provision of Internet-enabled computers in libraries, classrooms, and community centers, may not be as effective in addressing consumer health needs. Access may be needed outside normal business hours (especially if the Internet becomes a medium for providing care or real-time advice on medical emergencies), and many users might feel uncomfortable researching or discussing personal health problems in a public venue. They will increasingly want such access in their homes and may come to view it as a necessary component of a comprehensive health care system.

Access concerns will not be limited to consumers, however. Care providers, too, will need simplified, high-speed access to the Internet for a range of information services and to provide clinical services. In rural areas, in particular, it is not always easy for local care providers to access the Internet, and they often need more advanced networking services than consumers do. As health education shifts its focus from hospitals and medical centers to local offices and remote-practice clinics (both for training students and residents and for continuing medical education), ensuring broad access to Internet resources will become increasingly important. The same Internet that delivers health care to the home and curriculum materials to students in the health professions (e.g., doctors, nurses, and pharmacists) can also deliver information resources for patient and consumer education. The issue of access becomes even more complex for health applications that demand high bandwidth, because broadband technologies are still not widely deployed—especially in places such as rural areas that could most benefit from remote consultations. Access concerns also have implications for quality of service. High QOS cannot be provided cost-effectively to all Internet traffic, and many applications do not require it. Cost is one mechanism that is generally considered as helping to control demand for high QOS capabilities, meaning that QOS could be limited to those who are willing (or can afford) to pay for it.

Conclusion 5. Technical advances are needed across many areas of information technology (not just networking) if the potential of the Internet is to be achieved in support of health applications.

Networking capabilities are not the only technical impediments to many health applications of the Internet. For example, there is great concern regarding the quality of health information on the Internet and the inability of many consumers to assess adequately the credibility of information provided on different sites. While a number of nontechnical

approaches can be taken to rate the quality and provenance of information and to determine whether it comes from an authoritative source, there is considerable room for technological solutions as well. Other health applications demand different technological advances. Remote medical consultations (to rural medical clinics or the home, for example) will demand medical instruments (e.g., stethoscopes, blood pressure monitors, and respiration monitors) that are suitable for personal use and that can interface with home computing devices.

Conclusion 6. Health care organizations are ill-prepared to adopt Internet-based technologies and applications effectively.

A number of transformations in the health care industry are driving the use of the Internet as a medium for sharing information among providers, patients, and administrators. New Internet-based health companies are being established to offer consumers medical information, tools to help them to monitor their care more effectively, and other medical products and services. Existing organizations are using the Internet to alter their position, relationships, role, and power in the health care industry by moving into new areas, often ones that involve their reaching out directly to patients. The more innovative care organizations are providing customers with Internet-based resources to help them to better assess their medical needs and to seek appropriate advice. These trends reflect and reinforce ongoing attempts to reduce medical costs by reducing hospital stays and outpatient care, especially in the emergency room, and by promoting prevention through better monitoring of wellness, chronic disease, and behavior.

Nevertheless, the health care industry as a whole is ill-prepared to accommodate this change. Despite continuing consolidation among providers, insurers, and managed-care companies, health care is still largely a decentralized industry populated by diverse organizations with different motives, resources, and incentives—it is sometimes referred to as a “trillion dollar cottage industry.” It comprises thousands of hospitals and hundreds of thousands of physician’s offices and includes academic medical centers, community hospitals, large physician groups, solo physician practices, home health agencies, health centers, and rehabilitation hospitals. Other participants are pharmaceutical companies, managed care organizations, and traditional indemnity insurance companies. This diversity brings with it different degrees of sophistication and a diverse set of challenges, resources, and missions but makes it difficult to speak with a unified voice or to adopt a critical mass of technology. Recent consolidation of the industry has somewhat improved the industry’s ability to achieve critical mass more quickly, but much of the consolidation (e.g.,

the formation of integrated delivery systems that link different types of care provider organizations under a single organizational umbrella) tends not to remove the diversity as much as mask it behind single organizational identities.

Current fiscal constraints—on care provider organizations, in particular—further hinder the industry's ability to make major investments in information infrastructure and applications unless these investments can be shown to lead to significant and low-risk returns. Information technology will be adopted rapidly if it results in a material and obvious advance in medical practice (such as magnetic resonance imaging machines), but adoption is more difficult for technologies with less quantifiable benefits, such as the security technologies that control access to medical records. The Internet itself is such a new phenomenon that its eventual contribution to the delivery of care is poorly understood by the industry as a whole. It is unrealistic to expect that the industry will rapidly overcome its fragmentation and diversity to speak loudly and with one voice. The traditional conservatism of the health care industry in the face of information technology is likely to persist in the face of the Internet.

In addition to these larger structural problems, impediments exist within individual health organizations. Provider organizations lack information about the potential benefits and effectiveness of Internet-mediated health interventions in terms of both cost and quality, making it difficult for them to make decisions about Internet investments. Industry reference models have yet to be developed, and not much information has been developed or shared about the kinds of Internet-based systems that are demonstrably effective in improving care. Organizations also lack authoritative guidelines to help them develop formal policies related to Internet applications, such as rules for e-mail exchanges between patients and providers and for monitoring sponsored discussion groups. Reflective of a tight labor market overall, health care organizations also often have difficulty attracting qualified engineers and programmers to develop information technology systems for use in health-care settings.

Conclusion 7. A number of difficult public policy and regulatory issues constrain the adoption of Internet-based health applications by health organizations and consumers. Some of these issues are specific to the health sector; many others extend beyond the health sector but require the health community's active participation in their resolution.

Many of the same issues that have slowed the growth of traditional forms of telemedicine (using private networks) will continue to impede

the expansion of Internet-based medicine. For example, health care providers are currently licensed by individual states and are generally prohibited from providing care across state lines—a clear issue when a patient is in one state but the physician at the other end of a telemedicine link is in another. Liability claims are also handled at the state level, with considerable variation among states. Such policies have made it difficult, if not impossible, for physicians to practice telemedicine across state boundaries. So has the lack of means for paying for a full range of remote medical consultations in the Medicare program and in many traditional insurance companies.¹ These policy issues limit the returns that health care organizations can expect to receive from investments in network-based health care solutions. They also affect the willingness of physicians to adopt new technologies and practices.

Other policy concerns also threaten to slow the spread of health-related applications of the Internet. Concerns over the security of personal health information continue to dampen consumers' enthusiasm for Internet systems that share such information. The Department of Health and Human Services has promulgated draft regulations governing privacy and security of electronic health records, but these provisions have yet to be finalized and will not cover all exchanges of information among the many kinds of organizations that collect, process, and distribute health information. Beyond the issue of medical records privacy lies the larger issue of the privacy of online searches and transactions that has arisen in many forms of electronic commerce and Web browsing. What kinds of information can be collected about individuals online, and how can that information be shared? Recent accusations that several health-related Web sites have not adequately disclosed their information-gathering and data-sharing practices have brought renewed attention to this issue (Clausing, 2000). Other issues associated with the protection of intellectual property on the Internet also have repercussions for health care, especially in regard to the dissemination of professional literature and educational courseware designed for use over the Internet. Efforts are needed on the part of governmental, industrial, and provider groups to craft an effective solution.

RECOMMENDATIONS

Enhancing the Internet to support the health community more effectively will require active stewardship on behalf of health organizations, the information technology industry, the research community, and government. These groups must attempt to ensure that the Internet evolves in ways that support health care and that the health sector is prepared to incorporate the Internet into its processes for delivering care, paying for

care, conducting research in biomedicine and health services, improving public health, and providing health education. Action is needed in four areas, as outlined below: (1) ensuring that suitable technical capabilities are developed and deployed in the Internet, (2) demonstrating and evaluating Internet-based applications in the health and biomedical sectors, (3) educating the health community on ways to incorporate the Internet safely into their routine activities, and (4) resolving policy issues that impede use of the Internet in health.

Research, Development, and Deployment of Needed Technical Capabilities

To ensure that the Internet has the ability to support health applications, additional technical capabilities need to be developed and deployed. Although some existing technologies could be deployed in the Next Generation Internet (NGI) to facilitate experimentation with and evaluation of health and biomedical applications, other technologies will require further research. The health community must articulate its needs to the information technology research community and must actively engage in developing Internet-based systems. The goal is to assure that developers better understand the ways in which the requirements needed for health care applications of the Internet diverge from, or converge with, those needed to support Internet applications in other sectors. To this end, the committee makes four recommendations.

Recommendation 1.1. The health community should ensure that technical capabilities suitable for health and biomedical applications are incorporated into the testbed networks being deployed under the Next Generation Internet (NGI) initiative and eventually into the Internet.

The NGI testbed networks being deployed by federal agencies such as the Defense Advanced Research Projects Agency, the National Science Foundation, the National Aeronautics and Space Administration, and the Department of Energy will provide a basic communications infrastructure on which a number of high-end health applications could be demonstrated and evaluated. To help the health community experiment with new networked applications, these testbed networks must be deployed with the technical capabilities to support a range of high-end health applications, such as medical consultations at a distance, remote control of research equipment, surgical simulation, and collaboration among researchers and clinicians. The networks are expected to be deployed with bandwidth that can support a wide range of health applications,

although other technical capabilities will need to be deployed as the technologies become sufficiently stable:

- *Quality of service mechanisms.* QOS mechanisms should be deployed in the NGI to support applications that require, for example, the rapid transfer of digitized medical images (or medical records), the use of real-time video for remote medical consultations, and the control and receipt of imagery from remote experimental equipment, such as electron microscopes. Remote videoconferencing, for example, may require less than 1 Mbps of bandwidth to be effective, but that bandwidth must be predictably and reliably available for the duration of a clinical consultation—perhaps 30 minutes or more. In addition, low latencies will be needed to support more natural, real-time interaction among participants. The differentiated services (diff-serv) and integrated services (int-serv) models for QOS can each support a range of anticipated health applications. Both should be deployed across the NGI to allow further experimentation with their capabilities and their limitations. Deployment of the diff-serv model would be a valuable first step because it would allow users to experiment with applications that demand quantifiable QOS guarantees and to develop policies for determining which types of data traffic need a higher level of service and which can be satisfied with lower levels of service. Deployment of the int-serve model could expand the number of end-user organizations that take advantage of QOS capabilities because it would allow them to reserve premium services for use when needed, rather than having to subscribe to a premium service that they may leave unused much of the time. Deployment of int-serv across the NGI testbed networks would also allow researchers to better investigate the ability of the protocols to accommodate large numbers of users in an operational environment; issues of scalability have yet to be resolved with int-serv. Moreover, deployment of diff-serv and int-serv protocols in the NGI testbeds would enable the health community and ISPs to investigate effective business for financing the wide-scale deployment of QOS mechanisms throughout the Internet. Considerable work must be done to find a way of creating an incentive for ISPs to deploy QOS capabilities in a way that is consistent with the health community's needs.

- *IP Security.* Deployment of IP Security (IPSec) protocols would allow further experimentation with virtual private networks among health care organizations involved in the NGI program and would allow further testing of the ability of the protocol to scale sufficiently to support effective interchanges of information among health organizations. The use of IPSec across the Internet could allow secure communications between fixed sets of health organizations and permit the transfer of health records among affiliated care providers or the exchange of payment, enrollment,

and other administrative information between care provider organizations and third-party payers and administrators. While Secure Socket Layer (SSL) protocols can also support confidential transactions across the Internet, the use of IPSec could simplify security administration by providing a central point at which all traffic can be encrypted and authenticated as it leaves one site destined for another. It could further mitigate concerns regarding the authentication of individual users, as discussed below.

- *Public key infrastructure.* The use of SSL encryption across the Internet already enables a wide range of health applications, including remote access to medical records. What is now missing from such systems is typically a secure way of authenticating end users. While most online vendors use cryptographic certificates to validate their identities, most implementations of SSL rely on simple password schemes to validate the identity of end users—or they simply require a valid credit card number to be presented for billing purposes. The NGI program offers an ideal environment for experimenting on a limited basis with stronger forms of authentication, including the use of public key infrastructures to distribute and validate cryptographic certificates for individual end users. Effective and scalable PKI strategies could be developed to allow more routine sharing of information among more loosely affiliated participants in the health community.

Recommendation 1.2. To ensure that the Internet evolves in ways supportive of health needs over the long term, the health community should work with the networking community to develop improved network technologies that are of particular importance to health applications of the Internet.

It is difficult to predict which particular technologies will be needed, but based on the evaluation presented in this report, several areas seem reasonable as starting points for continued research:

- *More readily scalable techniques to provide bandwidth guarantees on demand.* Techniques are needed for ensuring that individual organizations (or individual end users) can receive high-level QOS across the Internet on an as-needed basis. Smaller health care organizations and individual users in particular are unlikely to be able to afford to subscribe to statically provisioned premium-level QOS offerings, but they may need on occasion to send large files quickly across the Internet or to engage in remote video consultations. The diff-serv model that has been standardized by the Internet Engineering Task Force (IETF) does not yet incorporate mechanisms for providing premium services on demand. Nor have

standards been defined for the provision of end-to-end QOS across multiple service providers using diff-serv. The int-serv model, which does provide such reservation mechanisms, may not, however, be sufficiently scalable for widespread utilization across the Internet backbones. There are also concerns that the level of QOS assurance provided by the diff-serv model are not strong enough for mission-critical applications. Research programs are under way to develop scalable mechanisms for providing QOS on demand across the Internet and must continue to be supported.

- *Stronger forms of authentication.* Additional research is also needed to develop methods for authenticating users of networked applications to ease the problem of reliably identifying individual consumers and allowing secure communications between parties that have not established relationships beforehand. Technologies and processes must be developed for key and certificate management systems that can be used by individuals and small organizations to authenticate their identities in online transactions. Solutions must be capable of scaling up to the size of the entire population and must address the issues of delegation of trust that certification hierarchies face. Work on smart cards, token-based authentication, and biometric authentication devices would seem to be especially useful in serving health needs.

- *Symmetric or dynamically reconfigurable broadband technologies for the last mile.* Techniques are needed to enable consumers to transmit information at data rates similar to those at which they can receive it. Current technologies for high-speed Internet access tend to be configured to allow much greater bandwidth into the home than out of it. For many health applications, however, it is conceivable that consumers will send as much information into the network as they retrieve from it (home-based remote medical consultation is an example). This capability might be achieved by deploying more symmetric forms of broadband access (e.g., symmetric DSL or cable modems that allocate more bandwidth to upstream capacity) or by developing ways to allow end users to use some downstream bandwidth for upstream communications when needed.

- *Hardened QOS guarantees.* Mechanisms are needed to ensure that critical health care applications do not lose their QOS guarantees as a result of link failures across the Internet, except in those cases in which the network suffers catastrophic failures that leave it without connectivity between desired points of communication. Part of this effort will require work on techniques for rapid reconvergence after link failures. Approaches are needed to ensure that Internet routers can determine updated communications paths across the network should a particular link fail and to ensure that data are not lost during the reconvergence period. Applications such as remote control of experimental or clinical equipment and

patient monitoring may not be able to tolerate the brief loss in connectivity (and perhaps loss of packets) that characterize reconvergence efforts on today's Internet. In particular, efforts are needed to address disruptions in service between two ISPs, where existing link recovery schemes may not be effective.

- *Disaster operations.* Techniques are needed to ensure delivery of high-priority traffic (including, but not limited to, health information) in the event of a natural or man-made disaster that disrupts large segments of the nation's information infrastructure.

Recommendation 1.3. The National Library of Medicine should forge stronger links between the health and networking research communities to ensure that the needs of the health community are better understood and addressed in network research, development, and deployment.

To date, interactions among members of the networking community and the health informatics community have been minimal. The health sector is generally underrepresented in standards-making bodies such as the IETF, and few research projects attempt to explicitly forge alliances between networking researchers and health researchers. Bridging this gap would help networking researchers and standards developers to better understand the challenges that health applications pose for networking infrastructure and would help the health informatics community to learn how to incorporate new networking technologies more effectively into their applications. One way of encouraging more cross-fertilization between these communities would be for the National Library of Medicine to encourage or require recipients of its contracts and grants for networking-related projects to participate in meetings and conferences of networking researchers and standards-making bodies. It could explicitly indicate that it would support travel and related costs for such endeavors. The NLM could also explicitly require collaboration among networking researchers and health researchers in some of its awards. It could also itself play a more aggressive role in establishing contact between the health community and the networking community. It could speak for the health community in conveying needs to networking researchers and in reaching out to ISPs, presenting health care as a cutting-edge example of a peer-to-peer application that demands networking capabilities beyond those now being deployed (e.g., QOS throughout the network and symmetrical forms of broadband access for residential use). The NLM could also establish an ad hoc task force to explore other ways of communicating health-related needs to the networking community.

Recommendation 1.4. The National Institutes of Health and its component agencies should fund information technology research that will develop the complementary technologies that are needed if the health community is to take advantage of the improved networking technologies that can be expected in the future.

As noted above, putting adequate network infrastructure in place is not sufficient for enabling Internet applications in the health domain; other information technologies are needed to enable network-based applications to develop. The medical service arena is a problem-rich one with considerable need for research into new technologies and new network capabilities. The Department of Health and Human Services (DHHS), by virtue of its broad purview over health-related efforts in the federal government, should play a leading role in funding research on information technologies of particular importance to the health community. At present, DHHS provides only a small portion of the federal government's funding for computing research (the vast majority of funding comes from the Defense Advanced Research Projects Agency and the National Science Foundation, with other significant contributions by the National Aeronautics and Space Administration and the Department of Energy). Two recent reports encouraged DHHS (or elements within it, in particular NIH, AHRQ,² and NLM) to become more actively engaged in information technology research. A report by the Working Group on Biomedical Computing (1999) of the National Institutes of Health, for example, calls for increased basic research on information technology tools to support biomedical computing and for the establishment of national programs of excellence to advance research in areas of biomedicine where computation is increasingly important. Similarly, the President's Information Technology Advisory Committee concluded that "the National Institutes of Health (NIH) should support biomedically motivated basic research in information technology and view it both as important information technology research and as fundamental biomedical research" (PITAC, 1999). It is not possible to identify in advance all areas in which the health community needs to become engaged, but the analysis in this report identifies the following topics as pertinent preliminary areas of inquiry:

- *Validation of online information.* Effective techniques are needed to help consumers to judge the quality of the health information they retrieve from the Internet and to give them more confidence in it. The goal of these techniques should not be to prevent dissemination of certain information via the Internet but to ensure that consumers have a way of judging the quality of what they find there. Different technical approaches

may be effective in automating what is now a more or less manual process of labeling content to indicate its provenance.

- *Tools for protecting anonymity online.* Technologies are needed to improve privacy on the Internet, so that users may pose queries, receive replies anonymously, and browse Web sites without necessarily divulging their identities. Use of the Internet to retrieve health information or to purchase products can reveal much about a person's health status and concerns. A provider's use of online clinical resources could also provide clues about his or her knowledge of a given subject area. Such information could be used in ways that harm Internet users. The health community therefore has a strong interest in developing tools to protect the identities of online users.

- *Access controls.* Mechanisms are needed to limit users' access to just the programs, databases, and data fields they need to perform their job function or play their role in an online transaction. Development of robust access controls is especially difficult in health care because there are many kinds of actors (care provider, educator, researcher, etc.), each of whom requires access to information, and many kinds of information, most of which should be privileged. These problems are further complicated by the dynamically changing nature of the relationships in health care, e.g., between a patient just admitted to an emergency room and the health care providers at the admitting hospital.

- *Controls on secondary distribution of patient health information.* Research is needed to develop effective mechanisms for managing the distribution of health information beyond the walls of the organization that originally collected it. A number of technologies, such as digital property-rights-management languages and encrypted containers, have been developed to control the distribution of digital media and to ensure the adherence of recipients to particular provisions or use of the information (e.g., payment for access, limits on redistribution). Additional work is needed, however, to extend such capabilities into the health environment, where the new technologies may be able to control the distribution and use of health records and other confidential information.

- *Improved audit capabilities.* Additional work is needed to develop tools for reviewing audit logs that health care organizations could compile on accesses to electronic medical records. These tools would need to automatically identify potential violations of confidentiality, drawing from external databases such as scheduling calendars and referrals to differentiate between legitimate and illegitimate accesses.

- *QOS policies that are suitable to health and health care.* Policies are needed to govern the provision of QOS in high-traffic environments to ensure that critical health-related messages receive their requested QOS in emergency situations. Mechanisms for defining and propagating QOS

policy across a network are currently being defined, and it is important that such mechanisms are developed in a way that meets the needs of the health community. Work is also needed to establish guidelines for determining which packets sent by particular institutions will receive the highest priority.

- *QOS-aware applications.* Work is needed on networked applications that are tightly linked to the underlying QOS capabilities of the network. Only by creating networked applications that leverage emerging QOS capabilities can the shortcomings of those capabilities for health applications be determined and requirements for new QOS capabilities be established. Furthermore, by demonstrating the utility of such capabilities in a health care context, it may be possible to influence the deployment of new QOS capabilities by ISPs in a direction that is useful to health care.

Demonstration and Evaluation of Health Applications of the Internet

Continued experimentation and evaluation is a central component of efforts to understand better the kinds of health applications that may become more widespread across the Internet and the technical capabilities they demand. By demonstrating health applications such as distant consultation, remote control of experimental equipment, and online access to electronic medical records, members of the health community will have a better opportunity to examine their relative costs and benefits, the business models needed to support them, and the kinds of organizational policies that are needed to govern their use. A number of public and private organizations have supported demonstration programs to allow such exploration. These efforts need to continue as new Internet technologies become available and new applications are envisioned. They will be increasingly important to the health community if it is to establish a dialog with the larger Internet community about its evolving needs. The efforts will allow continued identification of technical requirements that the networking community can address and of other problems and issues for the health community. To provide the kinds of information that will inform this dialog, a number of parallel efforts will be needed, as recommended below.

Recommendation 2.1. The Department of Health and Human Services should fund pilot projects and larger demonstration programs to develop and demonstrate interoperable, scalable Internet applications for linking many health organizations.

Public and private organizations have supported a range of pilot projects and testbeds to demonstrate the capabilities of information tech-

nology in health organizations. Most have tended to focus on stand-alone applications that operate within a single organization. Many of the challenges of Internet-based systems derive from the interconnection of many organizations to a large network. DHHS and its constituent agencies (NLM in particular) should fund pilot projects and testbeds that explicitly connect multiple organizations for purposes of information exchange. These projects are not intended to develop community-wide repositories of health information, as was attempted with community health information networks (CHINs), but should try to facilitate information exchange among limited sets of organizations, whether for clinical care (e.g., exchanges of medical records, the sharing of clinical guidelines, or remote consultations between an urban medical center and several remote clinics) or administration (e.g., payment of claims). Only by establishing testbeds on a large scale can the issues of scalability and the effects of decentralization be identified and evaluated.

These projects should explore the capabilities, limitations, and performance requirements for health applications in a highly networked environment with many participating users at different organizations. For example, projects could be supported in application domains such as (1) public health surveillance, (2) clinical care, (3) home-based care, (4) remote consultations, and (5) payment for services. Pilot projects that integrate entire vertical slices of the health care delivery process could also be tried. Such projects would help private and public organizations experiment with possible applications of the Internet and determine the ways the Internet can be used most effectively. By involving a large number of organizations, the projects will also aid in understanding issues associated with access to information resources, especially if the projects involve outreach to individual consumers. They would look at whether the research efforts outlined in Recommendation 1.1 had achieved the infrastructure required for health networking applications. Thus, demonstrations that show the effective use of advances in bandwidth, latency, QOS, and reliability would be most appealing.

Recommendation 2.2. Federal agencies such as the Department of Veterans Affairs, the Department of Defense, the Health Care Financing Administration, the National Institutes of Health, and the Indian Health Service should serve as role models and testbeds for the health industry by deploying Internet-based applications for their own purposes.

The information technology and health care industries need to do considerable experimentation, standards development, and integration work before the Internet can be used routinely for health care purposes.

The government could play a proactive, catalytic role in this effort. It operates a substantial health care operation of its own, in the form of the health care systems operated by the Veterans Health Administration (part of the Department of Veterans Affairs, or VA), the Department of Defense, and the Indian Health Service. The Health Care Financing Administration operates an enormous system to pay for care provided under the Medicare, Medicaid, and related programs. By partnering with industry, these government agencies could use the Internet aggressively in their health-related systems. HCFA could, for example, develop policies and standards for the electronic submission of Medicare claims and could act as the certificate authority for a public key infrastructure for authenticating people and organizations that submit claims. The DOD and the VA could use their telemedicine programs and efforts to exchange medical records as testbeds and demonstrations that could influence private sector initiatives. In effect, these would be pilot implementations and they would help industry develop appropriate standards and software. An open-source model could be required, thus making the fruits of this effort available broadly to industry. This is an area where governmental leadership is likely to be crucial and where passivity on the part of the government will cause many lost opportunities.

Recommendation 2.3. Health organizations in industry and academia should continue to work with the Department of Health and Human Services to evaluate various health applications of the Internet in order to improve understanding of their effects, the business models that might support them, and impediments to their expansion.

Health organizations will adopt health care applications of the Internet largely on the basis of their ability to improve the quality of care and reduce its costs. Deployment of needed infrastructure will be motivated by the development of business models for supporting the applications and paying for network services such as QOS. To date, the health industry has had little guidance in these areas, and its adoption of the Internet will be slow without better information. Work is therefore needed to determine the effect of Internet applications on care quality and costs; organizational performance; job skills; relationships between participants (e.g., changes in the role of patients and responsibilities of patients); the economic or business models of care (e.g., use of the Internet to establish contractual relationships between affiliated care providers rather than buying them outright); provider workflows; and confidentiality and liability concerns. Across all of these areas, evaluations should seek to understand whether Internet-based applications have a different impact on

quality, cost, and access to care than non-Internet based applications (e.g., are Internet-based implementations of electronic health records materially different from non-Internet implementations?). Similarly, studies should identify factors that hinder the expansion of health applications of the Internet from prototypes or demonstrations to broad organizational or national use. Additional work is needed to develop mechanisms whereby the health industry can afford the investments in information infrastructure needed to enable more sophisticated applications. Elements within the Department of Health and Human Services can continue to play a vital role in providing financial support for these evaluations, as has been done by the National Library of Medicine and Office of Rural Health Policy.

Recommendation 2.4. Public and private health organizations should experiment with networks based on Internet protocols and should incorporate the Internet into their future plans for new networked applications and into their overall strategic planning.

Even though several years may pass before the Internet can provide the QOS, security, and reliability needed for health-related transactions, health care organizations need to start preparing now so that they will be equipped to use the improved capabilities offered by future generations of the Internet. They must develop an understanding of the ways the Internet can support their missions, prepare their infrastructures to be compatible with Internet technologies, develop the human resources needed to design, develop, and deploy effective systems, and put policies in place to govern the use of the Internet and Internet-related applications. Individual organizations and professional societies have roles to play in this endeavor.

Because the health care industry has limited experience with Internet-based applications and because models for delivering and paying for health care continue to evolve, it is not yet clear how the Internet can best be used to improve health and minimize costs. In the end, some balance between Internet-based and private networks will probably emerge to meet the full spectrum of capabilities needed by health care organizations. In the meantime, these organizations should take steps to understand better the benefits of Internet-based systems and to have Internet-ready resources in place. They should establish institutional connectivity to the Internet and among their constituent organizations. They should establish network-based relationships with vendors that allow them to explore electronic commerce opportunities. They should begin using networks that incorporate Internet protocols such as TCP/IP, e-mail, FTP,

and HTTP for internal communications. And they should begin thinking about ways in which the Internet can enhance and extend their missions. These represent relatively inexpensive steps for gaining critical insight into the ways the Internet may become more fully engrained in health processes at a later date. Without such experimentation, health organizations risk missing out on future opportunities.

Addressing Educational Needs

In order for the Internet to achieve its full potential in health applications, not only must it provide adequate technical capabilities, but health organizations also must be capable of adopting it. Health organizations will need the internal capability to envision ways in which the Internet could support their missions and to design, develop, and implement systems that fulfill those visions. The experimentation outlined in Recommendations 2.1 through 2.4 will improve organizations' capabilities considerably, but additional efforts will likely be needed to bolster internal policy development and human resource development. Efforts in three areas are recommended.

Recommendation 3.1. Professional associations with expertise in health issues and information technology should work with health care organizations to develop and promulgate guidelines for safe, effective use of the Internet in clinical settings.

Improved information outlining best practices for using the Internet in health applications would help individual organizations benefit from each other's experiences and develop informed policies and procedures to guide their own efforts to harness the capabilities of the Internet. Professional associations have an important role to play because their membership spans large numbers of organizations that face common challenges. Professional associations with expertise in health care and information technology could help convene groups that would develop standards and guidelines based on the experience and expertise of their memberships. Some associations have already taken productive steps in this direction. The American Medical Informatics Association, for example, developed a set of guidelines for using e-mail in clinical settings, and the Association of American Medical Colleges has initiated programs to evaluate the information technology needs of academic medical centers that could also produce valuable guidelines. Similar efforts are needed to develop guidance in (1) monitoring and conducting health-related chat sessions, bulletin boards, and forums, (2) remote education of health professionals, (3) disseminating information to a broad audience, (4) appropriate and

inappropriate creation of provider/patient relationships, (5) assurance of the integrity and accuracy of patient-maintained health records, (6) means to assess trade-offs between security, confidentiality, and access, (7) direct marketing to patients of health care services (e.g., pharmaceuticals, prostheses), (8) communication across traditional boundaries (e.g., patient to provider), (9) clinical e-mail, (10) Web information services, and (11) privacy and security of electronic health information.

Recommendation 3.2. Government, industry, and academia should work together and with professional associations having experience in health and information technology to educate the broader health and health care communities about the ways the Internet can benefit them.

Part of the inability of health organizations to aggressively pursue Internet strategies derives from a lack of appreciation among health workers about the potential benefits of Internet-based applications. Many have had limited formal education in computing and communications technology and continue to have limited experience using it. Educational programs could go a long way to overcoming institutional resistance and helping workers to better use such systems. Academic health organizations and professional associations have important roles to play in educating the health community at large about the potential benefits of Internet-based systems in health care. Academic health organizations are among the leaders in applying the Internet to health applications and educate a range of health professionals. Associations can draw upon their large and diverse membership to pool ideas and reach out to individual organizations. They must also work individually to assemble information-technology-savvy staff who can envision and develop Internet-based health systems. Chief information officers and other high-level information systems professionals should have expertise in both information technology and health care.

Recommendation 3.3. The Department of Health and Human Services should commission a study of the health information technology workforce to determine whether the supply of such workers balances the demand for them, to identify the kinds of training and education that workers at different levels will need, and to develop recommendations for ensuring an adequate supply of people with training at the intersection of information technology and health.

For health applications of the Internet to be envisioned, developed,

and deployed, knowledgeable workers are needed who understand both the technical capabilities of the Internet and the nuances of operating in a health context. Information technology firms perceive a distinct shortage in the number of qualified information technology workers they can hire, although a formal assessment of the situation is still under way.³ The supply of workers with both information technology and health skills may be even tighter, but the situation has not been well investigated. Impressions of supply and demand tend to be based on anecdotal evidence about the demand for graduates of existing programs. All health organizations will be affected by the Internet and will need to develop competencies to work with it. To date, support for training in areas such as medical informatics has come almost exclusively from the National Library of Medicine, but it will soon need to come from other sources as well if the pool of qualified workers is to grow. It is not clear what kinds of skills workers at different levels in a health organization will need. Additional study is required to determine the extent of the problem and the best avenues for addressing it.

Addressing Policy Issues

A number of impediments to Internet-based health care must be addressed at a policy level. While it is not possible to identify or predict all the barriers that will arise as the Internet becomes more widely used in the health sector, there is a need to ensure that regulatory barriers do not unnecessarily impede application of the technology as it evolves and that all demographic groups are active participants in health on the Internet. The impediments include issues such as payment for services delivered via electronic networks, licensure, and malpractice. These impediments have slowed past attempts to deploy telemedicine services more broadly; they exemplify the kinds of mismatches that may result from attempts to use laws and regulations created in the past to govern a growing range of Internet-mediated services. What were minor concerns in the past may become more important as the capabilities of the Internet grow and its applications expand into health and health care. The committee notes that it is not possible to enumerate all the ways in which the legislative and regulatory environment may need to be altered to accommodate the Internet in health and health care, but some of the areas that have been identified hint at the larger set of issues to come. Although it is beyond the scope of its charge and its expertise to provide recommendations for remedying these policy concerns, the committee notes that they pose a significant barrier to the deployment of Internet-based applications in health and health care and makes the following recommendation to hasten their resolution:

Recommendation 4.1. The Department of Health and Human Services should more aggressively address the broad set of policy issues that influence the development, deployment, and adoption of Internet-based applications in the health sector.

Addressing the policy issues that are raised in this report will require strong leadership from federal health agencies. Not only does DHHS need to ensure that concerns and needs of the health community are reflected in attempts to address policy issues such as intellectual property protection, privacy, and access to the information infrastructure, but it can also help to ensure greater coordination of the efforts of federal health agencies in these areas. Elements of DHHS are involved in missions that are affected by these issues and have taken steps to address them, and the DHHS itself has taken steps to address issues such as the privacy and security of electronic medical records. Additional focus would help ensure that these issues are suitably addressed by the policy-making community. DHHS could play a number of roles:

- *Providing strategic leadership for Internet-related programs within the department and its constituent agencies and coordinating them with those of other federal agencies.* Because the Internet may transform a number of aspects of health care, including many of those overseen by federal agencies, government would be well served by a process that would assess, manage, and monitor the implications of the Internet for federal health activities. Many of the agencies within DHHS, including the Health Care Financing Administration, the National Institutes of Health, and the Centers for Disease Control and Prevention, have ongoing plans to evaluate Internet applications in some of their mission-critical operations, but little higher-level strategy exists within DHHS for better harnessing the capabilities of the Internet throughout the organization.
- *Convening public and private bodies to identify and examine issues related to the Internet and health care.* These bodies could help federal agencies identify issues that need to be resolved, provide guidance on the kinds of approaches that might be most effective, and ensure greater coordination of public and private efforts. The National Committee on Vital and Health Statistics (NCVHS) has been playing a similar role in the area of privacy and security of electronic health information and could serve as the model—or the seed—for other such groups.
- *Exploring cross-cutting issues that affect a number of government health agencies and developing programs for addressing them.* For example, DHHS could examine ways to implement a public key infrastructure that would support a range of federal health activities, from provision of care in DOD and VA facilities to payment of Medicare and Medicaid claims.

- *Encouraging sharing of information and perspectives among health-related agencies.* Development of a health information infrastructure will involve a multitude of participants with different responsibilities and interests: provision of care, payment for care, monitoring of care, health-related research, public health, and others. Because of its broad interest and activity in many of these areas, DHHS could serve as a focal point for encouraging dialog among these constituents and coordinating activities.

- *Advancing the national debate regarding key information technology issues that affect health care.* As noted throughout this report, considerable uncertainty exists about the ways in which the Internet is likely to influence the health sector and about the effectiveness of different applications. Based on experience with its own systems, its interactions with the private sector, and its ability to serve as a neutral meeting ground, DHHS could help the entire health community become better informed about use of the Internet in health care and about the technical, organizational, and policy issues that must be addressed.

- *Creating organizational structures to ensure that issues at the nexus of health and information technology are identified and addressed promptly and efficiently.* The speed with which the Internet and its applications are advancing requires proactive consideration of opportunities and challenges and the structures that can respond to a rapidly changing environment. Progress has been made along some of these lines by establishing the DHHS Data Council and redefining the charter of the NCVHS. The NCVHS has, in fact, begun to address the creation of a health information infrastructure (NCVHS, 1998), but additional effort along these lines will be needed to ensure that adequate attention is paid to this emerging area.

A FINAL WORD

The recommendations offered above are intended to set the nation on a course that will ensure that technology, organizational practices, and public policies converge in ways that will lead to broader deployment of Internet-based systems in health applications (see Box 6.1). Undoubtedly, this course will have to be recharted over time to reflect progress made along each of the fronts and as Internet-mediated health processes continue to unfold. Changes in the structure of the nation's health care will continue to drive the kinds of health-related systems that will operate over the Internet, and the Internet will, in turn, drive changes in the structure and nature of health care. Continued dialog between the information technology community and the health community will be central to ensuring that the Internet evolves in ways that meet the ever-changing demands and specialized needs of the health sector—and to ensuring that the Internet will support the health of the nation.

BOX 6.1 A View of the Future

Two scenarios demonstrate the kinds of capabilities that could be achieved if the recommendations outlined in this report are implemented.¹

Scenario 1: Georgia Johnson, a 64-year-old widow with hypertension and congestive heart failure, lives in Quincy, Pennsylvania. Her physician, Dr. Ramesh, is located in Baltimore, where she lived before her husband died. Ms. Johnson sees her doctor biweekly, through a videoconference visit that is hosted in her home by Marcus Brown, a student in the University of Virginia's distance-learning nurse-practitioner program. During a routine video visit, Dr. Ramesh tells Georgia that she needs to reduce her salt intake and gives her an information prescription. Marcus fills the prescription by adding two new features to Georgia's personal health Web page: an interactive diary that she will use to track her own sodium intake and a three-part multimedia series on living with congestive heart failure. During the visit, Georgia hears the doctor suggest that Marcus visit the NLM's Internet library of chest sounds to learn more about how to recognize congestive heart failure. At Georgia's request, Marcus shows her how to use this resource and links it to Georgia's Web page, too. After showing Georgia how to post a copy of the interactive diary to her electronic medical record, Marcus heads out for his next home visit. As he walks to his car, Marcus calls his e-transcription service, logs in, and dictates visit notes into Georgia's electronic medical record.

Scenario 2: Juanita and Santos Del Rios have lived in the United States for 3 years. They live in a large apartment complex in downtown Miami with their four children, who range in age from 2 to 9. They are learning English, but they also depend on their 9-year-old, Rosa, who learns English in school, to help translate for them. Last year, their health plan sent someone to deliver a home health kit consisting of Internet connectivity, a digital thermometer, a heart rate monitor, a stethoscope, and a videocamera. The installer, who spoke Spanish, showed them how to use the equipment and to connect to the Internet through their television. One night, Juanita is awakened by the baby, who is coughing, wheezing, and crying. She wakes up Rosa and asks her to call the HMO for advice. Hearing the

REFERENCES

- Clausing, Jeri. 2000. "Report Rings Alarm Bells About Privacy on the Internet," *New York Times*, February 7.
- Health Care Financing Administration (HCFA). 1997. *Telemedicine Report to Congress*, Department of Health and Human Services, Washington, D.C., December 4. Available online at <<http://www.hcfa.gov/pubforms/telemed.pdf>>.
- National Committee on Vital and Health Statistics (NCVHS). 1998. *Assuring a Health Dimension for the National Information Infrastructure*, October 14. Available online at <<http://www.ncvhs.hhs.gov/hii-nii.htm>>.

symptoms, the pediatrician on call asks to have a quick look at the baby. Rosa turns on the set, while Juanita sets up the health kit. Rosa establishes an encrypted session with the server and reserves a suitable level of bandwidth for a video-conference, using the "bandwidth wizard" on the server. When the connection is complete, Rosa selects the "habla Español" option for simultaneous subtitles, so her mother can communicate directly with the doctor, and pages the doctor. Remembering a recent news item about a rash of respiratory problems in this neighborhood, the pediatrician links to the local public health department e-channel to check on specifics. While waiting for the doctor to finish her research, Juanita scans the index of the HMO's self-help library and downloads two items, *Las Sintomas de la Asma* and *El Gripe y su Niño*, for later reading.

These scenarios require advances in technology, organizational capabilities, and public policy in order to become commonplace in the future. The first scenario requires dependable bandwidth on demand, authenticated remote access to patient records, and widely accessible Internet-based collections of resources like the multimedia series. It also requires cross-border licensing arrangements and health care reimbursement policies that cover this kind of service. The technical requirements for the second scenario include cable modems, reservable bandwidth, encrypted server access, digital instruments, and instantaneous language translation. The nontechnical requirements include a health plan that supports home telemedicine and online access to Spanish-language consumer health information. Most features of these scenarios exist now but are not widely available or easily accessible to those who may need them. The expanded capabilities for health care outlined in these scenarios could be achieved in the near future given action on some of the recommendations outlined in this report. Even more exciting are the applications that could be imagined if the nation were to begin to use the Internet to its full potential in health applications.

¹These scenarios were first described by Valerie Florance of the Association of American Medical Colleges in a public briefing to release a prepublication version of this report. Because of the interest they generated, they have been included here.

President's Information Technology Advisory Committee (PITAC). 1999. *Information Technology Research: Investing in Our Future*, National Coordination Office for Computing, Information, and Communications, Arlington, Va., February 24. Available online at <<http://www.ccic.gov/ac/report/>>.

Working Group on Biomedical Computing, Advisory Committee to the Director. 1999. *The Biomedical Information Science and Technology Initiative*, National Institutes of Health, Bethesda, Md., June 3. Available online at <<http://www.nih.gov/welcome/director/060399.htm>>.

NOTES

1. HCFA is already working to investigate means of reimbursing the costs of some types of remote consultation and has funded demonstration projects to explore alternative payment schemes. For information on HCFA's efforts on paying for remote health services, see HCFA (1997).
2. The Agency for Health Care Policy and Research was recently renamed the Agency for Healthcare Research and Quality (AHRQ).
3. The Computer Science and Telecommunications Board is studying workforce needs in information technology. Additional information on this project is available online at www.cstb.org.

Appendixes

APPENDIX A

Site Visit Summaries

As part of its data-gathering activities, the Committee on Enhancing the Internet for Health Applications visited eight sites that were either developing health-related applications of the Internet or engaged in health-related activities that could be transferred to the Internet in the future. These visits, conducted between December 1998 and February 1999, spanned half a day to one full day each. They provided committee members with a snapshot of the state of deployment of the Internet in the health community at that point in time and an opportunity to better understand the technical and other challenges associated with use of the Internet in support of health care.

This appendix briefly summarizes the committee's eight site visits, including four in California, two in North Carolina, and two in Washington State. The sites were the Laboratory for Radiological Informatics at the University of California at San Francisco (UCSF); Kaiser-Permanente of Northern California, Oakland; Stanford Center for Professional Development, Stanford University; the National Aeronautics and Space Administration (NASA) Ames Research Center, Mountain View, California; the Center for Health Sciences Communication at East Carolina University (ECU), Greenville, North Carolina; the University of North Carolina (UNC), Chapel Hill; and the University of Washington (UW) and Regence BlueShield in Seattle.

LABORATORY FOR RADIOLOGICAL INFORMATICS

Members of the study committee spent half a day at the UCSF Laboratory for Radiological Informatics (LRI) on December 16, 1999. H.K. (Bernie) Huang and his colleagues at LRI demonstrated systems designed to share three different types of telemedical images: digital mammograms; cardiograms; and neurological images made by magnetic resonance imaging (MRI) and computerized tomography (CT).¹ For the most part, these systems make use of a picture archiving and communications system (PACS) at UCSF that stores digital medical images in several terabytes of optical disk storage and an asynchronous transfer mode (ATM) wide-area network (WAN) that connects UCSF with nearby Mount Zion Hospital. Originally, the laboratory used a dedicated WAN, but then the university installed a synchronous optical network (SONET) ring, which now supports the network. A T1 connection—which supports data rates of 1.544 megabits per second (Mbps)—links UCSF with Stanford University Medical Center, an hour's drive away.

Telemammography

Steve Frankel of the Breast Imaging Section at UCSF and Andrew Lou of LRI provided an overview of the teleimaging system. UCSF now has two full-field digital telemammography systems that produce images of 40 to 60 megabytes (MB) compressed (using an acceptably lossless compression scheme). A typical study generates four such images, two of each breast, but also requires a historical set of equal size for comparison. Images can be transmitted across the WAN for remote interpretation and diagnosis or real-time reading by expert mammographers. In the demonstration, staff physicians at UCSF sent images to Mount Zion for interpretation. Physicians at both UCSF and Mount Zion used high-resolution (2,000 × 2,000 pixel) monitors to view the images. Using custom software developed at UCSF, the referring and consulting physicians could use an on-screen dual-pointer to identify objects of interest and change the brightness and contrast of the images to aid in interpretation. An electronic magnifying glass enabled mammographers to examine portions of the image in greater detail.

Expert reading of mammograms in real time is not necessarily needed for regular screenings, but it can be useful if potential abnormalities are discovered. In such cases, remote experts can provide faster diagnoses or request additional images before the patient leaves the mammography center. Busy centers may examine 80 to 100 women per day (20 per day is more typical for an average center), with the expectation that images will

be read the following day. Longer delays are not uncommon; many mobile mammogram centers have 2-week turnaround times.

Telemammography is viewed as a means of supporting increased demand for mammograms. The National Cancer Institute (NCI) recommendation that all women over age 40 have an annual mammogram could, if widely heeded, greatly increase the rate of mammography, UCSF system developers noted. They also said the nation has too few expert mammographers to handle the increased volume and, moreover, that many rural areas have no local experts. To provide teleradiology interpretation services for other health care organizations, UCSF has established a consortium with Emory University in Atlanta; Wake Forest University in Winston-Salem, North Carolina; Brigham and Women's Hospital in Boston; and the University of Pennsylvania. The consortium, named Telequest, accepts images over a dedicated T1 line and provides transcribed diagnoses. Participating physicians must be cross-licensed in the examination sites.

Cardiology Teleconferencing

Tony Chou, director of the catheterization labs, described a cardiology teleconferencing system linking UCSF and Stanford University. This application, which runs on a T1 line, is used as an educational tool to enable physicians to present cases to colleagues for postmortem reviews. It is not yet used for diagnostic purposes but it could be, once the institutions install digitized catheterization labs. The system has been used to review angiography and intravascular ultrasound images, which are currently captured as analog video and digitized. Digitized angiogram files are roughly 60 MB in size; intravascular ultrasounds are about 50 MB. The size of these files is expected to grow by a factor of 10 once the fully digital systems are installed. Videos are displayed at a rate of 25 frames per second.

In applications tested to date, digitized videos of angiograms and intravascular ultrasound have been transmitted across the T1 network and replicated on both sides of the connection, a process that takes approximately 5 minutes. In one case, images were transferred to a medical center in Germany overnight. Participants in the teleconferences can examine images simultaneously, using on-screen pointers to note items of interest and zooming in on particular portions of the video. Physicians report that the quality of the video is already sufficient for most diagnoses, but the system has been used only for reviewing outcomes and medical decision making.

Neuroimaging

Bill Dillon, chief of the Neuroimaging Department, provided an overview of the neuroimaging teleconsultation system linking UCSF and Mount Zion. The system is used primarily for MRI and CT scans, which usually are digitally captured and stored in the PACS. Before the PACS was established, some 15 to 20 percent of films were lost and hundreds went unread—meaning that the hospital could not bill for the service. The PACS allows immediate access to images and was accepted rapidly by physicians, who found that the system greatly increased the ease of locating needed studies before meeting with a patient.

The WAN has enabled UCSF to offer image interpretation services. In fact, UCSF now has a resident on call to read neuroimages taken at Mount Zion. Such centralization of interpretation skills helps accommodate the interests of the state of California, which is pressuring California hospitals and medical schools to reduce the number of specialists and trainees in specialty areas. Similar pressures are being felt at the national level as health care providers merge and attempt to consolidate services and eliminate duplicative capabilities across large health care delivery systems.

KAISER-PERMANENTE OF NORTHERN CALIFORNIA

The study committee visited Kaiser-Permanente of Northern California on the afternoon of December 16, 1999. Kaiser-Permanente is the nation's largest health maintenance organization (HMO), with more than 9.4 million members in 20 states. It has approximately 100,000 employees, including 10,000 physicians and 30,000 nurses and pharmacists, and it operates some 30 hospitals with more than 8,000 beds. The HMO is not a single entity but, rather, an affiliation of two separate organizations: the Kaiser Foundation Health Plan, which handles administrative and management functions, and Permanente Medical Group, which consists of 12 separate groups of practitioners. Kaiser-Permanente has experimented with network-based telemedicine applications, particularly in teledermatology and teleradiology. Its most notable success was with a system for retinal screening of diabetic patients, which began when a clinic bought a camera and began sending images electronically to an expert reader for interpretation. This simple procedure increased initial screening rates from 30 to 93 percent of all patients who met the criteria in the risk guidelines.

Kaiser-Permanente staff began evaluating different Internet applications at the request of its chief executive officer. The result was a unified, three-pronged strategy consisting of a provider-oriented system; a

customer-focused system; and a common, shared database. Development of the provider system and shared database—the Permanente Knowledge Connection (PKC)—has proceeded under the auspices of Kaiser-Permanente’s Care Management Institute. Development of the consumer component, the KPOnline system, has proceeded separately. The discussions during the site visit focused on both the institute’s efforts to develop internal applications for use by care providers and on the efforts of KPOnline, a customer-focused system.

Permanente Knowledge Connection

Peter Juhn, executive director, described the activities of the Care Management Institute (CMI), a national entity within Kaiser-Permanente that operates on behalf of both the Kaiser Foundation Health Plan and the Permanente Medical Group. CMI was established in 1997 and employs approximately 80 people, 30 of whom work in Oakland and the rest of whom are distributed throughout the Kaiser-Permanente system. Its principal function is to develop evidence-based approaches to care management. This work has three areas of emphasis: content, measurement, and implementation. Content work includes the development of management programs for conditions such as diabetes, asthma, and depression, as well as an overall compendium of clinical best practices. Measurement work encompasses large-scale national studies of health outcomes and has included studies of 200,000 diabetic patients, 320,000 cardiovascular patients, and 90,000 asthma patients in the Kaiser-Permanente system. The implementation work builds on the content foundation, using the collected information as a basis for clinical systems that can influence care at the point of delivery. The objective of these activities is to change care providers’ behavior to coincide with best practices developed throughout the Kaiser-Permanente system. Such systems can have dramatic effects on care. Over the previous year and a half, Kaiser-Permanente found a 10 to 15 percent gain in the use of practice guidelines in some areas where it had developed content.

The PKC is a network-based application that was developed to support the CMI’s objective of improving care. Its primary function is to allow care providers to access current CMI content on best practices. The CMI staff realized that each Kaiser-Permanente site was gathering useful information that could benefit other local and national care providers, but little of it was shared. Using Kaiser-Permanente’s national intranet, the PKC now has national and regional databases of best-practice information that is vetted before it is entered into the databases. A board of directors that represents all executive and physician groups approves information for inclusion in the national database. Twelve regional

groups establish approval processes for their own regions, and local offices establish processes for local approvals. This structure balances local freedom against the need for greater structure at the highest levels of the organization. The databases are linked together through the national office's intranet Web site, which makes the information searchable and reduces duplication of effort. In addition to linking care management information, the PKC provides a centralized outlet for other information resources of interest to care providers. It contains a section for continuing medical education (CME) that allows users to look up their CME credits. It also provides access to online text books and journals and supports discussion groups with threaded discussions. Its workgroup functions enable users to post material and conduct national meetings in a virtual manner. Workgroups can be designated for members only, with membership defined by the chair.

The CMI staff plan to build on the PKC's current capabilities to make PKC more useful to care providers. The staff will address topics such as improving search capabilities, tailoring information to provider needs, and expanding access to affiliated care providers. The PKC can support searches across the national intranet using a commercial search engine, but the CMI is looking into ways to improve search capabilities and help different types of users find information they need, perhaps by adding "metadata" capabilities. The staff also plan to use "push" (as opposed to pull) technologies and customization to provide users with information relevant to their immediate and long-term needs. Plans for this enhancement are tied to efforts to place a computer on every physician's desktop and provide access to a clinical information system (more than 40 percent of Kaiser physicians had desktop computers at the time of the site visit). Overall, however, the Kaiser-Permanente staff do not consider technology a limiting factor. They have modest aims for now, and technology is often a distraction from reaching other goals. They prefer to take known technology and find ways of using it to support their business, rather than pushing the technological envelope.

Nevertheless, determining how best to expand the PKC to affiliated providers will entail both technical and policy considerations. Kaiser-Permanente has agreements with approximately 40,000 affiliate providers in physician groups outside of California. How can Kaiser convince the affiliates to care for Kaiser patients according to Kaiser practices? Should the affiliates have access to all the available knowledge, or should some information be considered proprietary? How can information be layered and filtered easily to accommodate restrictions placed on affiliate access? How can firewalls be extended to include specific partners while still providing adequate protection for Kaiser's information systems? At present, the system does not contain patient-level data, which alleviates

some concerns regarding security, but Kaiser needs to evaluate alternative ways of providing security on an extranet and/or the Internet. It also needs to determine how best to blend public and proprietary information to benefit its providers. Kaiser officials do not consider their practice guidelines proprietary—and would even like to make them public—but the tools for implementing these practices are unique and will be kept proprietary. Kaiser would like to translate the guidelines into lay terms and make them available on its consumer-oriented Web site.

The Kaiser staff anticipate three types of outcomes from the PKC. First are the financial benefits. They want to leverage the size of the organization while avoiding duplication of effort—something the PKC can facilitate. The staff will try to determine how much money was saved by not starting new programs or sustaining existing ones because PKC indicated that similar work may prompt changes in clinical decision making that yield improved clinical outcomes and use of facilities. Other benefits will accrue from improved knowledge management, which should succeed in educating care providers about new diagnostic approaches and new techniques. The Kaiser staff hope to make the link between successful patient-provider interactions and the PKC system evident, to demonstrate the system's ability to support corporate objectives.

If the system is to be successful, then usage rates must increase. At the time of the site visit, 2,700 of Kaiser-Permanente's 10,000 physicians used the system, but all were expected to use it by the end of 1999. To attain and maintain that level of usage, the system will need to prove itself capable of educating physicians about the new diagnostic approaches and clinical techniques. The system will also need to demonstrate its capability to enhance the goals of the organization. Management will need to see a direct linkage between improved patient-provider interactions and the tools that support that interaction. Privacy issues will also need to be addressed so that providers know whether information will be collected on their searches of medical literature and whether such searches will be viewed as positive (i.e., the provider is engaging in continuous learning) or negative (i.e., searches indicate gaps in a provider's knowledge).

KPOne

Anna-Lisa Silvestre and Richard Leopold described Kaiser's consumer-oriented Web site, the primary component of KPOne. KPOne is a three-tiered system with a Web page interface that interacts with legacy systems through an intermediate object layer. Ms. Silvestre views the Web site as a service for interacting with Kaiser-Permanente members. It is not a marketing tool or a mechanism for providing content. Rather, it is intended to provide members with an alternative to telephone calls and

office visits. The site provides members with capabilities for messaging, scheduling appointments, and checking prescriptions 24 hours a day, 7 days a week. Eventually, the Kaiser staff would like to integrate KPOne with the PKC so they can take the information gathered from patients through KPOne and incorporate it into practice guidelines and, conversely, incorporate practice guidelines into chat rooms and e-mail discussions with patients. The goal is to help patients better understand health information, tailor it to them, connect them with providers, and to help them make sound, coordinated decisions regarding their care.

At the time of the site visit, approximately 16,000 registered users had logged on to KPOne more than once. This number is just a small fraction of Kaiser-Permanente's 9.4 million members, but the staff believes usage will increase as more of them gain Internet access. The recent biannual survey reported that 72 percent of members are adults, 53 percent of whom have Internet access at home, work, or both. Ten percent of all members have requested a personal identification number (PIN) to use with the system. Hence, Kaiser expects hundreds of thousands of members to access the system simultaneously in the near future (the target was 350,000 active users in 1999) and is in the process of procuring additional servers to handle the load. The organization is still trying to learn how members use the site and what types of services they seek. Kaiser is starting to collect data on site usage (it does not track an individual's movements within the site) but does not yet have adequate volume to examine usage by demographic category.

Consumers are coming online quickly, and KPOne is expected to become a basic utility that will benefit both consumers and Kaiser-Permanente. A basic evaluation will be performed to determine who uses the system, their level of satisfaction, and the overall utility of the system. More formal cost-benefit analyses will also be conducted. A tangible cost-benefit analysis will evaluate processes such as online pharmacy refills and automated appointment scheduling and compare them to more traditional, manual processes. The Kaiser staff expects that online pharmaceutical services have the highest benefits per unit cost because filling prescriptions becomes much less expensive when done in high volumes. Other tangible benefits not associated with cost reductions, such as helping members make good decisions, will also be considered. Some of the benefits will be difficult to quantify, but there may be ways to determine if online material helped to prevent an unnecessary visit to a Kaiser facility, improve the appropriateness of a subsequently scheduled visit, or increase membership retention rates.

Early experience with KPOne has uncovered numerous issues that need to be resolved. One of the main ones is the need for standards for measuring the quality of online transactions. For example, with respect to

online questions to nurses, it might be desirable to track the time of receipt, the time at which the question was answered, the time at which the member retrieved the response, and whether the answer given was valid. Similar standards are now in place to help train nurses who provide advice over the telephone, but such situations involve near-real-time feedback. In addition, nurses need training in how to provide care based strictly on text input (which creates a record of the interaction), with no voice or personal interaction with patients. At the time of the site visit, all of Kaiser-Permanente's care providers had e-mail accounts, but patients were not yet using them. The use of clinical e-mail raises several questions, as yet unanswered, regarding medical records.

Other issues include the determination of rules for intervening in sponsored chat groups. Kaiser-Permanente had a case in which a member's postings to a chat group suggested suicidal tendencies. The organization had to decide whether and how to intervene. Should it link the user's anonymous login with the medical record database to check the medical history and find contact information? In this case, the Kaiser staff did just that, and an advice nurse called the patient and arranged an appointment, which revealed that the patient was indeed suicidal. Events such as this prompt policy reviews.

Kaiser has established a set of technical measures and administrative processes to provide security on the Web site. The site uses Secure Socket Layer (SSL) encryption to protect messages between members and Kaiser, and members need a PIN to access chat groups. Members are required to provide their membership number and address to obtain a PIN. In addition, the manager of the business unit is a security trustee and has to ensure that policies and procedures are in place. These policies are reviewed regularly and upgraded as needed, and attempts are made to achieve consistency between online and off-line policies. For example, Kaiser dropped the authentication requirement for scheduling an appointment online because such authentication is not performed when scheduling appointments via the telephone.

Members of the Kaiser-Permanente staff identified several technical capabilities they would like to be able to incorporate into KPOnline:

- *Authentication technologies* that would allow multiple users within a single household to use the same computer but keep their information separate. In the current KPOnline system, if users forget to log out of a session, then other family members can see what they did and what information they retrieved. The Kaiser staff would like to obtain improved technologies to identify users and authenticate their identities. One possible solution is biometrics, but this approach would be costly to implement across many computers.

- *Higher-speed Internet connections for end users.* The objective is to enable members to have “a T1 experience” while using a 28.8 kilobit per second (kbps) modem. Kaiser has its own wide-area network that connects all Kaiser facilities. Should the organization become the equivalent of an Internet service provider (ISP) so that it can provide high-speed connections and ensure security? Should care providers act as ISPs themselves so that members can connect to the network using high-speed cable modems or digital subscriber line technologies?

- *A standardized Web browser.* The Kaiser support staff spends considerable time helping customers configure their browsers to work with KPOne. Users sometimes confuse the application with the infrastructure and call Kaiser’s support desk to report problems that are not associated with KPOne but rather with their ISP.

- *A simple telemedicine terminal.* A computer that supports videoconferencing, and sensors for collecting diagnostic information, would enable care providers (such as triage nurses) to interpret and evaluate cases based on more than just text-based information. A pilot program using this technology is ongoing in the mid-Atlantic region to monitor diabetic patients, and there are plans for another pilot program in the Northeast with congestive heart failure patients. Telephone lines are used to transmit information. The main problem experienced to date is not capturing the information but incorporating it into the medical record. Security issues also need to be addressed.

STANFORD CENTER FOR PROFESSIONAL DEVELOPMENT

Committee members visited the Stanford Center for Professional Development on December 16, 1998, to meet with Andrew DiPaolo, its executive director, and members of his staff and learn more about the Stanford Online system. Stanford Online is the current Web-based implementation of the Stanford Instructional Television Network, a system pioneered by Stanford University’s engineering school to teach courses at a distance, largely (in the past) to employees in local industry. The system allows students or employees to remain at their company sites while taking courses. Through the Honors Co-op Program, workers can maintain full-time employment while studying to earn a full-fledged master’s degree. The program focuses on engineering classes, but members of the University’s section on medical informatics have taught courses using the system.

The system started with live broadcasts over private systems between Stanford and the subscribing companies, generally with two-way audio supplementing one-way video so that students could ask questions of the instructors and participate in discussions. Subsequently, workers’ need

for flexibility in scheduling of studies led industry to call for videotapes that could be watched at the convenience of students rather than only during live video feeds. The program was broadened into the Stanford Center for Professional Development (SCPD) when it became clear that the distance-education model would be more than simply broadcast television. Anoop Gupta, a computer science faculty member, had developed Vxtreme, a technology for streaming video, which was experimentally adopted by the SCPD as a means of distributing "videotapes" by the Internet rather than by courier. Vxtreme was later acquired by Microsoft Corporation, and the SCPD continues to use the Microsoft streaming video products (which are incompatible with RealVideo).

Today, Stanford Online has become an important means of disseminating video-recorded classes along with slides and photographs of blackboards or projection screens. Videos are up on the Web within 3 hours of the class. Class videos are indexed into segments by undergraduate students, who manually perform this function immediately after the class. Methods are being explored for automating the indexing and matching the video to the master copy of slides. A previous year's videos are discarded because the courses are taught again and the lectures are continually updated.

Stanford Online students see three things in their browser window: a relatively small video of the instructor giving the class; an index of topics covered by the lecture (topics can be selected as the student wishes); and PowerPoint slides or, alternatively, photos of projected slides or blackboards, coordinated with the audiovisual track. Even with a 28.8 kbps connection, students can work through a lecture and skip around based on the indexing outline. The browser has a plug-in for streaming video. The system is intended to be used asynchronously; students cannot ask questions in real-time via the Internet, but they can send e-mail to professors and participate in online discussion groups.

Students can take courses whenever they want if they are not seeking credit. If they want credit, then it is best for them to take the course during the regular semester, at the same pace as on-campus students, to take advantage of resources such as teaching assistants who are available to grade homework. Class videos are also broadcast to dormitories at specific times over the campus television system, enabling Stanford students who live on campus to view lectures they missed or wish to view a second time. Television viewing does not allow user-controlled access to specific portions of the lectures, however, as does the Stanford Online version. With the availability of both scheduled replays of lectures and Stanford Online video on demand, it is now possible for students to take two courses that meet at the same time. Students like the asynchronous method because they can follow courses when they are away from the

campus. However, the online capability can reduce the number of students attending the live lecture; in fact, some students have opted not to attend class at all, especially early in the morning. The smaller number of live participants can reduce class interactions and force changes in the way professors deliver lectures.

File sizes for 10 minutes of video range from 1.6 MB to 38 MB depending on the bit rate at which they are sent (Table A.1). The size of the file is minimally affected by the size of the video image on the computer screen. The image quality of the frame degrades as the number of frames per second (fps) increases, even though the throughput (bit rate) may remain constant. The frame rate depends on the video capture card that converts video imagery into digital format and on the horsepower of the authoring workstation.

It is not clear whether there are unique networking requirements for medical education. Some technologies may be more important for the delivery of lectures dealing with medicine than of those dealing with other fields. For example, techniques are needed for tracking laser pointers on slides, an essential in teaching radiology classes. The SCPD is trying to develop techniques for overlaying the movement of a pointer on the screen and ways of noting changes between or within items on the screen. It is conducting formal studies of performance and has involved the university's education school in the evaluation process. At present, Stanford Online is a technical implementation of current models of teaching; little thought has been given to changing the educational model to fit the technology. This approach has at least one advantage: it allows lecturers to deliver courses online without changing their form or structure. Some minor changes have been made, such as the use of bigger chalk to make writing more visible.

The SCPD makes money for the university, enough to pay for a large studio and the operation of five channels. Net income is roughly \$4 million per year, some of which is divided among university departments. Faculty members have financial incentives to teach courses online: They receive a share of the tuition for online courses, as do their departments.

TABLE A.1 File Size of 10 Minutes of Video at Different Bit Rates

	Bit Rate (kbps)				
	22	50	100	256	512
File size (kB)	1,656	3,756	7,507	19,018	38,408

SOURCE: Michael Rouan, Stanford University, personal communication, January 18, 1999.

But new financial models may be needed to pay for new services. So far, off-campus services are offered only to "member companies." This approach would need to change if the university decided to allow individual students to pay for their own classes, as they would for CME.

The SCPD program—the online portion, in particular—is very popular. Online participation may cut into the broadcast model. Obviously, this could be a big business for Stanford, a trend that threatens small colleges. SCPD has the capability on existing servers to teach 2,000 simultaneous users; additional capacity could be brought online to expand its scope.

NASA AMES RESEARCH CENTER

The committee visited NASA Ames Research Center on December 16, 1999. Committee members were hosted by Muriel Ross, then director of the Ames Center for Bioinformatics, and her team of biocomputing researchers. NASA's space exploration mission, particularly the possibility of crewed flights over long distances, poses particular challenges in network communication. As space vehicles venture further from Earth and for longer periods of time, it is unlikely that the crew will have all the needed medical expertise. Accordingly, the NASA team is focusing primarily on how to provide health care using telemedicine over long distances with high latency. The emphasis is on technologies that consume little power (such as personal computers) and that can leverage the power of more costly hardware through network connections. They are addressing two particular aspects of this challenge: the need to support remote, collaborative medical assessment and treatment using the Internet (as an approximation of the network that would be available using wireless communication) and the need to support rapid, accurate three-dimensional rendering of organs for assessment at a distance by medical professionals.

As part of its preparation for remote medical consultations, NASA Ames has established a collaboration among in-house staff and Stanford University scientists to experiment with the transmission of images, in particular complex three-dimensional models of hearts, skulls, and other structures. The challenge for NASA networking scientists is to send the images simultaneously to multiple sites (i.e., multicasting, as opposed to point-to-point communication). They have created a testbed network that links NASA Ames with Stanford University Medical Center, the Cleveland Foundation Clinic in Ohio (through NASA's Glenn Research Center), Salinas Valley Memorial Hospital (through the University of California at Santa Cruz), and the Navajo Nation at its Northern Navajo Medical Center in New Mexico. The network is extremely complicated, linking local-area networks (LANs) at the participating institutions to a variety of

high-speed WANs, including the NASA Research and Educational Network; Abilene, a high-speed network backbone project of the University Consortium for Advanced Internet Development; and the very high performance backbone network service (vBNS), a research network launched through a 5-year cooperative agreement between MCI and the National Science Foundation. Connections to the Navajo Nation run through satellite and high-speed ground links.

The three-dimensional medical images used in NASA's tests contain several million polygons and are too complicated for typical computer workstations to render. NASA therefore planned to perform volume rendering on a powerful central graphics computer and then transmit the images (or changes in the images) to peripheral workstations at the remote sites. Because the files would be too large for straightforward transfer on today's Internet, NASA planned to use the Abilene network, which operates at 9,920 gigabits per second (Gbps). Limitations in bandwidth in various parts of the network, however, precluded even this design. Consequently, NASA designed a hybrid system in which high-resolution static images (containing roughly 1.2 million polygons) are transmitted from the centralized site and real-time rendering of object manipulations is handled at lower resolution (approximately 20,000 polygons) by the individual sites. Changes are made visible to all participants, and at the end of a manipulation, the centralized NASA computer sends an updated, full-resolution image back to each of the remote sites using multicast technologies. Users at the collaborating sites wear special glasses to view the three-dimensional images. The images for cranial-facial surgery consist of skin and skull only, but future modeling will include brain tissue and blood vessels.

The testbed provides simulations of remote medical care for Earth-bound patients as well as astronauts. If, for instance, it was suspected that a child at the Navajo Nation might have a congenital birth defect of the heart, a local physician and technician would send three-dimensional images obtained with ultrasound or MRI technology to consultants at Stanford and Salinas for opinions about treatment. The three-way collaborative environment would allow the specialists and caregivers to assess the situation quickly and weigh two options: local management or transfer to a medical center.

EAST CAROLINA UNIVERSITY

The visit to East Carolina University (ECU) on February 2, 1999, featured extended discussions of the university's Telemedicine Program as well as tours and demonstrations. The study team spent the day with David Balch, director of ECU's Center for Health Sciences Communica-

tion and director of the Telemedicine Program, and Gloria Jones, the Telemedicine Program coordinator. Other staff members and affiliated clinicians and educators participated in portions of the meeting.

Between August 1992 and February 1999, the ECU Telemedicine Program conducted about 2,500 real-time consultations in 31 different fields, using about 60 different doctors. The top five areas are dermatology; cardiology; neurology; gastroenterology; and allergy, asthma, and immunology. In addition, the program has conducted numerous radiological consultations, many on a store-and-forward basis. The Telemedicine Program also runs distance learning programs, the first of which began in 1989 over a statewide network that links major universities in North Carolina.

Between 1991 and 1999, ECU spent about \$8 million on local facilities and infrastructure for telemedicine and teleducation. Physical space is provided by the School of Medicine, and most equipment has been provided through grants from the Office of Rural Health Policy (ORHP) and Health Care Financing Administration (HCFA). The site visit committee toured the facilities and observed a telemedicine consultation to a rural site.

Telemedicine

East Carolina University has two telemedicine suites—an older one and a newer one—with a total of eight rooms. Each room has a telephone, a personal computer for displaying electronic medical records, a display of the remote site, and a display of the consulting physician. The newer rooms are 6 by 8 feet and have a “sound dome” overhead to direct audio to the physician; the older rooms are slightly larger to accommodate a larger display and are soundproofed. The newer rooms were built for about \$8,000, of which \$4,500 was spent on equipment (e.g., computers, monitors, stethoscopes) and \$3,500 was spent on furnishings and construction. Remote sites feature standard arrays of diagnostic equipment that can plug into the network.

Physicians typically schedule three or four consultations in a 2-hour session. The technical staff makes sure that all the equipment is working and that patients’ online records (if there are any) are available. After each consultation, the physician completes a report, which is entered into the clinical database. Cases are presented at remote sites by a nurse or a physicians’ assistant. A distant site provides a half-time dedicated employee who is trained by ECU to be a telemedicine presenter. The presenters generally train with the ECU specialists for a week and come in to the center about every other week; in this way, physicians and remote presenters learn to collaborate before being placed at opposite ends of a

telemedicine link. Physician training includes a 1-hour interview with a simulated patient. ECU charges outsiders for this training; the university needs to subsidize training for its own people. Distant-site physicians are not involved in the telemedicine consult unless they wish to be; only rarely do they participate.

ECU centralizes the management of store-and-forward consultation messages, which typically consist of medical images attached to e-mail. This process allows managers to control bandwidth, monitor usage and flow of information, maintain the equipment, and capture the demographic and survey data they need for their own purposes and to satisfy requirements of the HCFA. ECU tried switching from motion Joint Photographic Experts Group (JPEG) to wavelet compression, but physicians who were already trained did not like the resulting interface change.

At distant sites, ECU pays for all equipment and 3 years of line charges, taking this as a marketing and referral expense. Standard installation costs \$150,000 per site, including remote camera control, two-way audio and video, and continuing medical education content. Installation and bandwidth are sized to support echocardiograms and cineangiograms, the most bandwidth-intensive images. ECU tries to install identical capabilities at all sites, to avoid creating second-class sites with less expensive equipment.

Each telemedicine application supported by ECU has a different business model and a different networking infrastructure. Some of the applications are moneymakers, whereas others are seen as loss leaders for future services and still others may not be sustainable after federal funding expires. For example, the state fully funds Medicaid telemedicine sessions, but this arrangement is in jeopardy because of new national policies. ECU has a waiver from the HCFA and is receiving Medicare reimbursement for consultation sessions, but the paperwork burden is heavy—up to a dozen forms filled out for each telemedicine session. Many of these consultations are store and forward, and new applications in primary care (e.g., wellness) are not considered consultation, so there is no reimbursement potential for these sessions.

The following five subsections describe the primary telemedicine services offered by ECU. The sixth subsection outlines how these services differ in terms of networking infrastructure.

Prison Telemedicine

ECU has a contract with the state of North Carolina for services to the state's central prison. The most popular applications include dermatology and endocrinology. This was ECU's first telemedicine program, and it is a moneymaker. ECU bought the equipment and charged back for it

over 4 years on a flat-rate basis that included equipment costs and maintenance plus 10 free consultations. The break-even point is 800 consultations per year. The state was able to document a per prisoner cost saving from avoiding transportation costs (it formerly cost about \$800 per year to transport a prisoner to another site for health care). Officials also noticed a decline in the number of medical complaints by prisoners, who can no longer view a visit to the doctor as a chance to leave the prison. ECU probably will expand this program to cover the entire North Carolina prison system.

Hospital to Hospital

The hospital-to-hospital telemedicine link was funded initially by the HCFA but now is covered by ECU and the participating hospitals. There are 20 hospitals involved, each with a T1 connection. No more than half of the T1 link is needed for video, so the other half carries data. Originally, the system used so-called triple integrated services digital network (ISDN) at 384 kbps, but after the telephone company switched its billing practices, ECU converted to T1 lines for both conferencing and data. For echocardiography, triple ISDN might work in some situations, but it is not adequate for cine-angiography.

The ECU hospital funds the program at \$150,000 per partner. ECU purchases and loans the equipment to the site; sets up the router, computer, and coder-decoders (codecs); orders the T1; installs the equipment; and trains and supervises the half-time staff person (usually a nurse or physician's assistant) who works the equipment at the remote site. The partner has to agree to fund this position, provide space and insurance for the equipment, and pick up the line charges after 3 to 5 years. The service is not a moneymaker for ECU—no one makes money from consultations. It is viewed as a marketing effort that may establish relationships that bring in patients and eventually lead to procedures, which would generate income. For instance, the hospital might get more cardiac referrals as an indirect result of a well-baby telemedicine program. Such benefits would be difficult to track, however. The financial gain would be virtually impossible to calculate because any lost business opportunities would need to be figured in.

Home Health

The home health program is funded by a grant from the ORHP within the U.S. Department of Health and Human Services. The goal is to determine if hospitalizations can be avoided through better home monitoring of health indicators, such as blood pressure and heart rate. The program

connects 14 homes through ordinary telephone lines (sometimes referred to as "plain old telephone service" or POTS), using two different vendors. At first ECU used two phone lines per home (one for transmitting audio from the stethoscope, the other to carry the patient's voice), but they have since found ways to compress the audio sufficiently to enable the use of a single phone line. Participants are selected on the basis of their frequent use of hospital facilities (e.g., high-risk obstetrics patients or patients who are seen more than twice a week in the emergency room) and their ability to have the system installed in their home for at least 6 months. ECU has approval from Medicaid for the program, but home health is not included in the new regulations promulgated by HCFA regarding government payment of health services delivered over electronic media.

Health Screening

Another ORHP grant is supporting two distributed networks (in contrast to the hub-and-spoke architecture of the other networks) that link mental health centers and schools to ECU. Over one network, ECU provides medical/psychiatric consultations to four mental health centers. The other network links four schools with the county health department, a mental health center, a local pediatrician, and ECU. The networks use 128-kbps ISDN lines (so-called single ISDN, which has proven to be more reliable than the triple-ISDN lines used in other applications). It is hoped that the networks will reduce travel for care providers.

The connection to the schools is used to screen students for mental health problems (25 percent were found to have such problems) and to hold wellness clinics and more general disease screenings. The project has a number of aspects: student health awareness (through focus sessions on topics such as peer relations, nutrition, and physical fitness); clinical consultations; continuing education in cultural sensitivity for teachers; clerkships for students in health professions; and health assessment in the schools. The school-based program sets up live videoconferences between a doctor and school students for discussions of health-related topics.

Store-and-Forward E-mail Consultations

A store-and-forward system is used for multimedia e-mail consultations (using the VisiTran MD application). The system uses the Internet and POTS dial-up to a central server to connect a number of sites on the Outer Banks of North Carolina with ECU. Remote sites use standard desktop computers (Pentium class) to transmit image files averaging 3 MB in size, a figure driven by technical limitations. Nurses could easily take

more images or add more video, but most gateways cannot handle files larger than 5 MB. The system loses money because of the HCFA's reimbursement schedules, but income is generated by procedures that result from referrals.

Variations in Network Infrastructure

The ECU program began in 1988 with a microwave network designed for educational purposes. Since then, the networking infrastructure has evolved into an amalgam of components cobbled together over the years as telecommunications companies have offered different technologies. Old networks have not been upgraded to provide greater uniformity because of the costs involved. Instead, ECU officials have developed a bridge to link their disparate systems. Codecs pass signals from one medium to another in a manner transparent to both ends. For example, the prison contract uses the state microwave system, but the signal is digitized for delivery to ECU's desktop workstations. The digital signal is converted back to microwaves for delivery to the prison. All of ECU's telemedicine activity takes place over leased lines rather than the general Internet, although North Carolina has built a fiber optic superhighway that links many of the state's universities.

The infrastructural variations are not solely technology-driven, because, as David Balch noted, one size does not fit all applications in telemedicine. Different specialties require different amounts of bandwidth. Echocardiography and cine-angiography require the highest bandwidth, 784 kbps; cardiology, emergency room consultations, and neurology require roughly 384 kbps; and psychiatric consultations require 128 kbps. Static image specialties such as radiology, dermatology, and pathology are generally bandwidth-independent because images can be sent in a store-and-forward mode. Home health requires just one telephone line. Hence, various communication media are used, including microwave, T1, single ISDN (128 kbps), triple ISDN (384 kbps), and POTS. Table A.2 summarizes the networking infrastructure for each of the systems ECU currently has in place.

Teleeducation

East Carolina University has eight teleclassrooms on campus. The network was originally built to deliver distance education in engineering and is currently used for nursing classes. (The classrooms are too large for doctor-patient telemedicine applications.) A standard room has all networking modalities available: the North Carolina Information Highway (NCIH) at 45 Mbps; T1 networking at 1.5 Mbps, ISDN dial-up capability;

TABLE A.2 Network Infrastructure Used for Different East Carolina University Telemedicine Services

Telemedicine System	Network Infrastructure
Prison	1/2 T1 line, by microwave
Hospital to hospital	1/2 T1 line or 384 kbps ISDN dial-up
Home health	POTS (one line)
Distributed	Single ISDN (128 kbps)
Store and forward	Volume-based; Internet, POTS, ISDN to Raleigh, T1 from Raleigh to Greenville

SOURCE: David Balch, director of Telemedicine at ECU, site visit by committee on February 2, 1999.

POTS lines; and CU-CME (a software package for sending low-quality video across the Internet). The rooms hold 10 to 15 people each and have cameras aimed at each chair for video.

There is also a larger lecture hall, similarly equipped, that is used for grand rounds sessions. ECU provides grand rounds three times per week to 20 hospitals. These sessions can be carried out using one 15-person conference room and one small auditorium, both of which can connect to 37 sites on the North Carolina Research and Education Network (NCREN) at T3 speeds (or 45 Mbps), 127 NCIH sites at T1 speeds, and other sites using ISDN. PowerPoint slides are converted to video. Grand rounds are seen as a loss leader designed to acclimate remote users to the equipment and style of interaction, in the hopes that they will later try clinical consultations at the remote sites. Scheduling is complex; ECU uses scheduling software developed in-house.

ECU is using Internet-based streaming media (real-time audio), a chat channel, Web pages, and animations (Shockwave) for a small number of nursing courses. Both asynchronous and synchronous instruction are available (i.e., students can participate in lectures in real time or download the video for later viewing). The network was designed to accommodate 28.8 kbps modem speeds so that it could be accessed by a large number of users. That bandwidth is suitable for transmitting lecturers' slides, but it forces audio to be compressed into 6.8 kbps, making the video slightly choppy. The network determines the speed at which to deliver data based on the receiving student's access bandwidth.

The asynchronous system incorporates 20 kbps video, which consumes most of the available bandwidth. PowerPoint slides are converted to HTML for use by both the nursing instructors and the students at home. Audio is captured and sent in a 6.8 kbps stream to leave room for the graphics. Lecturers have added some small animations; it takes about

1 hour to create 10 seconds of animation. Materials are available online within 24 hours after a class session. The system can handle up to 25 users per class; 5- to 10-second delays are considered acceptable.

The synchronous system uses Microsoft NetMeeting or ICQ for interactivity. It takes advantage of the chat feature alongside PowerPoint slides for interaction between instructors and students, but little interactivity has been observed. In a large class, the professor might have an assistant answer the chat questions. Problem areas include scheduling, preparing auxiliary materials, marketing, and meeting assessment deadlines. Presenters must act as moderators and must learn to teach to an empty room or to a screen audience and a live audience at the same time.

Next Steps: Switching to the Internet

The Internet, especially the Next Generation Internet (NGI), would offer increased bandwidth and networking capabilities that would in turn enable a range of scenarios for telemedicine. One scenario might be physician-free medical practices in which nurses and physicians' assistants handle many patient complaints and consult with specialists only when needed. In another scenario, centers such as the one at ECU could become brokers between practices and specialists rather than hubs that provide the specialists. The center at ECU also could become a means of providing services to more affluent clients who opt to pay for their own health care rather than relying on their health plan. Real-time treatment planning, with its increased bandwidth requirements, is seen as unlikely. In the future, Internet Protocol (IP) video and Web TV will be close to providing telepresence.

ECU is not currently slated to receive a connection to the NGI, but researchers are trying to find a way to get a connection. In the meantime, telemedicine system developers have been thinking through the challenges inherent in moving their telemedicine network to an all-IP environment, which would provide more of a plug-and-play capability. Mr. Balch has been working with the American Telemedicine Association to sort out which disciplines are most likely to be effective using telemedicine and the bandwidth and throughput required for each. Table A.3 summarizes his conclusions, which divide telemedicine consultations into three broad categories: (1) those in which high-resolution, static images are needed (such as in pathology, dermatology, or radiology), (2) those in which medium-resolution imagery is needed, with little motion (such as in psychiatry or internal medicine); and (3) those in which medium-resolution images are needed, but with a high degree of motion (such as in cardiology).

ECU officials would consider switching from leased lines to the

TABLE A.3 Bandwidth Required for Different Types of Telemedicine Applications

Telemedicine Application	Needed Bandwidth
High resolution, no motion	Store and forward
Medium resolution, low motion	128 kbps
Medium resolution, high motion	384 kbps

SOURCE: David Balch, director of Telemedicine at ECU, site visit by committee on February 2, 1999.

Internet if there were evidence that it would reduce costs and lead to increased access, which it currently would not in rural eastern North Carolina. They would like to outsource the networking part of the business to a company experienced in dealing with telephone companies (i.e., an ISP). They are moving to IP-based video, in spite of the cost and access issues, to overcome the problem of heterogeneous systems. Beyond these issues, use of the Internet would require attention to a variety of administrative and technological challenges.

One challenge is rapid change in both health care models and technology. Modes of health care delivery and payment are in flux, and provider plans are becoming more integrated and national in scope. This is not currently a problem for ECU because it serves only North Carolina, so it is able to use physicians' assistants on the remote end and a physician on the central end. But working across state lines would require a physician on each end and perhaps some type of cross-licensing arrangement. Meanwhile, as noted earlier, rapid changes in technology have left ECU with a system that is cobbled together rather than designed for optimum overall functioning.

A variety of standards and quality control measures would have to be instituted. A technical service broker would need to document the sessions, keep records, and keep track of the time. For the system to be efficient, participants would need to use standard history-taking procedures, standard clinical protocols, and standard tools for teleconsultation. Triage and quality control would be needed on the store-and-forward systems. A certification system would also have to be created for consultants, presenters, and, perhaps, for overall telemedicine programs.

Security also would need to be improved. Security is an issue for both data transmission and medical records, which must be kept secure from unauthorized viewers, including technicians and hospital employees. Currently, ECU faxes medical records and uses Proshare, which provides encrypted mail, for the store-and-forward system. Each mes-

sage is checked to be certain that all the parts are there, but this model will not scale up. The NCREN fiber network cannot be used for telemedicine because it passes through the telephone company, creating too many security issues.

Quality of service (QOS) is not an issue for ECU today because the program uses leased lines; service-level agreements (pacts between ISPs and users on minimum bandwidth, maximum latency, etc.) to provide equivalent QOS would be needed if the Internet were used.

UNIVERSITY OF NORTH CAROLINA, CHAPEL HILL

The site visit to the University of North Carolina (UNC) in Chapel Hill on February 3, 1999, featured two main activities: a visit to the computer graphics laboratory for demonstrations of telepresence systems with potential applications to medicine and a meeting to discuss Internet-based education programs in the School of Medicine.

UNC's Computer Graphics Lab

Guided by Henry Fuchs, the committee saw four demonstrations of computer graphics and virtual reality systems with potential applications to medicine: the Telepresent Office, augmented-reality ultrasound, wide-area tracking/walking, and image-based rendering/depth extraction. All of these systems attempt to convey telepresence.

Telepresent Office/Office of the Future

The Office of the Future is an attempt to create an immersive environment that enhances the sense of participation in teleconferences by using multiple cameras and multichannel audio. The remote participant sits at a corner desk, and a 270-degree image of the central conference facility (or operating room) is projected on the walls of the room. Real-time video and audio are transmitted to the remote location. The setup requires about a dozen live video streams, which require substantial bandwidth (although bandwidth demands might be reduced). Persons at the central location see the remote participant on a video monitor and receive an audio feed.

Much attention has been focused on image registration by cameras and projectors and on the mapping of images onto room geometries that differ between sending and receiving locations. Flight simulator data show that occlusions and breaks must be kept at 50 msec to avoid interfering with a person's conscious attention to the task (100-msec breaks create a distraction). The 50-msec figure is used as a point of reference in build-

ing smoothness into the collaborative environment. The system has been used with multiple IP video streams over the vBNS. Empirical observations suggest that it does not work very well in spite of apparently adequate bandwidth; there is not much emphasis, at present, on the efficient network transport of video and audio.

Latency and time-stamping of different video streams are important, especially during the current transition from a two-dimensional to a three-dimensional system. A true sense of presence can be provided only if there is no latency. The system needs to factor in the distance between viewer and viewed and provide a 360-degree view without breaks. In the demonstration, two or three cameras were added to compensate for distance, and two or three more to ensure a 360-degree view. Then, all the video streams need to be delivered together and reconstructed into a seamless image. Perfect time-stamping, down to microseconds, is needed for reconstruction. The time stamp could be a property of either the network or the packet. There are problems with either approach: a time stamp in the packet helps with synchronization, for example, but it creates a latency problem.

Augmented-Reality Ultrasound Visualization

The augmented-reality ultrasound system consists of a head-mounted display that provides an “X-ray vision” view of a surgical patient’s abdomen. The user can visualize hidden structures reconstructed from CT images and/or an ultrasound probe. The system is able to convert sequential two-dimensional ultrasound into three-dimensional objects within the abdominal cavity. It was developed in collaboration with a local UNC surgeon who specializes in minimally invasive procedures. It is not envisioned as a network application.

Wide-Area Tracking

The wide-area tracking system uses a three-dimensional helmet display of a virtual architectural space. Guided by position sensors in the helmet, the system renders appropriate changes in perspective as the person wearing the headgear walks through a room. It is supplemented by tactile feedback from Styrofoam objects that correspond to three-dimensional images. The system is tied to high-speed position detection and three-dimensional rendering of local objects, but it could be converted to a network application.

Image-Based Rendering/Depth Extraction

The image-based rendering/depth extraction technology is a three-dimensional modeling system in which architectural surfaces are scanned by a laser to acquire depth information, upon which surface information from two-dimensional images is mapped. The system produces photo-realistic rendering on a computer display (requiring tens of millions of image primitives per second) with interactive joystick control. Users can move around the scene. It runs on a pixel planes computer and is bound to a central processing unit and local memory. The system is not envisioned as a network application.

Educational Technologies

John Loonsk, head of information systems for the School of Medicine and the Division of Medical Informatics in the Department of Biomedical Engineering, discussed the UNC approach to technology-based curriculum support. The goal is to create an Internet-based learning environment for medical education for both resident and nonresident students. The system incorporates standard off-the-shelf software and a browser-based interface. It uses flat HTML for all pages, with URL interlinking for navigation. A staff of 45 handles all of the work; few faculty members are involved directly in creating materials.

All medical students are required to purchase a specific laptop computer, which has standard software and preconfigured ISP accounts for remote access. The staff transfers course notes from the live network environment so that students can download the previous, current, and following week's syllabus materials and use them off-network. A recent study of system use found that 136,000 pages were retrieved and videos were played 4,300 times in a month (Box A.1). Students also are required to take a medical informatics course delivered online. Before each class session, the students receive an e-mail message with HTML links to basic materials for that topic. Then a lab session is held in the five rooms in which student carrels are located; each student has a carrel with network access and power.

The preclinical version of the educational system has five components. One is reference materials based on the university's UNCLE system, which includes links to (1) MEDLINE; (2) products from Ovid Technologies that contain links to 60 full-text journals and textbooks; and (3) some UNC-specific links to content. A selection/editorial committee, with members from various UNC schools and run by the library, chooses and deselects materials.

BOX A.1

Use of UNC's Web-based Educational Suite

Student use of the University of North Carolina (UNC) Web-based educational environment is very good and growing, according to data for September 1998. Usage increased substantially over the previous year, as first-year students appeared to use electronic resources much more often than any previous class. Actual use is undoubtedly even higher than the figures suggest because UNC enables students to download the most frequently used material (e.g., histology images, anatomy images, weekly syllabi), and downloaded material is not tracked. The September 1998 data are as follows:

- Roughly 136,000 electronic syllabus pages were retrieved (if typical Web hits were counted, then this figure would be about five times higher).
- Computer-assisted instruction (CAI) and reference programs (separate from UNCLE) launched from the Web were run 6,000 times.
 - Digital videos and audio files were run 4,300 times.
 - 1,300 searches were conducted.
 - UNCLE sessions numbered 33,400 (each session represents multiple pages).

The users were varied. For electronic Web pages, 74 percent of retrievals were from student laptops and the remainder were from "public" student workstations and faculty/staff machines. Almost all of the videos were run from student laptops. For CAI and reference programs, about 35 percent of access was from student laptops and the remainder was from public or faculty/staff machines.

SOURCE: John Loonsk, University of North Carolina, site visit by committee on February 3, 1999.

A second component is teaching materials, beginning with the syllabus, which is distributed to students in both an online and a printed version. Traditional instructional materials (e.g., lecture notes and slide images) are captured and integrated with common sources of published information to support reference/retrieval and problem solving. Search engines allow students to find material easily. Typically, course materials are provided two weeks in advance and left on the system until updated or replaced. Because the full 2 years of the preclinical syllabus are online, the School of Medicine has abandoned its curriculum management database; faculty members simply search the syllabus using keywords.

The third component is an image repository—a collection of medical images with some text descriptors and titles. Staff are trying to convert these into flat HTML format. A directory structure exists, but there is no

database structure for generating dynamic HTML pages. Staff are already building an image database and case repository using a standard format; they hope faculty will use the repositories to build instructional materials with reusable components. For legacy applications (PC-based, computer-aided instruction materials), a launcher was built that can be reached from the browser.

The fourth component is a case repository, which contains specific problems that provide an “evocative” presentation of a particular case. Different classes may use the same cases. Most cases are text-based; there are no simulations yet.

The fifth component is communications technology, consisting of a standard e-mail client. Dr. Loonsk wants to make more use of e-mail with embedded HTML and store-and-forward capabilities. There was an unsuccessful attempt to provide threaded discussion lists; the students complained of too many distracting messages.

The clinical learning environment was developed for third-year students. It contains quick references on UNCLE as well as access to full-text articles; clinical support tools so that students can document patient encounters (e.g., problem lists); and learning frames, which provide background and instruction on particular clinical problems. A learning frame is developed for each topic on a problem list that students use as they complete clerkships. The frame keep the basic didactic information up to date and also includes canned searches against basic references and a number of other custom views of the curriculum resources. Access to the clinical database, a noncommercial application, is provided to both local and remote students.

There are several constraints on the system. One is bandwidth—the so-called last mile connectivity to students’ homes. There is no digital subscriber line (a digital telecommunications protocol for sending data at high rates over copper telephone lines) or cable modem access in the area, and there are many different phone companies. If bandwidth were much higher and universally available, then UNC would use teleconferencing for student mentoring activities in which students start as a group on campus and then disperse to community sites. Another constraint is security, especially for clinical data. The hospital uses audit trails to deter improper access to medical records. A related problem is the absence of unique and secure IDs for user authentication. On the other hand, the use of IP authentication by vendor products limits off-campus access to some reference materials, a growing problem because 50 percent of a student’s clinical time is spent outside the hospital.

Other constraints include the demands of managing virtual private network (VPN) and extranet services; administrative structures are needed to establish VPNs among sites and to distribute and revoke encryption

keys as needed. There is also inadequate QOS to enable the provision of video and public-network-based telemedicine to academic health education centers. Users cannot be assured of getting the requisite 128 kbps or more consistently across the Internet without some sort of service-level agreement. There is also a need for higher level standards (e.g., XML) to supplement URLs. Finally, there is some faculty inertia in moving toward the use of new communications methods (e.g., newsgroups, listservs).

UNIVERSITY OF WASHINGTON

The visit to the University of Washington (UW) on February 10, 1999, consisted of a series of briefings on, and demonstrations of, a range of projects related to medical informatics and human-computer interfaces, a key factor in making the NGI more accessible to health care. Since the 1970s, UW has engaged in a series of computing, informatics, and Internet technology projects that have built upon one another. The site visit team heard about many of these projects as well as about changes in the Pacific Northwest brought about by information technology.

Computing Activities at the University of Washington and in the Seattle Area

Ed Lazowska, chair of UW's Computer Science Department and a member of the Computer Science and Telecommunications Board, provided an overview of computing and Internet-related activities at UW and in the Seattle area generally.

Over the past few decades, Seattle has been transformed from a lumber town into a community dominated by Boeing Corporation and then into a much more diversified city. The UW Medical Center has played a key role in that transformation, helping to spur the development of the local biomedical electronics, biotechnology, and software industries. Washington State now has the largest concentration of high-tech employees in the nation (i.e., workers employed in companies with higher than average levels of research and development). Although the aerospace industry has leveled off in both dollar and employment terms, other high-tech areas are growing quickly.

A recent survey by the Washington Software Alliance found that software is a \$20 billion industry in the state, with 2,500 firms employing 47,000 permanent workers. Employment has grown by a factor of four over the last decade, and the software industry has an employment multiplier of 5.5 (meaning that each software job creates 5.5 additional jobs in other sectors), which is twice as high as that of the technology sector as a whole. The average annual wage in Washington's packaged-software

industry in 1997 was \$198,000 (salary plus exercised options), not including the pay of 640 company officers. The average salary before options was \$67,000. The software industry has 7,300 current vacancies and anticipates 64,000 new hires over the next 3 years. If these jobs can be filled, then the industry will generate an additional \$12.8 billion in revenues over the 3-year period. Seventy percent of these jobs require a bachelor's degree or higher. These data have implications for local universities; namely, they find it difficult to recruit top people for academic jobs, and there is limited practical value in 2-year academic programs, whose graduates fail to meet the employment criteria of much of local industry.

The ARPANET was brought to the Pacific Northwest in 1979 by the UW Computer Science Department. The NWNnet (part of the NSFNet) later formed around UW. As of August 1997, the vBNS was not designed to extend to the Pacific Northwest, even though there were somewhat redundant points of presence (POPs) in regions with supercomputer centers. (A POP is a dial-in site where a backbone network connects to access networks and where Internet service providers house switching hardware and transmission equipment.) Now, the region has a vBNS connection to Denver and San Francisco. Plans call for the development of regional collaborations, which are key to developing regional expertise and will allow for regional peering and traffic aggregation through a single point of connectivity. UW was one of the four original sites for Abilene and worked with Qwest, which is laying the fiber for Abilene as it lays its own national network, to get hooked up early, so the university now has a 2.5 Gbps fiber-optic connection. The Pacific Northwest gigaPOP, which includes Abilene and vBNS, is operated by UW, which received money from the state legislature for equipment. (A gigaPOP is a point of presence for accessing high-speed networks, sometimes called gigabit networks.)

A 10 Gbps link connects UW to the Westin Building, where the gigaPOP is located, in downtown Seattle. UW is working with Microsoft and U.S. West on a regional bandwidth experimentation program. UW also has requested National Science Foundation (NSF) funds for connections to regional universities, including Oregon Health Sciences and Oregon State. The remaining challenge is to convince agencies other than the NSF to connect to the gigaPOP. Private networks (such as those run by the departments of Defense and Energy and NASA) tend to serve the same geographic areas, often missing the Pacific Northwest. If agencies carried each other's traffic, then aggregate demand could be considered in determining the need for network/POP connections. However, shared networks might pose availability and security problems.

Informatics and Internet-related Activities

Brent Stewart provided a brief overview of the School of Medicine's informatics and Internet-related activities. A number of early projects were not discussed in detail during the site visit. For instance, UW was one of the first sites in the Integrated Advanced Information Management Systems program of the National Library of Medicine (NLM). It also participated in the Advanced Communications Technology Satellite program, a NASA effort that started in 1972 with a satellite providing a connection to Alaska. A project on ultrasound telemedicine, sponsored by the Defense Advanced Research Projects Agency's Technology Reinvestment Program, was designed to develop dual-use ultrasound telemedicine technologies that could be used in ambulances in both battle-field and civilian situations.

Other efforts included the Bench-to-Bedside project, which extended the reach of the UW School of Medicine's Internet resources to community hospitals and libraries. It encompassed research, deployment, and testing activities. Bench to Bedside and Beyond (B3): Building and Testing a Regional Telemedicine Testbed was sponsored by the NLM; it expanded the original program to allow the sharing of clinical information among participants in the Washington, Wyoming, Alaska, Montana, and Idaho (WWAMI) Rural Telemedicine Network. It includes a secure Web interface to medical records, secure clinical e-mail, and access to medical library resources.

Projects that were discussed in some detail at the site visit are summarized in the following subsections. They include the distributed radiology oncology network, the MINDSCAPE interface to a clinical data repository, the WWAMI network, medical projects at the Human Interface Technology laboratory, the gigaPOP, NGI projects, Biomedical Library programs, and the Digital Anatomist electronic repository of anatomical images and teaching programs.

Distributed Radiation Oncology Network

Ira Kalet, of the Radiation Oncology Division, discussed radiation planning software that he helped develop. It uses a client-server architecture to enable collaborative planning of radiation treatments by physicians in different locations. The software creates detailed three-dimensional visualizations of cancerous regions of the body by building up sequences of two-dimensional CT images. The images can then be sent across the network to be viewed and manipulated by the collaborating physicians. Radiation oncologists can use the system to identify tumors and plan treatments in collaboration with dosimetrists. The collaborators can view

various radiation portals and beam trajectories to aid in planning treatments; the software synchronizes the images shown on each user's screen. Such collaboration formerly was possible only by faxing images to and from collaborating physicians. In its current configuration, the system allows sharing of images across a LAN. In the future, it will allow true remote collaboration.

System development poses a number of technical challenges. The difficult programming task, carried out in collaboration with the Computer Science Department, resulted in several papers published in computer science journals. Latency is a significant concern, as it currently ranges from 1 to 2 seconds and could increase across a wider area network. Obtaining funding for continuing research on such projects is difficult. While health organizations such as the National Cancer Institute understand the need for research on ways to calculate radiation doses, for example, they generally fail to see why they should fund research on collaboration software or new software design tools, according to Dr. Kalet. Industry resists supporting work like this, too, viewing it as too risky and as taking too long to produce results that can be commercialized.

MINDSCAPE

Tom Martin, director of systems development for UW Medical Center Information Systems, demonstrated MINDSCAPE, a Web-based interface for viewing a clinical data repository. MINDSCAPE is based on the Medical Information Networked Database (MIND) data repository developed at the medical center between 1991 and 1994. The system contains about 60 gigabytes of data, including clinical data (e.g., electronic medical records), and offers links to library reference materials (e.g., drug databases, MEDLINE, laboratory tests reference data, and clinical guidelines) as well as other decision support tools. MINDSCAPE generates reminders about exams that patients will need soon, lists of medications, and dynamic reports for several measures, such as hemoglobin levels in diabetic patients. Other reports also can be generated on a clinic- or provider-specific level. For example, the system can generate information about compliance rates of patients in a particular clinic or under a particular physician's care. Users can pull up patient records so they can call and remind them to come in for appointments.

Terminals and/or computers with access to MINDSCAPE are located in all exam rooms at the medical center. The system is also accessible through dial-up modems. Because it is primarily text-based, MINDSCAPE does not currently generate unusual requirements for bandwidth; however, images are being added to records, and that will increase bandwidth

requirements. MINDSCAPE uses commercially available Web-server encryption technologies (128-bit SSL) and server authentication for security, along with access controls. Access to system information is granted on a need-to-know basis, and overrides are in place to allow access in emergencies. All access can be audited to ensure compliance with confidentiality policies. The confidentiality and security policies are derived from policies in place for paper records. All physicians and other staff members with access to the system must sign a confidentiality agreement and complete training on confidentiality policies.

WWAMI Rural Telemedicine Network

UW serves as the tertiary care center for the five member states of the WWAMI (Washington, Wyoming, Alaska, Montana, Idaho) Regional Medical Program and is the regional medical school for those states. WWAMI links the UW School of Medicine, the UW Medical Center, the Harborview Medical Center, and the Children's Hospital sites and clinical teaching sites throughout WWAMI. The participating states—which together cover about one-fourth of the U.S. land mass but have sparse populations—created an affiliated education and care program. Students complete their first year of medical school in their home states and then come to Washington for their second year. The third year is decentralized rotations; about half of these students go out into the field. The residency program places students in member states.

WWAMI has undertaken many telemedicine activities since the 1970s, especially to Alaska, to support its education programs. Funds from the Rural Health Policy Agency allowed the consortium to set up six telemedicine consultation sites in small communities. The local primary care physicians are also preceptors for medical students, and an average of one or two consultations, mainly specialty consultations, are held at each remote site every month. Psychiatry, cardiology, and dermatology are the most common telemedicine specialties, although there is some use of the system with peripheral instruments and for trauma consults. Telemedicine has not worked for rheumatology. The network also is used as an administrative link among the UW sites.

The system uses a frame-relay setup with 56-kb switched lines to the rural sites, which lack either the capacity or funding to get digital lines. Communications are imperfect at this speed and produce video artifacts, for example. The local Seattle locations have T1 linkages. For tele-dermatology, the participants use store-and-forward capabilities; other consultations are synchronous. The telephone lines cost, on average, \$3,000 to \$5,000 monthly, which could not be covered without grant money. The system uses PictureTel equipment.

The local physicians are satisfied with the program; they use the library's resources and learn to like the Internet. But several issues must be resolved before the telemedicine program can be expanded. First, the program crosses state lines and requires licensed physicians at either end, so beginning students cannot present patients. When students do clinical rotations in rural areas (as half of the students must do), they must get licensed in both states—unless they are participating in telemedicine programs operated by the federal government (such as the Veterans Administration), which are exempt from state licensure requirements. Second, insurance is an issue; some patients will not use this medium for consultations because their insurance will not cover it. Third, telemedicine can affect local referral patterns. Finally, reimbursement is a major problem. Consulting doctors are currently paid from the grant money. The HCFA and the Rural Health Policy Agency support only 80 percent of a normal payment, and the doctor has to split the fee with the referrer. Montana approved telemedicine for Medicaid because it saves money normally spent on patient travel.

Medical Projects at the Human-Interface Technology Lab

The Human-Interface Technology (HIT) laboratory is a research unit in the UW College of Engineering. It has a roster of 108 people; 18 are regular staff members and the rest are made up of faculty associates, visiting scholars, graduate students, and so on. The income of the laboratory is about \$17 million, two-thirds of which comes from grants and contracts. Forty companies participate as members of the Virtual Worlds Consortium, providing a little less than one-sixth of the revenue. Spin-offs from the laboratory include 13 different companies; 10 of these are still active, including 3 created recently.

Suzanne Weghorst, assistant director, introduced the HIT laboratory, which has a goal of developing and demonstrating mission-transferable technology. Tom Furness, the director, elaborated on the laboratory's history. He worked for many years at Wright-Patterson Air Force Base, designing human-machine interfaces in airplane cockpits. He came to Seattle in 1989 to found the laboratory, by which he hoped to extend the scope of his work, with the general goal of providing better coupling of humans to advanced machines.

Furness emphasized his view of the future as technology that simulates "being there," improving humans' ability to transport themselves by moving their eyes to different places and times. Such experiences range from teleconferencing to "transport" through an endoscope to view hidden portions of the body. For instance, a videotape made 5 years ago showed an interactive teleconference in which participants in the United

States and Japan wore virtual reality (VR) helmets and cooperated in the task of herding virtual creatures across a conference table with paddles. The telecommunications link consisted of four ISDN lines. The images displayed were somewhat flat and cartoonish, but the interaction was successful.

The committee was shown three active laboratory efforts. In the virtual operating room, the participant wears a VR helmet and holds a control stick. The helmet's location and orientation are sensed by a device mounted on a fixed stand over the space, and this information is used to drive the displays to the video helmet (much in the manner of current VR games). The environment is based on photographs and equipment from Harborview Medical Center. There is a patient on an operating table; by using the control stick, the participant can manipulate displays, such as the electrocardiogram output, and instruments, such as an endoscope inserted in the patient's lower abdomen. The displays can be controlled so that they are visible regardless of the participant's perspective, or they can be fixed at various positions in the virtual room. There were some noticeable lags in the following speed of the display if the participant turned quickly, and the overall precision of the location sensors appeared to be on the order of inches rather than millimeters.

The second active effort involved simulation of surgical suturing through a computer-controlled force-feedback device. The user interacts with the environment through a pair of scissors holding a virtual needle. A standard video monitor displays the position of the needle relative to a wound in a small area of skin. A finite-element model of the skin—it has about 200 nodes, with a relatively higher concentration of nodes near the wound—simulates the restoring force of the skin against the needle and controls the force feedback. If the user inserts the virtual needle into the skin orthogonally to the skin surface, for example, then little force is felt, but if he or she holds it at an oblique angle, then considerable force is needed to pierce the skin. The force-feedback device requires updates about 1,000 times per second; the visual display runs at a standard 30-cycles-per-second refresh rate.

The third effort is the Virtual Retinal Display. The idea is to paint an image directly on the retina with photons instead of projecting the display elsewhere and requiring the person to follow it visually. The image appears only on the participant's retina. The lab bench setup included low-power red, green, and blue laser light sources fed through an optical fiber, with the fiber's output scanned mechanically over the retina by a moving mirror. Although the lab setup seemed cumbersome, a company called Microvision was spun off in 1993 to commercialize this technology and evidently has had some success in reducing it to a practical size and weight for portable use. Because the image can be focused so that it

passes through only a small part of the user's cornea and can potentially be focused on a specific part of the retina, the technology holds promise for assisting persons with impaired vision and providing bright, high-precision displays for VR applications. A number of surprising results have been found in experimental use. For instance, users do not perceive flicker in the static images even at relatively low refresh rates (e.g., 15 frames per second), meaning that the system offers twice the resolution of other displays at the same bandwidth.

Seattle/Pacific Northwest GigaPOP

Jim Corbato discussed the evolution, current architecture, and expected future evolution of high-bandwidth IP network infrastructures for UW and its Seattle-area health care partners, which include the Harborview Medical Center, Children's Hospital, and the Fred Hutchinson Cancer Research Center. Network interconnections in Seattle are simplified by the physical proximity of data networks from various interexchange carriers (e.g., Qwest and US West) and ISPs at the gigaPOP facility in downtown Seattle. The network connections have relatively poor site security (because they are in a general-use office building) and present a potential single-point-of-failure for Internet access for the entire Pacific Northwest.

Next Generation Internet Projects

UW received a phase 1 award from the NLM to examine biomedical applications that would benefit from the NGI and has since received a phase 2 award from NLM to further this work. This project is titled Patient-Centric Tools for Regional Collaborative Cancer Care Using the NGI. Brent Stewart outlined the project, in which a high-performance metropolitan area network is being designed to transmit clinical data, including radiology images (with a capacity to deliver eight simultaneous 10-MB image sets simultaneously, using a 622-Mb channel), to a planned cancer care facility on the south shore of Lake Union. The center is scheduled to be completed within a year. The bandwidth requirement is based on fully digital radiology, with interpretation of images by radiologists at UW and a digital archive at UW rather than at the clinic site.

The development of this center is based on three hypotheses: that health care is becoming highly distributed and differentiated; that health care is operating in a resource-limited environment; and that the NGI will enable more collaborative practice, regardless of where patients are located at a given time. The NGI will enable the formation of the cancer care alliance; facilitate teaching and research; enable a fully integrated

team approach to diagnosis, treatment, and management of cases; and accelerate the discovery and dissemination of knowledge.

The underlying technology will consist of the local gigaPOP as well as a virtual, enterprise-wide multimedia electronic medical record based on MINDSCAPE. There will be a backup line in place, perhaps a leased DS-3 line. The system will transmit real-time video (e.g., ultrasound, fluoroscope, and synchronous telemedicine consultations), store-and-forward video, and interactive radiation oncology treatment planning (e.g., graphics, images, and video). Additional multimedia knowledge resources, such as the Digital Anatomist and streaming video for patient education, also will be available.

Technical requirements are based largely on the needs of remote radiological image archiving and display. To allow the simultaneous downloading of eight different 10-MB images within one second, the system needs bandwidth of 640 Mbps. The UW and Harborview radiology departments are all digital now, but they use computed radiography (in which an imaging plate is scanned by a laser) rather than flat-panel digital. The centers will become fully digitized once the technology comes down in price. Stewart envisions that a radiologist covering at a remote site might use the system to perform work that he or she would have done at the home site, downloading images remotely.

There are no formal plans for evaluating the technical infrastructure. Because the Internet today provides no QOS, UW will put in place as much bandwidth as possible and use whatever service level results. Dr. Stewart viewed medicine's demand for bandwidth as similar to that of other industries; he compared multisite telemedical collaboration to automobile companies linking together their remote research and development sites.

Biomedical Library Services

The UW Health Sciences Library is moving all of its services to Web-based delivery (see <healthlinks.washington.edu>). Geographic issues related to supporting the WWAMI program (i.e., mountains, small towns, long distances between towns) make this transformation necessary. In addition to the WWAMI medical education program, the pharmacy, nursing, public health, and social work programs have distance education programs. Faculty want to deliver digital video to their WWAMI-based students and provide them with access to course materials and the clinical digital library resources. This program will require high bandwidth and many servers. At present, the schools use scanned PDF format to deliver interlibrary loan and course materials (e.g., course notes and

reserve materials) to distant students. The documents are an estimated 300 kb in size each.

The Health Sciences Libraries offers over 1,400 full-text online electronic journals to UW faculty, staff, and students. A major issue is compliance with licensing agreements. UW librarians negotiate aggressively for licenses allowing digital materials to be available to faculty, staff, and students at any location. They control access with user IDs and passwords and UW IDs and are increasingly making materials available through a proxy server. Access is also a problem, because not all materials are locally loaded, and network latency for materials stored at remote sites (NLM, journal publishers, etc.) is an issue. The plan is to support nomadic computing because students, faculty, and residents move around constantly. Network latency is a serious problem for accessing remotely stored full-text journals and other resources.

UW was part of the NLM test on access time for PubMed over the Internet that was published recently in the *Journal of the American Medical Informatics Association*. NLM has tried to track the latency problem and 3 years ago performed a minor test for the Utah link for the online journals (because it intended to provide all resources remotely). Users see real degradation of performance from 11 a.m. to 3 p.m. Pacific Standard Time, but that latency is due in part to issues within the UW network (in other words, the latency formerly observed in NLM's MEDLINE for that time period now is observed on the Internet.) However, the latency depends on the connection. UW is upgrading the network to 10 Mbps Ethernet (10Base-T) to improve throughput; moving to 100 Mbps Ethernet (100Base-T) on every public library machine is a much more costly venture (about \$400 per station). The UW Health Sciences Library has many public stations, including about 150 in the microcomputer lab it manages for the Health Sciences Center and over 100 public workstations in the three Health Sciences Library sites. Remote users may have trouble with commercial connections to library online resources (e.g., through MSN or AOL), which can impose latency problems during certain time periods.

Digital Anatomist

The Digital Anatomist is an NLM-sponsored project to develop an electronic repository of anatomical images. Anatomy is, of course, fundamental to health sciences education, and it provides a framework for organizing other biomedical information. Jim Brinkley presented the work of the UW Structural Informatics Group. The group works in three areas: representations of structural anatomical information, from the level of individual cells to gross anatomy; methods for accessing and using structural information; and practical applications of their tools for

research, education, and clinical work. It tries to exploit opportunities for online systems dealing with anatomical information. The key data structure for its work is an ontology of anatomy, developed by Cornelius Rosse, that serves as a common data structure for most of the applications. They call this the foundational model of anatomy.

The system demonstrated for the committee provides authoring tools using both symbolic information (e.g., names, semantics, structures) and spatial images (typically three-dimensional images) of anatomical structures. It consists of a symbolic information database and a separate three-dimensional image database accessed through a single server. A number of intelligent agents have been developed to assist in retrieving and assembling data sets and images. The agents have knowledge of both the information available on the system and the user's level of sophistication. In response to the command "Show me the structures of the left lung," for example, the system will check the symbolic database to find out what structures are in the left lung, then go to the image database to determine what images are available, and then use a scene-generator to assemble the pieces properly. The user then can highlight particular elements of interest, rotate or zoom in on the image, and remove objects that block the view of other objects of interest. All processing is done on the server.

On the authoring side, the system contains a knowledge-builder for adding information to the symbolic database. It is a relational database containing 25,000 terms describing all structures with dimensions of 1 mm or greater in a particular set of organisms, including humans. For the spatial database, the system can create volumetric, three-dimensional models of anatomical structures from two-dimensional images. It also can create and perform animations. A brain image demonstrated for the committee superimposed vascular structures onto the brain and contained some 100,000 polygons. Images can be annotated for clinical and educational purposes.

On the user side, the system contains an annotated image server. Users can call up images and click on individual structures within the image. The system outlines the selected structure and generates its name. This system is used in anatomy education, but the images are so large and the networks so slow that students tend to use a CD-ROM rather than access the database through either campus or remote networks. The system's quiz mode can test students' knowledge of anatomy. A tutorial system can embed images from the atlas into other documents. The Digital Anatomist and interactive atlas get an estimated 10,000 hits per day.

In a separate effort called the Brain Project, a digital neuroscientist system is being built that will overlay neurological images on images of the brain. The system will allow cutaways of volume-rendered images. These data are collected operatively by neurosurgeons who do real-time

stimulation and mapping of critical regions for various neurological functions, such as speech. NGI technology could streamline this process by enabling real-time superposition of neurological data on an open brain; the system then could sense the surgeon's probes and automatically maintain information about the location of the probe and the result. This would require capabilities similar to telesurgery, but the system would also be linked to databases for documentation and postmortem analysis.

The NGI offers several other opportunities for applying this technology. One is anatomical education, in the form of virtual dissections and intelligent scene generation. Another opportunity is brain mapping for either research (e.g., language mapping to identify correlations between brain structures and language skills/development) or clinical purposes, such as surgical planning. The technology also could provide structure-based visual access to biomedical information. Indeed, some have suggested that the ideal user interface to biomedical information resources is a model of human structure, which can serve as the organizing principle for this information.

REGENCE BLUESHIELD

The visit to the corporate offices of Regence BlueShield, the largest health care insurer in the Pacific Northwest with annual revenues of approximately \$2 billion, took place on February 11, 1999. The committee heard presentations on Internet-related activities within Regence as well as related activities in the Seattle area. The related activities include programs operated by the Foundation for Health Care Quality (FHCQ); the Washington State Department of Public Health; and the Community Health Information Technology Alliance (CHITA), working with a group called Agora.

Regence Web-based Services

Steve Moe, manager of electronic business practices for the Regence Group, presented a Web-based interface application called Network Data Express (NDEX) for determining beneficiary eligibility and making referrals. The Web-based system offers claim status inquiries, provider directories, reference materials (such as the formulary), e-mail, and managed care data and reports. It processes about 20,000 transactions per month (peak times are early in the day and during lunch), doing the work of two or three full-time employees who otherwise would give the same information out by phone (Regence processes millions of claims a month). Regence has deployed 1,500 workstations (including 800 intranet and 700 dial-up systems linked to a private Web server), of which about 30 per-

cent are in use on a regular basis. All users are assigned an ID and password and sign a confidentiality agreement. Patient and customer information is indexed by social security number.

Kirk Bailey, manager of security policy, stated unequivocally that the Internet is considered unsafe and will not be used for Regence's electronic commerce (e-commerce) transactions until it has sufficient security measures and functionality to meet the company's business requirements. The Internet raises concerns about security, privacy, and reliability. Other factors that have slowed the adoption and use of NDEX are a lack of content sponsors, especially among payers; the lack of Web browsers in many provider offices (they will get browsers during their next hardware upgrades, but Regence does not fund such upgrades); and user behavior (e.g., administrative workers are accustomed to using the phone instead of the computer to get information).

Foundation for Health Care Quality

Rick Rubin, president of the FHCQ, gave an overview of the foundation, a not-for-profit entity created in 1988 to meet the shared health information needs of the Seattle region. The foundation serves as a neutral meeting ground for providers, payers, plan purchasers, consumers, and others involved in health care. It participates in or sponsors programs in three areas. One area is e-commerce pilot projects, including a multistate effort funded by the Robert Wood Johnson Foundation to define eligibility and referrals, and CHITA, described in further detail below. The second area is performance measurements for health plans and providers, and the third is consumer affairs.

The FHCQ, which views itself as an economic development agency for the region, has learned a number of lessons about operating in the highly competitive health care marketplace. These lessons emphasize the importance of (1) enabling instead of mandating standards, because mandates may change willingness but do not affect capabilities; (2) making a business case that differentiates needs (i.e., things that stakeholders are willing to pay for) from wants (i.e., things they are not willing to pay for); (3) the Internet, which is widely viewed in the region as a plausible means of achieving long-held visions of seamless integration of information across organizations and which allows organizations to assume that networking capabilities will be in place so they can concentrate on higher order functionality; (4) information security and privacy, which can be either a barrier or an enabler, depending upon the circumstance; (5) the widespread sharing of expertise and information; (6) education as a means of facilitating the migration of information technology into health care, especially through efforts to reengineer the way organizations operate (a

process that can be more important than the technology itself); (7) working to refine national standards and develop implementation manuals; and (8) balancing competition and cooperation (firms can cooperate on some subsets of issues but not on others that are seen as having greater proprietary value).

Current or recent projects include an effort to standardize eligibility information, for which there is agreement on data items but not on presentation. Other regional projects are aimed at exchanging data on pediatric immunizations, referrals, claims, and lab transactions.

Community Health Information Technology Alliance Project with Agora

Peter B. Summerville, director of CHITA, and Kirk Bailey, manager of security policy for the Regence Group and founder of Agora, presented an overview of the Three-State Model Security Prototype. CHITA was chartered in 1997 and has 60 member organizations, including providers, payers, and state agencies. It is part of the FHCQ but has a separate board of directors. Agora, a local group interested in computer security, has about 450 members representing 120 Pacific Northwest region corporations. It was formed by chief information officers and security officers who became increasingly concerned about network vulnerabilities as their companies began to move online.

CHITA's early work focused on eligibility and referral transactions—negotiating agreements on data fields and standards to facilitate the electronic interchange of information. CHITA and the FHCQ worked with organizations in Massachusetts and Minnesota on a three-state project focusing on electronic security. The goals were to determine how electronic security could be implemented affordably and to develop a business case for a community-wide, secure infrastructure for electronic business. The group worked with Science Applications International Corporation (SAIC) to develop a security and risk management plan for business-to-business health information networks.

The plan identifies seven levels of increasing health care security. Together with Agora, CHITA is working to implement health security level 6 (HSL 6) within participating organizations. HSL 6 includes specifications for three network-based information services: authenticated, secure messaging; authenticated, secure file escrow and transfer; and authenticated, role-based access. The security model has been developed and published, and CHITA is in the process of identifying a bridge operator organization that will function as a trusted intermediary to oversee a prototype implementation, followed by a wider pilot project in the region. Issues to be addressed include the identification of a certificate authority,

which might be a nonprofit organization, the government, or a private corporation such as Verisign.

CHITA has no plans to attempt to change the Internet or its directions but rather will attempt to accommodate whatever weaknesses it exhibits with respect to information security. While asserting that businesses need to move too quickly to wait for the NGI, Mr. Bailey wondered if it would be possible to allocate part of the Internet 2 (perhaps one or two frequencies) for health care. He also would consider the formation of a separate health information network as a means of avoiding some of the security concerns associated with the Internet.

According to him, security officers in health care have responsibilities that differ from those of their counterparts in other industries. The applicable state and federal laws are different, the privacy and security concerns are greater, and health care organizations must meet requirements for successful electronic data interchange. At the same time, the health care industry is driven by economics, not privacy.

Washington State Laboratory Reporting Project

Jac Davies, representing the Washington State Department of Health, described the Electronic Laboratory Reporting System (ELBRS) project, which involves the electronic submission and tabulation of reportable events within the state, of which there are fewer than 100,000 every year. (Physicians and testing laboratories are required to report certain conditions to their county health department.)

Such reports generally are sent by regular mail, fax, or voice mail. Public health officials then are required to follow up with the doctor and patient to further investigate possible causes, paths of contagion, and so on. Often, reports are sent to the wrong county and/or are not subsequently forwarded to the state. Furthermore, different states and counties tend to have their own lists of reportable conditions, which are tied closely to local concerns (the conditions vary, for example, between urban and agricultural counties), and they have different rules for where to send the information. As laboratories (and health organizations generally) consolidate into national entities, tracking different reporting requirements has become time-consuming. SmithKline Beecham, for example, operates a number of clinical laboratories and has three or four people dedicated to tracking different reporting requirements.

Under Washington's planned system, lab reports would be sent directly to the state rather than to local health departments. The state then would process the reports and forward information down to local communities and up to the Centers for Disease Control and Prevention (CDC), as necessary. Such centralization would allow the state to better

track incidents across county lines. Planners hope that the system will encourage greater communication between the state and local communities or the CDC and that it will improve compliance with reporting requirements. Several issues have informed the planning for this proposed system. One is the use of the Internet, which is not only a logical choice but also the only viable option. Another issue is the sensitive nature of the data; there is, for example, a state requirement for reporting AIDS cases. A third issue is privacy, which is a major concern of the governor and residents of Washington.

A pilot program is under way with Group Health of Puget Sound. Labs encrypt their test reports and send them to the state health department's file transfer protocol server, which sits outside a firewall. State personnel move the file behind the firewall, check for errors, run it through an HL-7 formatter, put the data on an SQL database server, and send them to the county. They use a public key cryptography system (Pretty Good Privacy) described as minimal. There is no formal program in place for changing keys. According to a preliminary evaluation, the pilot program improved the completion and timeliness of reports. The time required to send information to the local health office improved modestly (to less than 1 day) and the time required to send information to the state improved by an average of 40 days (to about 1 day).

NOTE

1. Bernie H.K. Huang relocated to the Children's Hospital of Los Angeles and the University of Southern California as professor and director of Informatics effective January 1, 2000.

APPENDIX B

National Library of Medicine Awards to Demonstrate Health Applications of the Next Generation Internet

The National Library of Medicine (NLM) announced a new, three-phase program in 1998 to develop innovative medical projects that demonstrate the use of the capabilities of the Next Generation Internet (NGI), such as improved quality of service, security, network management, and support for nomadic computing. Phase I awards were announced on October 14, 1998, and included 24 contracts totaling \$2.3 million that were intended to improve understanding of ways the NGI can affect health care, health education, and health research systems in such areas as cost, quality, usability, efficacy, and security. Phase II awards were announced in late 1999 and consisted of 15 projects aimed at implementing capabilities in local testbed settings. Some of the Phase II awards build on projects begun under Phase I of the program, while others build on work originally conducted under other research programs. Summaries of each of the Phase I and Phase II projects announced to date are provided below. Additional information regarding these NGI awards and NLM's telemedicine evaluation program is available on the NLM home page at <http://www.nlm.nih.gov>.

PHASE I AWARDS

1. Pathology Image Database System

Yale University is planning a pathology image database system, Pathmaster, accessible via the World Wide Web. When a pathologist is con-

fronted with a slide containing a cell whose nature is uncertain, a digital image of the cell can be submitted to Pathmaster, along with certain clinical information about the specimen. Pathmaster will automatically compute descriptors and pass back images to the user, along with their cell types and diagnoses.

Contact: Perry L. Miller, M.D., Ph.D.
Yale School of Medicine
Center for Medical Informatics
333 Cedar Street
P.O. Box 208009
New Haven, CT 06520-8009
203-785-6753

2. Networked 3D Virtual Human Anatomy

The goal is to build a virtual human cadaver based on the Visible Human data set. An online virtual cadaver would be available to a wide range of students who could explore the virtual cadaver with a variety of tools. High-end applications will have a haptic interface.

Contact: Victor M. Spitzer, Ph.D.
University of Colorado Health Sciences Center
4200 East Ninth Avenue
Denver, CO 80262
303-274-0501

3. Rural Health Science Education

This project will develop a plan to evaluate the use of computer and interactive compressed video technologies to support rural health science education. It will enable delivery of interactive educational programming, such as grand rounds and continuing medical education, clinical information systems, library services, and consultation. Beneficiaries will be students, residents, and health care professionals.

Contact: Dr. Leo Bairnsfather, Ph.D.
Louisiana State University Medical Center
1501 Kings Highway
Shreveport, LA 71130-3932
318-675-6536
318-675-7757 fax

4. Biomedical Teleimmersion

By combining teleconferencing, telepresence, and virtual reality, teleimmersion enables teachers and students to interact with three-dimensional models. Teleimmersion combines several virtual reality systems with advanced network capabilities for learning especially in surgical education. NGI guarantees data privacy and security and will allow tele-immersive environments derived from models of patient data.

Contact: Jonathan C. Silverstein, M.D.
University of Illinois at Chicago
School of Biomedical and Health Information Services
1919 W. Taylor
Chicago, IL 60612-7249
312-996-5112
312-996-8342 fax

5. National Emergency Medicine Information Extranet

The National Emergency Information Infrastructure Consortium (EIIC) will create a plan for implementation of a secure National Emergency Medicine Information Extranet to improve emergency care across the nation. The primary application to be developed will enable interlinked standards-based emergency encounter registries, then feed back to providers just-in-time multimedia educational and treatment protocol services. The project will create an open architecture to enable other layered applications in the future.

Contact: Edward Barthell, M.D.
Infinity Healthcare, Inc.
1251 Glen Oaks Lane
Mequon, WI 53092
414-290-6700
414-290-6781 fax

6. Personal Internetworked Notary and Guardian

The Personal Internetworked Notary and Guardian (PING) project is designed to address the control of a personal record that can be integrated with more traditional sources of clinical information for patient use in the home, at work, and at school. In particular, PING is focused on (1) reconstitution of the patient longitudinal records from both provider-based information systems and portable, personal record systems, on the

Internet; (2) providing simple and secure authentication mechanisms; and (3) evaluation of the impact of PING upon the process of health care.

Contact: Isaac S. Kohane, M.D., Ph.D.
Director, Children's Hospital Informatics Program
300 Longwood Ave.
Enders 150
Boston, MA 02115
617-355-7821
617-730-0456 fax

7. Implementation to Serve Visible Human Datasets

This project plans to implement an NGI production system to interactively serve Visible Human data sets and anatomical data evaluation software. The image and knowledge data objects will be accessed by NGI-enabled World Wide Web users and evaluators. The system will provide to the user multi-resolution, anatomically labeled images within these Visible Human data sets as requested.

Contact: Brian D. Athey, Ph.D.
The University of Michigan Medical School
4771 Medical Science Building II
Department of Anatomy and Cell Biology
1335 Catherine St.
Ann Arbor, MI 48109-0616
734-763-6150
734-763-1166 fax

8. G-CPR and the NGI

The Louisiana State University (LSU) Medical Center proposes to implement a system of longitudinal electronic health records over the NGI that will integrate its ten public hospitals. This project is based on the G-CPR, or Government Computer Based Patient Record, a collaborative effort between the Department of Defense, Department of Veterans Affairs, Indian Health Service, and the LSU Medical Center. The objective of this project is to enable secure access and sharing of clinical information.

Contact: Richard Ferrans, M.D.
Louisiana State University
Medical Center Department of Public Health
1600 Canal Street, Suite 800

New Orleans, LA 70112
504-588-3507
504-588-3938 fax

9. Secure Radiologic Collaboration on the Next Generation Internet

The goal is to plan the implementation and deployment of a suite of collaborative medical applications to provide a secure, real-time, interactive environment for viewing, analyzing, and comparing radiological images in a clinical environment. This will provide clinicians and technologists the ability to share, in real time, diagnostic imagery and medical data.

Contact: Douglas L. Long, Sr., Principal Scientist
Odyssey Research Associates, Inc.
Cornell Business & Technology Park
33 Thornwood Dr., Suite 500
Ithaca, NY 14850-1250
607-257-1975
607-257-1972 fax

10. Open Architecture Multispecialty Data and Telemedicine Integration on the Next Generation Internet

The purpose of this project is to plan the implementation of a multi-specialty telemedicine testbed using NGI. The plan will identify existing and new multispecialty applications in patient care, continuing medical education, and patient education to be integrated into this platform. The planning activity is to be conducted by a team of scientists and clinicians from all pertinent parts of the proposing organization.

Contact: Joseph C. Kvedar, M.D.
Corporate Director Partners Telemedicine
1 Longfellow Place, Suite 216
P.O. Box 8941
Boston, MA 02114
617-726-4447
617-726-7530 fax

11. Patient-centric Healthcare Management over NGI

This project will demonstrate a patient-centric approach for health care management over the NGI. The demonstration will build upon the Elec-

tronic House Call system developed by Georgia Tech and the Medical College of Georgia to allow patients to videoconference with their health care providers and to monitor medical measurements over a secure network. A simple graphical user interface enables patients to control the system themselves. The system combines videoconferencing, vital signs measurements, patient education resources, and medical records, and enables patients to participate in their own health care.

Contact: Mr. John W. Peifer
Senior Research Scientist
Biomedical Interactive Technology Center
Georgia Institute of Technology
250 14th St., NW
Atlanta, GA 30332-0200
404-894-7028
404-894-7025 fax

12. Adopting the NGI as a Tool for Healthcare and Information Access

This project will assemble a team of medical informatics users and networking advisors to analyze biomedical and healthcare information processes and select those that best demonstrate the application of NGI technologies and tool sets, while simultaneously providing demonstrable benefit to healthcare practitioners and end users. Many information processes in health care clinical services, biomedical education, and research will be assessed. Once applications have been identified, the assessment team will select viable candidates, then formulate an implementation strategy for one application area.

Contact: Brent K. Stewart, Ph.D.
University of Washington Grant and Contract Services
3935 University Way NE
Seattle, WA 98195
206-616-1314
206-543-3495 fax

13. The Empathy Network: Improved Healthcare Delivery for Survivors of Mild Traumatic Brain Injury

The objective of the Empathy Network is to employ virtual reality (VR) technology, high performance computing centers, and NGI capabilities to dramatically improve the health care delivered to mild traumatic brain injury (MTBI) patients. VR technology will allow clinicians to construct a

virtual world that simulates the cognitive and perceptual deficits experienced by an MTBI patient. VR and NGI technologies will then enable a patient's other health care providers, family, friends and co-workers to experience the MTBI patient's problems in coping with everyday life. This will engender empathic insight, support, and understanding that are crucial elements of an MTBI patient's recovery and adaptation.

Contact: David L. Zeltzer
Sarnoff Corporation
201 Washington Road
Princeton, NJ 08540
609-734-2975
609-734-2662 fax

14. Remote, Real-Time Simulation for Teaching Human Anatomy and Surgery

This project plans to demonstrate remote, real-time teaching of human anatomy and surgery, using the NGI. A simulator architecture will be developed to deliver real-time simulation and visualization technologies to a diverse audience. The client component is a desktop PC or workstation. The simulation server receives sensor and control input from the client and transmits response streams. The NGI network-based architecture will allow for a heterogeneous mix of client configurations ranging from simple mouse and color displays to multiple high-resolution stereographic displays and haptic devices.

Contact: Parvati Dev, Ph.D.
Stanford University
School of Medicine SUMMIT
1215 Welch Road, Modular A
Stanford, CA 94305-5401
650-723-8087
650-498-4082 fax

15. Interactive Medical Data on Demand: A High-Performance Image-Based Warehouse Across Heterogeneous Environments

The goal of this project is to determine the requirements of a system for intuitive, real-time access to patient-specific data records based on multimodal images and multimedia. They will evaluate and select system architectures, software, and network configurations to provide access over different network bandwidths and platforms. This design will include

scalability of the system and extensibility to other health care applications.

Contact: Donald L. Stredney
Ohio State University
Research Foundation
Health Sciences Offices, B-030 Graves Hall
333 West Tenth Avenue
Columbus, OH 43210
614-292-9248
614-292-7168 fax

16. NGI-Aware, Scalable, Secure, and Adaptive Technology for Rural Telemedicine

The goal of this project is to develop a plan to demonstrate telemedicine applications that will utilize NGI infrastructure. Telemedicine scenarios include (1) nomadic clinics; (2) a public health station; and (3) a consulting health station in rural clinics and hospitals. These systems will be configured with a set of videoconferencing, diagnostic, and patient monitoring equipment.

Contact: Y.V. Ramana Reddy, Ph.D.
West Virginia University
Research Corporation
886 Chestnut Ridge Road
Morgantown, WV 26506
304-293-7226
304-293-7541 fax

17. Medical Nomadic Computing Applications for Patient Transport

The goal of this project is to transmit multimedia diagnostic information in real time from ambulances to receiving physicians using NGI technologies, thus enabling diagnostic and treatment opportunities during transport.

Contact: David M. Gagliano
TRW, Inc.
One Federal Systems Park Drive
Fairfax, VA 22033
703-345-7497

18. Distributed Revolutionary Medical Education Environment

The objective of this project is to develop a plan to implement and evaluate a distributed, medical education environment on a network testbed that simulates the characteristics of the NGI. These applications will be delivered across the spectrum of medical instruction, from undergraduate to postgraduate to continuing education.

Contact: Lael C. Gatewood, Ph.D.
University of Minnesota
Office of Research and Technology
1100 Washington Avenue So., Suite 201
Minneapolis, MN 55415
612-625-4909
612-625-7166 fax

19. Radiation Oncology Treatment Planning/Care Delivery Application

The goal of this project is to develop, implement, and evaluate NGI capabilities for radiation oncology treatment planning and care delivery. The application will provide diagnostic support, treatment planning, and remote verification of proper operation of treatment equipment from the Comprehensive Cancer Center to a remote Johns Hopkins University treatment facility. The proposed project will have a strong evaluation component focused on quality of service, security, privacy, and data integrity.

Contact: Joseph S. Lombardo
Johns Hopkins University
Applied Physics Laboratory
11100 Johns Hopkins Road
Laurel, MD 20723-6099
240-228-6287
240-228-6834 fax

20. Applications Layer Security Solution for Stationary/Nomadic Environments

This project will evaluate extant security techniques within the context of an open security architecture. The solution is based on security shared among collaborating parties, nomadic computing, and the privacy of medical information. The architecture includes user authentication, remote access to medical databases, nomadic computing, and confidentiality of data.

Contact: Brenda Garman
Motorola Space and Technology Group
1190 Winterson Road
Airport Square #14, Suite 350
Linthicum, MD 21090
410-859-4761
410-859-0787 fax

21. Human Embryology Digital Library

The goal of this study is to develop a research and education network for medical image acquisition and analysis. A high-performance optical network testbed will link government labs and universities with traditional medical research facilities. The focus of the project is on the analysis and delivery of digital histopathology image data. The proposal includes the definition of a set of demonstration projects that use a collaborative consultation system for research, surgical planning, and basic research.

Contact: George S. Michaels, Ph.D.
George Mason University
Office of Sponsored Programs
4400 University Drive
Fairfax, VA 22030
703-993-1998
703-993-1993 fax

22. Integration of Security Mechanisms for Internet Applications

The goal of this project is to develop a plan to integrate the PCASSO (Patient Centered Access Secure Systems Online) with biomedical applications. It will be demonstrated through a testbed involving medical treatment facilities in Delaware, Pennsylvania, Maryland, and New Jersey and the Frederick (Md.) Biomedical Supercomputer Center in an information technology infrastructure. The NGI infrastructure for this region is being developed under the HUBS (hospitals, universities, business schools, and communities) Initiative.

Contact: Raymond E. Cline, Jr.
Science Applications International Corp. (SAIC)
1710 Goodridge Drive, M/S 2-3-1
McLean, VA 22102
703-749-8648
703-821-1134 fax

23. Telemammography Using the NGI

The goal of this project is to plan and implement a testbed to demonstrate the feasibility of a national breast imaging archive and network infrastructure to support telemammography using NGI technologies. The proposed infrastructure would support traditional breast screening; provide the opportunity to maintain and apply standard image processing and computer-aided diagnosis software; permit access to breast imaging experts for primary and secondary interpretations; and provide an opportunity to study and understand epidemiologic issues in breast cancer.

Contact: Mitchell Schnall
University of Pennsylvania
Research Services
133 S. 36th Street, Suite 300
Philadelphia, PA 19104-3246
215-662-7238
215-662-3013 fax

24. Teletrauma and the NGI

The goal of this project is to plan the implementation of an integrated system of trauma care for Southern Louisiana using an NGI telemedicine network. This network will provide instant access to the Trauma Team at the Medical Center of Louisiana at New Orleans, which will provide online assistance. Distance education training for emergency personnel, network management, and quality of service issues are all elements of the project.

Contact: Richard Ferrans, M.D.
Louisiana State University
Medical Center
Department of Public Health
1600 Canal Street, Suite 800
New Orleans, LA 70112
504-588-3507
504-588-3938 fax

PHASE II AWARDS

1. Personal Internetworked Notary and Guardian

The Personal Internetworked Notary and Guardian (PING) proposal aims to provide a patient-controlled personal medical records system. The

PING record is available to the patient from any Internet-connected device. It is encrypted and accessible only to authorized parties for health care and/or research or public health purposes. It will include integration of data from two birth hospitals, a tertiary care pediatric hospital, a pediatric practice network, public health authorities, and the patients and their families. The goals of the PING project include (1) providing access for highly mobile postpartum mothers at work, school, and home to their infants' records; (2) enabling patients and families to manage a fundamentally collaborative process of clinical documentation over the Internet; and (3) ensuring that all PING transactions provide the highest available confidentiality of the patient's data, under their control.

Contact: Isaac S. Kohane, M.D., Ph.D.
Director, Children's Hospital Informatics Program
300 Longwood Ave., Enders 150
Boston, MA 02115
617-355-7821
617-730-0456 fax

2. Biomedical Teleimmersion

By combining teleconferencing, telepresence, and virtual reality, teleimmersion enables teachers and students to interact with three-dimensional models, point, gesture, converse, and see each other. Teleimmersion combines CAVE and ImmersaDesk virtual reality systems with advanced network capabilities to make learning environments so compelling that people will use them even when they are in the same room. They plan to demonstrate and assess teleimmersive environments for surgical education.

Contact: Jonathan C. Silverstein, M.D.
University of Illinois at Chicago
School of Biomedical and Health Information Services
1919 W. Taylor
Chicago, IL 60612-7249
312-996-5112
312-996-8342 fax

3. Patient-Centric Tools for Regional Collaborative Cancer Care Using NGI

This project plans to investigate the application of collaborative tools in a distributed and differentiated medical enterprise, the Seattle area Cancer

Care Alliance (CCA). The applications should (1) enhance the CCA partners' existing clinical care programs into new highly collaborative patient-centered interdisciplinary efforts; (2) allow for a fully integrated team approach to cancer, i.e., state-of-the-art diagnosis, treatment, and management of cancer patients through collaboration of distributed cancer care clinicians and researchers; and (3) accelerate the dissemination and application of new knowledge related to the diagnosis and the treatment of cancer, both inside the enterprise and throughout the region. They propose to examine the application of collaborative technologies to the three areas of physician interaction with patient information in the diagnosis, management, and treatment of cancer: consultations between referring physicians and CCA physician, including the patient; tumor board conferencing; and radiation oncology treatment planning.

Contact: Brent K. Stewart, Ph.D.
University of Washington
Grant and Contract Services
3935 University Way NE
Seattle, WA 98195
206-616-1314
206-543-3495 fax

4. Connectivity, Security, and Performance of an NGI Testbed for Medical Imaging Applications

The objective of this project is to implement an NGI testbed in northern California's San Francisco Bay Area for medical imaging applications. The two regional sites are the University of California at San Francisco (UCSF) and Stanford University. This NGI testbed will be built on two existing high-performance networks. The goal is to provide insight into NGI capabilities with respect to performance in a regional environment, potential for extension to the national level, and improvements needed. The clinical applications to be evaluated include telemammography consultation service in a regional compared with a local environment and how real-time interactive teaching in breast imaging would improve the confidence level of general-practice radiologists. The two characteristics of NGI that will be utilized include file size capability and near-real-time transmission.

Contact: H.K. Huang, D.Sc., FRCR (Hon.)
Children's Hospital of Los Angeles/University of Southern California
Department of Radiology, Mailstop #81
4650 Sunset Boulevard

Los Angeles, CA 90027
818-889-9411 (telephone and fax)

5. Indianapolis Testbed Network for NGI Applications to Telemedicine

Indiana University proposes to convert the Indianapolis Network for Patient Care (INPC) into a testbed of NGI technologies including IP Security (IPSec), quality of service (QoS) in televideo applications at a nursing home, and IP roaming capabilities with portable wireless workstations in clinical settings. The project plans to conduct randomized trials to test the effects of nursing home televideo and nomadic computing in the clinical environment. They plan to perform a trial to determine the effects of patient-physician videoconferencing and batch video applications (linked to Web-based electronic medical records) on health services utilization and physician/patient satisfaction at a 250-bed remote nursing home. The project will also perform a crossover trial of handheld personal computers, evaluating the effects on physician behavior by time-motion studies, physician satisfaction, and patient encounter data. These handheld computers will be equipped with capabilities for computerized order-entry, access to patient data, task-list management, and e-mail.

Contact: Clement MacDonald, M.D.
Regenstrief Institute for Health Care
101 West 10th Street, RG 6th Floor
Indianapolis, IN 46202
317-630-7070
317-630-6962 fax

6. Internet Protocol Video Telemedicine and Patient Cardiology Education

The purpose of this project is to address the technical issues impacting the delivery of telemedicine and sophisticated medical education using IP video over the Next Generation Internet (NGI). IP video over NGI has the potential to provide a common telecommunication infrastructure for real-time high bandwidth medical applications that cannot be supported by the commodity Internet. NGI solutions for real-time telemedicine with high-bandwidth video and audio requirements could eventually eliminate the need for expensive dedicated telemedicine networks and give broader access to these services. As part of the project, extensive evaluation, including impact on patient care, will be done. Technical and clinical protocols will be developed for all applications.

Contact: Susan S. Gustke, M.D.
East Carolina University School of Medicine
Center for Health Science Communication
Brody Medical Science Building, 1S-10
600 Moye Blvd.
Greenville, NC 27854
252-816-5219
252-816-8596 fax

7. A Multicenter Clinical Trial Using NGI Technology

NGI technology will be applied to provide the infrastructure of a multicenter clinical trial of new therapies for adrenoleukodystrophy (ALD), a fatal neurologic genetic disorder. This project involves the formation of a worldwide imaging network of clinical institutions to evaluate ALD therapies. This network is required to provide a sufficient number of patients for evaluating ALD therapies. This can serve as a model for many other disorders. Three centers will collaborate on this project. The Imaging Science and Information Systems (ISIS) Center at Georgetown University Medical Center, the Kennedy Krieger Institute, and the Department of Radiology at Johns Hopkins University. NGI technology will be used to speed the transmission and evaluation of high quality MRI images. Another important feature of this proposal is to gain insight into procedures that will ensure medical data privacy and security.

Contact: Hugo W. Moser, M.D.
Kennedy Krieger Research Institute, Inc.
707 North Broadway
Baltimore, MD 21205
410-502-9405
410-502-9839 fax

8. PathMaster: A Web-Accessible Cell Image Database Indexed by Mathematical Descriptors and Supported by Parallel Computation

The project will develop the PathMaster computer system as a testbed. PathMaster is designed to help the pathologist with the process of making a diagnosis in a cytologic specimen. Phase II focus will be on the analysis of lymphoma touch preparations and thyroid aspirates. To use PathMaster, the pathologist creates digitized images of a selected set of cells from a specimen and submits these to PathMaster over the Web. Each image is automatically subjected to a computational analysis to

determine more than 2,000 mathematically derived descriptors. Each image will then be compared to a database using network-based parallel computation. The analysis will produce ranked sets of images from specimens whose diagnosis is known. Images will be returned to the user to help in making a diagnosis. A variety of NGI testbed evaluations will be performed.

Contact: Perry L. Miller, M.D., Ph.D.
Yale University School of Medicine
Center for Medical Informatics
333 Cedar Street, P.O. Box 208009
New Haven, CT 06520-8009
203-785-6753
203-785-6664 fax

9. Remote, Real-Time Simulation for Teaching Human Anatomy and Surgery

Stanford University proposes to develop two teaching applications and a local NGI testbed network for evaluating their effectiveness. The first application will support instruction in human anatomy and the second the performance of surgical manipulations. Both applications will support synchronous collaboration through a shared virtual workspace and use haptic feedback to augment the visual sense. This technology will allow the definition of new curricular elements including the repeated dissection of anatomical structures, the visual segmentation of raw data sets, the creation of three-dimensional organ models, and the practice of fundamental surgical skills. The investigators anticipate that a wide community of teachers and users will, through a distributed client-server system, share online, image-rich data and professional experiences.

Contact: Parvati Dev, Ph.D.
Stanford University
1215 Welch Road, MOD B
Stanford, CA 94035-5401
650-723-8087
650-498-4082 fax

10. Human Embryology Digital Library and Collaboratory Support Tools

George Mason University proposes to develop and demonstrate technologies to enable collaboration between multiple, distributed research-

ers and to make progress toward advanced clinical and educational goals. The offeror plans to integrate existing data capture and analysis procedures at the National Museum of Health and Medicine (NMHM) into a high performance testbed network that will include a petabyte archive and analysis capability. The project will use an existing, government-funded gigabit network to connect the NMHM to key sites across the nation. The testbed requires a minimum data transport rate of 622 Mbps in the key regional networks and quality of service.

Contact: J. Mark Pullen, Ph.D.
George Mason University
Computer Science MS 4A5
4400 University Drive
Fairfax, VA 22030
Phone: 703-993-1538
703-993-1710 fax

11. Medical Nomadic Computing Applications for Patient Transport

The objective of this project is the real-time transmission of multimedia patient data from an incident scene and during transport to a receiving center enabling diagnostic and treatment opportunities prior to arrival. The offeror will use the diagnosis and treatment of challenging clinical models—including acute ischemic stroke and trauma scene response—to define a range of quality of service (QOS) requirements for multiple critical care applications, evaluate the effectiveness of the system, and derive principles of nomadic computing applicable in other time sensitive emergency care models in which treatment options are constrained by the delay between onset/injury and definitive diagnosis. TRW and the University of Maryland, Baltimore had previously developed a mobile telemedicine system for remote, real-time diagnosis using narrow bandwidth wireless technologies, but suffered from QOS problems. Phase III will extend the trial in a larger regional setting.

Contact: David M. Gagliano
TRW, Inc.
One Federal Systems Park Drive
Fairfax, VA 22033
Phone: 703-345-7497

12. Remote Treatment Planning System

This proposal addresses the development, implementation, and evaluation of an application to support remote treatment planning for radiation therapy. This application, Remote Treatment Planning System (RTPS), relies on network infrastructure technology for collaboration; on high bandwidth and QOS to support interactive review sessions; and on data privacy and security to protect patient privacy, confidentiality, and data integrity. Review sessions provide a collaborative environment for dosimetrists at the planning site, the oncologists at the care delivery site, and peer reviewers. It utilizes video teleconferencing and a shared view of the images to support treatment planning. The evaluation will measure outcomes at the care delivery site, process improvements at the treatment-planning site, and estimate cost impact on the remote treatment planning process. Phase III is proposed to deploy the application and testbed features to Peninsula Regional Medical Center in Salisbury, Maryland. Connectivity will be provided by the Maryland State Asynchronous Transfer Mode (ATM) backbone, NetWork.Maryland.

Contact: Joseph S. Lombardo
Johns Hopkins University Applied Physics Laboratory
11100 Johns Hopkins Drive
Laurel, MD 20723-6099
240-228-6287
240-228-5026 fax

13. Next Generation Internet Implementation to Serve Visible Human Datasets Phase II: Development of Testbeds

The University of Michigan (UM) Visible Human (VH) project team will develop an NGI production system to serve visible human data sets. These include a comprehensive set of interactive 2D and 3D VH browsers with arbitrary 2D cutting and 3D visualizations. An interactive Web navigation engine will be deployed to create and visualize anatomic fly-through, under haptic control of the user, and to deliver fly-through developed by expert anatomists and clinicians. Anatomical labels will enhance these visualization sequences and enable real-time links with appropriate resources on the Web using XML. The UM NGI VH system will complement and extend currently deployed passive Web information systems with active computational services. This will allow for delivery of several simultaneous high-quality digital streams, creating structured medical knowledge using the VH datasets. An evaluation team will continually

respecify and focus the testbed deployments and measure performance and educational effectiveness.

Contact: Brian D. Athey, Ph.D.
University of Michigan School of Medicine
Ann Arbor, Michigan 48109-0616
734-763-6150
734-763-1166 fax

14. Networked 3D Virtual Human Anatomy, Phase II

The University of Colorado Health Sciences Center proposes to demonstrate and assess the use of Web-based, 3D-explorable virtual humans to enhance traditional anatomic teaching. This will be accomplished with audio, graphic, and haptic interfaces. The application will be assessed in anatomy curricula developed for undergraduate to postgraduate levels of education. Modules teaching the anatomy, function, and pathology of the knee will be used for this demonstration. The investigators will also demonstrate an extension of the virtual environment to include surgical simulation applied to arthroscopy.

Contact: Victor M. Spitzer, Ph.D.
University of Colorado Health Sciences Center
13001 East 17th Place, PO Box 6508, Mail Stop F-435
Aurora, CO 80045-0508
303-724-0501
303-724-0911 fax

15. Mammography for the Next Generation Internet, Phase II

The University of Pennsylvania proposes to develop a testbed to demonstrate the feasibility of a national breast imaging archive and network infrastructure to support digital mammography using NGI technologies. They plan to improve access and performance of breast cancer screening with an imaging archive that supports storage, retrieval, and distribution of breast images for clinical and research purposes and ensures privacy and confidentiality with multilevel security embedded throughout the system. The proposed infrastructure would (1) support traditional breast screening through the maintenance and distribution of a digital record of prior breast examinations and relevant medical history for primary interpretation and expert consultation; (2) provide the opportunity to maintain and apply computer-aided diagnosis (CAD) software at central, well-maintained computing resources to studies from all women; (3) provide

unique tools for creating educational and training programs; and (4) create an unparalleled opportunity to study and understand many epidemiologic issues in breast cancer through searches of a national breast screening database. NGI technologies will be used to transfer large data files, execute real-time queries, and access information securely. The testbed will demonstrate that quality of service, medical data privacy and security, nomadic computing, network management research and development, and infrastructure technology for collaboration are NGI technologies that are integral to widespread deployment and optimal utilization of digital mammography.

Contact: Mitchell Schnall, M.D.
University of Pennsylvania
Radiology Department
1 Silverstein
3400 Spruce Street
Philadelphia, PA 19104
215-662-6470
215-662-3013 fax

APPENDIX C

Biographies of Committee Members

Edward H. Shortliffe (*chair*) recently moved from Stanford University to Columbia University, where he serves as professor and chair of the Department of Medical Informatics. He also holds appointments as professor in the Medicine and Computer Science Departments. Dr. Shortliffe is a member of the Institute of Medicine and has served on the Computer Science and Telecommunications Board, the Federal Networking Advisory Committee (National Science Foundation), and the Biomedical Library Review Committee (National Library of Medicine) and was the recipient of a research career development award from the NLM. In 1993, he co-chaired a CSTB planning meeting on the role of information infrastructure in health care. He currently sits on the President's Information Technology Advisory Committee (PITAC). Dr. Shortliffe combines expertise in medicine and computer science. He received an A.B. in applied mathematics from Harvard College in 1970, a Ph.D. from Stanford in medical information sciences in 1975, and an M.D. at Stanford in 1976. During the early 1970s, he was principal developer of the medical expert system known as MYCIN. After a pause for internal medicine house-staff training at Harvard and Stanford between 1976 and 1979, he joined the Stanford internal medicine faculty, where he directed a research program in medical expert systems development. Dr. Shortliffe is interested in a broad range of issues related to integrated decision-support systems and their effective implementation. He spearheaded the formation of a Stanford degree program in medical informatics and at Columbia is continuing to divide his time between clinical medicine and medical

informatics. Dr. Shortliffe is a member of the American Society for Clinical Investigation, the American Association of Physicians, and the American Clinical and Climatological Association. He has also been elected to fellowship in the American College of Medical Informatics and the American Association for Artificial Intelligence. He sits on the editorial boards of several medical computing and artificial intelligence publications. In addition, he received the Grace Murray Hopper Award of the Association for Computing Machinery in 1976 and has been a Henry J. Kaiser Family Foundation Faculty Scholar in General Internal Medicine. Dr. Shortliffe has authored over 180 articles and books in the fields of medical computing and artificial intelligence. Volumes include *Computer-Based Medical Consultations: MYCIN* (Elsevier/North-Holland, 1976), *Readings in Medical Artificial Intelligence: the First Decade* (with W.J. Clancey; Addison-Wesley, 1984), *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project* (with B.G. Buchanan; Addison-Wesley, 1984), and *Medical Informatics: Computer Applications in Health Care and Biomedicine* (with L.E. Perreault, G. Wiederhold, and L.M. Fagan; Addison-Wesley, 1990; 2nd ed., Springer-Verlag, Spring 2000).

Russ Biagio Altman is associate professor of medicine (and computer science by courtesy) at Stanford University. His primary research interests are in the application of computing technology to basic molecular biological problems of relevance to medicine. He is currently developing techniques for collaborative scientific computation over the Internet, including novel user interfaces to biological data. Other work focuses on the analysis of functional microenvironments within macromolecules and the application of nonlinear optimization algorithms for determining the structure and function of biological macromolecules, particularly the bacterial ribosome. He is on the executive committee (as Molecular Science Thrust leader) for the National Partnership for Advanced Computational Infrastructure (NPACI), the NSF-sponsored program at the San Diego Supercomputer Center. Dr. Altman holds an M.D. from Stanford Medical School, a Ph.D. in medical information sciences from Stanford, and an A.B. from Harvard College. He has been the recipient of the U.S. Presidential Early Career Award for Scientists and Engineers, a National Science Foundation CAREER Award, and the Western Society of Clinical Investigation Annual Young Investigator Award. He is a fellow of the American College of Physicians and the American College of Medical Informatics.

Patricia Flatley Brennan is Moehlman Bascom Professor at the School of Nursing and College of Engineering at the University of Wisconsin. Dr. Brennan's research is in the area of nursing informatics and examines

ways to use the Internet, home-based computer systems, and specialized information resources to promote self-care and disease management skills among patients and their family caregivers. She earned a Ph.D. and an M.S. in industrial engineering from the University of Wisconsin at Madison, an M.S. in nursing from the University of Pennsylvania, and a B.S. in nursing from the University of Delaware.

Bruce Davie works at Cisco Systems in Chelmsford, Massachusetts, where he is a Cisco Fellow. He received his B.E. (electrical) from the University of Melbourne, Australia, in 1984. He completed his Ph.D. in computer science at the University of Edinburgh, Scotland, in 1988 before starting work in the Computer Networking Research Department at Bell Communications Research (Bellcore) in Morristown, N.J. While at Bellcore, he worked on the Aurora gigabit testbed, the first wide-area gigabit network. At Bellcore he held the positions of director of Internet networking research and chief scientist and led research efforts on the Next Generation Internet Protocol, IP-over-ATM, and the support of real-time applications over the Internet. Dr. Davie joined Cisco Systems in 1995. He leads a group working on the development of multiprotocol label switching and quality-of-service capabilities for IP networks. He is the author of numerous journal articles, conference papers and book chapters, and three books on computer networks. He is an active member of both the Internet Engineering Task Force and the End-to-End research group.

William M. Detmer is president and chief executive officer of a new Internet-based digital publishing company and is adjunct assistant professor of health evaluation sciences at the University of Virginia. He was formerly vice president of clinical information products for Ovid Technologies, Inc. Dr. Detmer's interests lie in the characterization of the knowledge needs of clinicians and the development of content, information science methods, and technologies that collectively meet those needs. In 1994 he developed WebMedline, the first World Wide Web interface to Medline, and he since has developed a variety of Internet-based applications and methods that bring medical knowledge to the point of care. Dr. Detmer holds a B.A. degree from the University of North Carolina at Chapel Hill, where he was a Morehead Scholar, an M.D. from the University of California San Francisco, and a master's degree in medical information science from Stanford University. He is board certified in internal medicine.

Valerie Florance is project director for `better_health@here.now`, the information technology futures initiative of the Association of American Medical Colleges (AAMC), and principal investigator for IAIMS: The Next

Generation, a state-of-the-art review of information management in academic health sciences centers. Before joining AAMC, she was director of academic information at the University of Rochester Medical Center. Dr. Florance has over 17 years of experience in health sciences libraries and information services at the University of Utah, Johns Hopkins, and the University of Rochester. She holds a B.A. in anthropology and an M.A. in medical anthropology from the University of Utah, an M.L.S. from Brigham Young University, and a Ph.D. in library and information science from the University of Maryland. She has served in a number of professional leadership roles, most recently as editor of *Annual Statistics of Medical School Libraries in the United States and Canada*. In 1995 and 1997, she won the Ida and George Eliot Prize from the Medical Library Association. Her scholarly interests center on technology, infrastructure, and policy issues in health sciences information management.

Andrew Friede is a physician executive with the Cerner Corporation, a major supplier of health care information systems and services. Dr. Friede writes expert systems, designs information systems for medical researchers, and advises health care organizations on strategy. He was formerly at the Centers for Disease Control and Prevention (CDC), where he led the development of CDC WONDER, an integrated information and communications system that provides online access to over 30 public health databases, and has specialized features for CDC surveillance programs. His epidemiological research interests focused on maternal and child health conditions, including immunizations, injuries, infant and maternal mortality, teenage pregnancy, and obstetrical complications. He has been a consultant to the World Bank and the U.N. for public health and clinical information systems in China and Madagascar. He is the author of numerous publications and served as the chief editor of *CDC Prevention Guidelines*.

Mark Frisse is vice president of clinical information services at Express Scripts, Inc., a pharmaceutical benefits management company. Until December 1999, he served as director of the Bernard Becker Medical Library; director of the Medical Informatics Laboratory; associate dean for Academic Information Management; professor of medicine, and academic director of the Health Services Management Executive M.B.A. Program in the John M. Olin School of Business at Washington University. In these roles he was responsible for a broad array of medical informatics research and teaching activities.

John Glaser is vice president and chief information officer for Partners Healthcare System. He was founding chair of the College of Healthcare

Information Management Executives, is past president of the Healthcare Information Management Systems Society, and was the 1994 recipient of the John Gall award for health care chief information officer of the year. Dr. Glaser previously managed the health care information systems consulting practice at Arthur D. Little. He is the author of more than 40 publications on health information systems and holds a Ph.D. in health care information systems from the University of Minnesota.

John Huffman is chief technology officer at Stentor, Inc., a company specializing in digital medical applications. He was previously in the Corporate Research and Development Group of Silicon Graphics Inc. (SGI) and managed its medical market technology and business development efforts. Mr. Huffman has concentrated on medical imaging applications for the last 10 years. Additional responsibilities while at SGI were the design and specification of next-generation processor and graphics architectures, development of image- and signal-processing methods, and numeric computation methods. Prior to joining SGI, Mr. Huffman was the cofounder and chief technology officer for Aware, Inc., where he developed image and video compression systems. Mr. Huffman holds several of the early wavelet patents in this field and wrote the only FDA-approved package for medical image compression. In addition, he worked on the development of adaptive signal processing methods, numeric solution of nonlinear partial differential equations, and video compression methods. Prior to joining Aware, Mr. Huffman spent two years working in the artificial intelligence group at the Microelectronics and Computer Technology Corporation under Doug Lenat, developing CYC, the largest, most comprehensive artificial intelligence system integrating all known methods into one reasoning system. Before that, Mr. Huffman was one of the original employees at Thinking Machines Corporation, where he designed the VLSI chip and contributed to microcoding the Connection Machine. Mr. Huffman also developed automated circuit placement methods that are in use in the industry today. Mr. Huffman began his professional career at the National Security Agency as a VLSI design engineer while performing graduate work in computer science and electrical engineering at the University of Maryland. He completed his undergraduate work in chemical physics at the University of Chicago in 1979.

Isaac Kohane is director of the Children's Hospital Informatics Program in Boston and assistant professor of pediatrics at Harvard Medical School. Dr. Kohane was one of the lead developers of the W3EMRS system, a system that allows secure Internet-based sharing of medical information among emergency room doctors in five Boston-area hospitals. From 1995 to 1997, Dr. Kohane was leader of the Boston Collaborative Group for

Web access to electronic medical systems. He is currently working on the Health Information Identification and De-Identification Toolkit to enable the specification of health information systems with multiple trade-offs in confidentiality. He is also principal investigator of the Personal Internetworked Notary and Guardian project to enable patients to fully control their own medical records over multiple institutions from the Internet. Dr. Kohane earned a joint M.D./Ph.D. from Boston University in 1987 and a B.Sc. in biology from Brown University. He conducted his doctoral research in collaboration with the Clinical Decision Making Group in MIT's Laboratory for Computer Science, with which he maintains an affiliation.

Carl E. Landwehr is a senior fellow and acting director of the Information Security Center at Mitretek Systems, a nonprofit center conducting research and development in the public interest. For many years, Dr. Landwehr headed the Computer Security Section of the Center for High Assurance Computer Systems at the U.S. Naval Research Laboratory. He has led a variety of research projects to advance technologies for computer security and high-assurance systems and has served on review panels for high-assurance research and development programs at the National Aeronautics and Space Administration and the National Security Agency. Dr. Landwehr serves as an expert consultant to the North Atlantic Treaty Organization and for 10 years chaired an international defense panel on secure information systems. The International Federation for Information Processing (IFIP) awarded him its Silver Core for his work as founding chair of IFIP Working Group 11.3 on database security, and the IEEE Computer Society awarded him its Golden Core for his work on behalf of its Technical Committee on Security and Privacy. He has served on the editorial boards of the *High Integrity Systems Journal*, *IEEE Transactions on Software Engineering*, and the *Journal of Computer Security*. Dr. Landwehr also served on the CSTB committee that produced the report *For the Record: Protecting Electronic Health Information*. He received a B.S. (engineering, 1968) from Yale University, an M.S. (computer and communication sciences, 1970) from the University of Michigan, and a Ph.D. (computer and communication sciences, 1974) from the University of Michigan.

Daniel R. Masys is director of biomedical informatics at the University of California, San Diego School of Medicine, and associate clinical professor of medicine. An honors graduate of Princeton University and the Ohio State University College of Medicine, he completed postgraduate training in internal medicine, hematology, and medical oncology at the University of California at San Diego and the Naval Regional Medical Center, San

Diego. He served as chief of the International Cancer Research Data Bank of the National Cancer Institute, National Institutes of Health, and from 1986 through 1994 was director of the Lister Hill National Center for Biomedical Communications, which is the computer research and development division of the National Library of Medicine. He also served as the NIH representative to the federal High Performance Computing, Communications, and Information Technology committee, which advised the President's Office of Science and Technology Policy in the area of advanced computing and national information infrastructure. Dr. Masys is a diplomate of the American Board of Internal Medicine in medicine, hematology, and medical oncology. He is a fellow of the American College of Physicians and a fellow of the American College of Medical Informatics. He is a founding associate editor of the *Journal of the American Medical Informatics Association* and has received numerous awards, including the NIH Director's Award, the Public Health Service Outstanding Service Medal, and the U.S. Surgeon General's Exemplary Service Medal. Dr. Masys' research interests are in Internet-accessible health information and information systems support for clinical research. He is co-principal investigator of the Patient-Centered Access to Secure Systems Online (PCASSO) research project funded by the National Library of Medicine, which is developing and evaluating a secure Web-based access method for clinical data.

Jane E. Sisk is a professor in the Department of Health Policy and co-director of the Center for Evidence-Based Medicine and Aging at Mount Sinai School of Medicine in New York City. She was formerly a professor in the Division of Health Policy and Management and director of the Master's Program in Effectiveness and Outcomes Research at Columbia University School of Public Health. Before joining Columbia, she directed projects at the Congressional Office of Technology Assessment as a senior associate and project director in the Health Program. She also served as president, International Society of Technology Assessment in Health Care, of which she was a founding member. Dr. Sisk is a fellow of the Association for Health Services Research and sits on the editorial boards of *Health Services Research* and the *International Journal of Technology Assessment in Health Care*. She served on the IOM committee that produced the report *Telemedicine: A Guide to Assessing Telecommunications in Health Care* and is currently a member of the IOM Committee on Immunization Finance Policies and Practices and the NRC-IOM National Cancer Policy Board. Dr. Sisk earned a B.A. in international relations (Phi Beta Kappa, magna cum laude) from Brown University, an M.A. in economics from George Washington University, and a Ph.D. in economics from McGill University.

Thorsten Von Eicken is assistant professor of computer science at Cornell University. His research has focused on high-performance communication in clusters of workstations and developed the U-Net user-level networking architecture to close the dramatic gap between the bit rate of high-speed networks and the communication performance seen by applications. He recently started a new project: the Safe Language Kernel (SLK), which is an operating system infrastructure for customizable Internet servers and application-specific gateways. The primary goal of SLK is to allow users to download custom services into servers in the network in a secure yet flexible manner. Just as Java enables Web browsers in which users safely download applets, SLK will enable safe Internet servers into which users can upload servlets. Dr. Von Eicken received his Ph.D. from the University of California at Berkeley. He was the recipient of an NSF CAREER Award in 1997. He has published widely on high-performance computing architectures, parallel computing/programming, and low-latency data communications.

APPENDIX D

Individuals Who Participated in Site Visits or Briefed the Study Committee

SITE VISIT PARTICIPANTS

December 16, 1998

California (San Francisco, Oakland, Palo Alto, and Mountain View)

University of California at San Francisco, Laboratory for Radiological Informatics: Fei Cao, Tony Chou, Bill Dillon, Steve Frankel, David Hoogstrate, H.K. (Bernie) Huang, Andrew Shyh Liang Lou, Laura Snarr, Johannes Stahl, Albert Wong, and X.Q. Zhou.

Kaiser-Permanente: Peter Juhn, Richard Leopold, Anna-Lisa Silvestre, and Valerie Tolous-Shams.

Stanford Center for Professional Development (SCPD): Andrew DiPaolo, Aubrey Harris, Jay Kohn, and Mike Rouan.

NASA Ames Research Center: Cynthia Bruyns, Shirley Burg, Rei Cheng, Muriel Ross, and Xander Twombly.

February 2-3, 1999

North Carolina (Greenville and Chapel Hill)

East Carolina University Center for Health Sciences Communication, Telemedicine Program: David C. Balch, Doug Barnum, Christi Brewer, Thomas Feldbush, Susan Gustke, Gloria Jones, Marc Krien, Ted Kummer, Lori Maiolo, Lance Rogers, and Ron Rouse.

University of North Carolina at Chapel Hill Computer Graphics Laboratory: Andrew Ade, Henry Fuchs, and Lars S. Nyland.

University of North Carolina at Chapel Hill School of Medicine: John Loonsk.

February 10-11, 1999

Seattle, Washington

University of Washington: Thomas Anderson, James Brinkley, Steve Corbato, Sherrilyne Fuller, Tom Furness (Human Interface Technology (HIT) Laboratory), Harold Goldberg, Ira Kalet, Debra Ketchell, Edward Lazowska, Henry Levy, James LoGerfo, Tom Martin, Jean O. Nelson, Tom Norris, Peter Oppenheimer (HIT Laboratory), Brent Stewart, and Suzanne Weghorst (HIT Laboratory).

Regence BlueShield: Magdalene Aliu, Kirk Bailey, Jac Davies (Washington State Department of Health), Steven Moe, Donn Morse, Richard Rubin (Foundation for Health Care Quality), Peter Summer-ville (Community Health Information Technology Alliance), and Jerry Tonkovich.

BRIEFERS DURING COMMITTEE MEETINGS

September 13-14, 1998

Washington, D.C.

Michael Ackerman and Donald Lindberg, National Library of Medicine; George Strawn, National Science Foundation.

December 16-18, 1998

Palo Alto, California

Alan Hannan, Frontier/Globalcenter; Hon Hso, SBC Communications; Gary Leiber, WebTV Education; Milo Medin, @Home; Jackie Parker, Intel; Geoff Rutledge, Healthcon; J.J. Singh, Caresoft; Mark Stefik, Xerox Palo Alto Research Center; Hal Varian, dean, School of Information Systems & Management, University of California at Berkeley.

March 1, 1999 - March 2, 1999

Washington, D.C.

Donald Detmer, former chair of the National Committee on Vital and Health Statistics; Michael Fitzmaurice, Agency for Health Care Policy Research; Margaret Hamburg, Department of Health and Human Services; Thomas Kalil, National Economic Council; James Ostell, National Center for Biotechnology Information; Richard Satava, Yale University; and William Yasnoff, Centers for Disease Control and Prevention.

Index

A

Access controls, 64-65, 145, 157-160, 242-243, 256, 302

Adrenoleukodystrophy, 328

Aetna U.S. Healthcare, 60

Agency for Healthcare Research and Quality (AHRQ), 120, 223, 228, 229, 234 (n. 25), 255

Agora, 311

AIDS, 55 (n. 12), 60, 95, 103, 313

Alternative medicine, 60

American Medical Association, 155-156

American Medical Colleges, 261

American Medical Informatics Association, 14, 22, 261, 307

American National Standards Institute, 85

Anatomy, 48, 102, 108, 114-115, 244, 296, 300, 306, 307-309, 315, 317, 320, 329, 331, 332

Anonymity, 167-173

anonymous e-mail, 169-170

anonymous data, 172-173

anonymous payment, 172

protected Web browsing, 170-172

Association of American Medical Colleges, 130 (n. 30), 261

Asynchronous transfer mode, 39, 56 (n. 19), 72, 117, 134, 163-166, 272, 281, 290, 331

AT&T, 172

Audio, 31, 50, 100, 104, 105, 106, 157, 164, 220, 280, 285, 286, 288, 290, 293, 294, 296, 327, 332

see also Video and teleconferencing

Authentication, 17-18, 63-65, 70, 145, 253, 279, 297

biometric, 18, 55 (n. 13), 279

certification authorities, 151-152, 155-156, 253

encryption and, 151-153

protocols, 150, 151-156

remotely controlled medical devices, 69

token-based, 17-18, 83, 84

Availability (network and system), 9, 11, 15, 24, 40, 45, 87-88, 93, 122, 124, 144, 160-162, 191, 237, 242-244

defined, 10

digital divide, 14-15, 71, 125, 195, 209, 210, 211, 245-246

e-mail between patients and providers, 63, 70

in consumer health applications, 63, 70, 87-88, 124

in professional education, 107-108

in public health, 101

see also Data security; Intellectual property; Privacy; Ubiquity of access

B

Backbone networks, 11, 41, 42, 49, 50, 78, 134, 135-137, 140, 253, 284, 299, 331
Very high performance Backbone Network Service (vBNS), 49, 51, 115, 116, 284, 294, 299
Balanced Budget Acts, 92, 219, 220-221, 224
Bandwidth, *ix*, 6, 9, 11-12, 39, 41, 51, 123, 132-134, 240, 244-245, 252-253, 289, 305
cable modems, 12, 18, 34, 39, 40, 49, 67, 132, 163-166
costs of increases, 135, 166-167, 209
defined, 10, 39
digital subscriber line services, 132, 164, 165, 166
for administrative and financial applications, 93
for biomedical research, 110, 111, 121-122, 123, 240-241
for clinical care, 72-73, 75, 77-79, 82-83, 87, 88
for consumer health applications, 61, 65, 67-68, 69, 71, 280
for e-mail between patients and providers, 63
for emergency care, 133
for home care, 67-68, 71, 267
for images, 73-74, 77-79, 80, 123, 134, 136, 162, 164, 241, 244, 272, 288-289, 301-302
for patient records, 65, 82-83, 320-321
for professional education, 104, 105, 107, 123, 126, 241, 281, 282, 290-291, 297-298
for public health, 100
for video and teleconferencing, 11, 73-74, 100, 132-133, 241, 244, 290
integrated services digital networks (ISDN), 74, 287, 288, 289, 290, 304
local loops, 61, 162-167
National Health Alert Network, 99-100
Next Generation Internet, 16, 49, 250-251
overprovisioning, 137, 138-139
private networks, 43
reconfigurable, 18, 164, 253
satellite telecommunications, 166-167
Very high performance Backbone Network Service (vBNS), 49, 51, 115, 116, 284, 294, 299

Biomedical research, *ix*, 1, 2, 3, 5, 16, 57, 108-123, 255, 306-307, 319
bandwidth needed, 110, 111, 121-122, 123, 240-241
clinical trials, 55(n.2), 120-121, 328
collaborative research, 118-120
data access controls and, 110-111, 122, 159, 316
databases, 109-112, 240-241, 295-298, 331-332
Digital Anatomist, 307-309
MEDLARS, 33, 55 (n. 5), 135
MEDLINE, 3, 33, 37, 55 (n. 5), 103-104, 109, 110, 215, 295, 301, 307
pathology image data system, 314-315
private sector, 109-110
digital images, 211, 245-246
expert systems, 227, 228-229
patient records, 65, 169, 173-174
publication of results, 116-118, 215-217, 239; *see also* Intellectual property; MEDLARS; MEDLINE
remotely controlled equipment, 5, 113-116, 240
simulations, 112-113, 118, 320
technical requirements, 6, 8, 108-123, 159, 240-241, 255
Biomedical Teleimmersion, 316, 325
Biometric authentication, 18, 55 (n. 13), 279
Blue Cross/Blue Shield, 89
Breast cancer, 55 (n. 12)
telemammography, 48, 78, 79, 142-143, 236, 244, 272-273, 324, 332-333
Bureau of the Census, 172-173

C

Cable modem technologies, 12, 18, 34, 39, 40, 49, 67, 132, 163-166
Cancer, 55 (n. 12), 62, 300-301, 305-306, 322, 325-326, 328-329
National Cancer Institute, 120-121, 273, 300-301, 305, 322, 324, 325-326
telemammography, 48, 78, 79, 142-143, 236, 244, 272-273, 324, 332-333
Cardiology, 30, 66-67, 73, 74, 128, 158, 266, 272, 273, 275, 280, 285, 286, 287-288, 289, 291, 302, 304, 327-328
Centers for Disease Control and Prevention, 95, 96, 99-100, 225, 264, 312-313

- Certification authorities, 151-152, 155-156, 253
- Chat groups, *ix*, 5, 60, 191, 192-193, 196, 239, 261, 290
see also Support groups
- Chronic illness, general, 13-14, 34, 55, 58, 60, 191
home care, 68
risk factors, 3, 31, 34, 58, 59, 60, 95-96, 274, 288
- Cities, *see* Urban areas
- Clinical care, 2, 55 (n. 2), 55 (n. 11), 57, 71-88, 120-121
and e-mail, 62, 261, 262
patient/provider exchanges, 14, 22, 37, 62-64, 66, 70, 189, 192, 197, 248, 279
bandwidth needs, 72-73, 75, 77-79, 82-83, 87, 88
data security, 65, 68, 82-84, 88
expert systems, 227, 228-229
latency needs, 81-83, 87-88, 124
Next Generation Internet, 48, 319
organizational factors, 182-184, 188, 189, 192, 195, 196
remote consultations, 6, 8-9, 16, 17, 18, 20, 30-33, 36, 57, 71, 72-76, 88, 220-221, 228-229, 236-237, 238-239, 244, 247, 258, 266-267, 280, 285-286, 288-289, 292, 323
standards/guidelines, 5, 14, 22, 57, 238-239, 275, 278
technical requirements, 4-5, 71-76, 87-88, 238-239, 246, 258
see also Diagnosis; Emergency care; Hospitals; Primary care physicians; Patient records; Specialists
- Clinical trials, 55 (n. 2), 120-121, 328
- Collaboratory for Microscopic Digital Anatomy, 114-115
- Collections of Information Antipiracy Act, 218
- Community Health Information Technology Alliance (CHITA), 90, 91, 309, 311-312
- Comprehensive Telehealth Act, 224
- Compression technology, 72, 77, 78-79, 80, 114, 115, 116, 129, 272, 286, 288, 290, 315
- The Computer-Based Patient Record*, 204
- Computed tomography (CT), 4, 65, 72, 78, 134, 164, 238, 274, 300
- Computer Science and Telecommunications Board, 1-2, 144, 173, 229
- Conference on Fair Use, 217
- Confidentiality, *see* Encryption; Privacy; Security
- Consumer Assessment of Health Plans Survey, 90
- Consumer health, 2-3, 27, 35, 37-38, 57, 58, 71, 152, 235, 245-246, 262
alternative medicine, 60
availability of access, 63, 70, 87-88, 124
bandwidth needs, 61, 65, 67-68, 69, 71, 280
chat groups, *ix*, 5, 60, 191, 192-193, 196, 239, 261, 290
data security for, 61, 63, 64-65, 68, 70-71, 238-239, 249
e-commerce and, 9, 10-11, 35, 60-61, 63, 71, 85, 168, 179-182, 310
e-mail, patient/provider exchanges, 14, 22, 37, 62-64, 70, 189, 192, 197, 248, 279
foreign language speakers, 119-120, 266-267
latency needs, 62, 69-70, 279
organizational factors, 9, 10-11, 179-182, 183, 191, 195
primary online activities, ranked, 61
quality control, 2, 14, 61-62, 246-247, 255-256, 292
remote consultations, 6, 8-9, 16, 17, 18, 20, 30-33, 36, 57, 71, 72-76, 88, 220-221, 228-229, 236-237, 238-239, 244, 247, 258, 266-267, 280, 285-286, 288-289, 292, 323
risk factors, 3, 31, 34, 58, 59, 60, 95-96, 274, 288
support groups, *ix*, 5, 60, 61, 196, 239
technical requirements, 4-5, 6, 8-9, 66-68, 69-71, 238-239, 245-246
- Web sites, 9, 10-11, 59-62, 64-65, 69, 190, 196-197, 216, 256, 262, 277-280
pharmaceutical companies, 60, 71, 189, 221
see also Clinical care; Data security; Home care; Insurance; Managed care; Privacy; Records of patients; Telemedicine; *cross-references under* Demographic factors

Copyright, *see* Intellectual property
Cost and cost-benefit factors, 8, 13, 37, 179-180, 184, 185, 186-187
 administrative, 89, 91-92, 97-98, 279
 bandwidth increases, 135, 166-167, 209
 clinical trials, 120-121
 data security, 144, 279
 demonstration projects, 257
 digital images, 211, 245-246
 electronic publishing, 216
 e-mail, 62, 63-64
 integrated delivery networks (IDNs), 35-36
 integrated services models, 140-141
 managed care, 59, 184
 medical imaging, 77, 79, 186
 organizational factors, 59, 179-180, 184, 185, 186-187, 194-195, 196
 patient monitoring and home care, 66
 patient record sharing, 81-82, 144
 price of Internet services and PCs, 209-210
 private networks *vs* Internet, 44-45
 professional education, 282-283
 property, urban areas, 210-211
 publication of biomedical research, 216
 public health surveillance and, 97-98, 101-102
 remote consultation via private networks, 74
 remotely controlled devices, 76
 spot market purchase of health care, 191-192
 telemedicine, 184, 204, 287, 302, 303
Cryptography, *see* Encryption
CT, *see* Computed tomography

D

Defense Advanced Research Projects Agency, 46, 47, 49, 51, 250, 300
Demographic factors, *see* Educational attainment; Geographic factors; Language factors, human; Rural areas; Socioeconomic status; Urban areas
Demonstration projects, 15, 19-22, 250, 257-261
 Internet 2, *ix*, 1, 50-51, 52
 National Library of Medicine awards, 314-333
 payment methods, 221
 private sector, 20, 260-261
 scalability of protocols, 16-17
 testbed networks, 2, 16, 17, 20-21, 28, 47, 250-252, 259, 283-284, 323, 326-327
 see also Next Generation Internet
Department of Agriculture, 212, 215
Department of Commerce, 33, 209-210, 212, 225, 226
 Bureau of Census, 172-173
Department of Defense, 20, 21, 24, 43, 161, 221, 223, 225, 226, 258, 259, 264, 299
Department of Energy, 43, 46, 49, 225, 226, 250, 299
Department of Health and Human Services, 14, 20, 23, 24-25, 62, 86, 92, 208, 215, 221, 223, 225-227, 255, 257, 259-260, 262-265
 Agency for Healthcare Research and Quality (AHRQ), 120, 223, 228, 229, 255
 Indian Health Service, 20, 21, 24, 258, 259
 see also Health Care Financing Administration; National Institutes of Health; National Library of Medicine
Department of Veterans Affairs, 20, 24, 221, 223, 225, 258, 259, 264
Dermatology, 73, 184, 191, 274, 285, 286, 289, 291, 302
Diabetes, 55 (n. 11), 61, 66, 68, 275
Diagnosis, 8, 37, 48, 106, 306, 321, 322, 326, 328-329, 330
 see also Digital images; Specialists
Differentiated services model (*diff-serv*), 12, 18, 45, 137-140, 143, 151, 251, 252-253
Digital Anatomist, 307-309
Digital divide, 14-15, 71, 125, 195, 209, 210, 211, 245-246
Digital images, 5, 8, 55(n.4), 65, 71, 76-79, 82, 134, 140, 236, 272-274, 300-301, 304-305, 320-321, 323, 328-329, 332
 bandwidth, 73-74, 77-79, 80, 123, 134, 136, 162, 164, 241, 244, 272, 288-289, 301-302
 biomedical research applications, 114-115, 239
 cost factors, 77, 79, 186

CT scans, 4, 65, 72, 78, 80, 134, 164, 238, 272, 274, 300
Digital Anatomist, 307-309
latency, 115-116
MRI, 4, 72, 77, 80, 113, 134, 238, 248, 272, 274, 284, 328
Next Generation Internet, 48-49, 305-306, 309, 317, 320-321, 323, 324, 326, 329, 331-333
Picture Archiving and Communications System (PACS), 76, 79, 80, 274
telemammography, 48, 78, 79, 142-143, 236, 244, 272-273, 324, 332-333
ultrasound technology, 77, 273, 284, 293, 294, 300, 306
X rays, 4, 36, 65, 74, 78, 80, 84, 113, 134, 164, 238, 244, 294
see also Video and teleconferencing
Digital Millennium Copyright Act, 217
Digital subscriber line services, 132, 164, 165, 166
Disasters, 18, 99-100, 239, 254
see also Terrorism
Discrete multitone, 164, 166
Diseases and disorders, general, *ix*, 30-31, 34, 57, 98, 182
surveillance, 3, 4, 8, 57, 94-99, 101-102, 238-239, 258, 312-313
see also Chronic illness; Risk factors
Distance education, 6-7, 214, 217-218, 238, 240-241, 249, 261, 266, 280-283, 289-291, 293-298, 306-307, 324
Distance medicine, *see* Telemedicine
Drugs, 5, 10, 32, 37, 38, 61, 182, 189, 190, 239
clinical trials, 120-121
data security, 11, 71
insurance claims for, 89
intellectual property, 168-169
online ordering, 60, 71, 189, 221
public health surveillance, 8, 97, 98

E

East Carolina University, 72-73, 76, 222, 284-293
E-Biomed, 116
E-commerce, 9, 10-11, 35, 60-61, 85, 179-182, 310
encryption, 63, 71
privacy and, 168
Economic factors, 13, 248, 298-299
market forces, 33-35, 185-186, 191-192
see also Cost and cost-benefit factors; E-commerce; Funding; Health care financing; Socioeconomic status; Underserved areas
Education, professional, *see* Professional education
Educational attainment, 14, 61, 195, 210
Elderly persons, 83
E-mail, 35, 169, 260-261, 286
anonymous, 169-170, 256
cost factors, 62, 63-64
encryption, 156-157, 167-168, 169, 300
latency, 39
patient/provider exchanges, 14, 22, 37, 62-64, 66, 70, 189, 192, 197, 248, 279, 300
Embryology, 199, 323, 330
Emergency care, 3, 11, 18, 26 (n. 1), 133, 210, 288, 324
data security as impediment to, 144, 159
disasters, 18, 99-100, 239, 254; *see also* Terrorism
patient records, 81, 83, 144, 159, 160, 206
patient transport, 321, 330
videoconferencing, 4-5
Encryption, 17, 30, 32, 40, 62, 63-65, 70, 79, 145, 147, 150-155, 325
authentication and, 151-153
bioterrorism response, 100
e-commerce, 63, 71
e-mail, 156-157, 167-168, 169, 300
firewalls, 149
home care, 68, 267
in administrative and financial transactions, 91, 172
MINDSCAPE, 302
patient record sharing, 82-83, 84, 292-293
Secure Socket Layer (SSL) encryption, 17, 63, 64, 154-155, 252
time factors, 150-151
European Union, 208-209
Evaluation issues, 15, 19-22, 193, 250, 257-261
biomedical research, 117-118
consumer health information, 61-62, 255-256
e-mail between patients and providers, 62

quality of information, 2, 14, 61-62, 246-247, 255-256, 292
scalability of protocols, 16-17
technical considerations, 38-41
see also Demonstration projects;
Standards
Expert systems, 227, 228-229
Extensible Markup Language (XML), 85, 86, 198, 331

F

FDA Modernization Act, 120
Federal Communications Commission, 211
Federal government, *ix*, 20-21, 24-25, 29, 96, 202-203, 211-212, 223-227, 254-260
National Aeronautics and Space Administration, 24, 46, 47-48, 49, 221, 250, 283-284, 299, 300
National Center for Biotechnology Information, 109, 111, 117
National Center for Microscopy and Imaging Research, 114-116
National Committee on Vital and Health Statistics, 25, 264, 265
National Institute of Standards and Technology, 46, 47
National Telecommunications and Information Administration, 212
see also Funding; Next Generation Internet; Regulatory issues; Legislation; *terms beginning* Department of
Financial issues, *see* Health care financing
Firewalls, 147, 148-150
For the Record: Protecting Electronic Health Information, 204
Foundation for Health Care Quality, 91, 310-311
Funding, 51, 186, 202-203, 223-227, 254, 287-288, 303
certification authorities, 156
clinical research, 120
Department of Health and Human Services, 20, 23, 257-258
distance learning and telemedicine, 214
Next Generation Internet, 46-48
payment mechanisms, 221
remote consultations, 36, 287-288, 301
underserved areas, 211-212, 213-215
workforce study, 262-263

G

Geographic factors, 125, 195
biomedical research, bandwidth, 110
clinical care, 71, 72
Kaiser-Permanente, 275-276
payment policies, 220
professional education, 102, 203; *see also* Distance learning
see also Home care; Remotely controlled equipment; Rural areas; Ubiquity of access; Underserved areas; Urban areas
Government role, 27
see also Federal government; Legislation; Public health; Public policy; Regulatory issues; State government

H

Hackers, 4-5, 11, 144, 145, 203, 240
anonymous e-mail, 169-170
firewalls, 147, 148-150
viruses, electronic, 145, 160
Health Alert Network, 99, 225
Health care financing, *ix*, 2, 57, 88-94, 125, 186-187
payment mechanisms, 2, 10, 24, 28, 32, 44, 89-92, 186, 194-195, 219-221, 235, 264
anonymous, 172
encryption, 91, 172
integrated services model, 140-141
patient monitoring, 66-67
secured transactions, 61, 190-191, 205, 242, 243
technical requirements, 4-5, 13, 90-91, 93-94, 238-239
see also Insurance; Medicaid; Medicare
Health Care Financing Administration, 13, 20-21, 89-90, 91-92, 219-221, 225, 258, 264, 286
see also Medicaid; Medicare
Health Care Personal Information Nondisclosure Act, 206
Health Data in the Information Age: Use, Disclosure, and Privacy, 204-205
Health Insurance Portability and Accountability Act, 86, 92, 185, 205, 208

Health Level Seven (HL7), 84, 85
Health maintenance organizations, 89, 103, 179, 186, 195, 196, 266-267
Health Plan Employer Data and Information Set (HEDIS), 90
Health plans; *see* Insurance; Managed care; Medicaid; Medicare
Health records, *see* Patient records
Home care, 57, 66-69, 75, 76, 128 (n. 7), 196, 209, 235, 240, 258, 287-288, 290
 bandwidth requirements, 67-68, 71, 267
 encryption requirements, 68, 267
 patient/provider e-mail exchanges, 14, 22, 37, 62-64, 66, 70, 189, 192, 197, 248, 279, 300
 technological issues, 66-68, 267
 time factors, 68, 70
 video, 129 (n. 9)
 vignettes, 29-32, 266-267
 see also Telemedicine
Hospitals, 13, 188-189, 287, 290
 consumer information, 5
 insurance claims for, 89, 195
 integrated delivery networks (IDNs), 35, 44
 patient record sharing, 81-82
 public health surveillance and, 97-98
 see also Emergency care; Surgery
Hypertext transfer protocol, 84, 147, 261
Hypertext markup language (HTML), 86, 104, 196, 290, 295, 296-297

I

Images, *see* Digital images
Indian Health Service, 20, 21, 24, 258, 259
Infrastructure, *see* Technological issues
Insurance, 5, 8, 11, 12, 35, 37, 44, 89-90, 186, 195, 239, 247
 data security, 11, 145, 241
 Health Insurance Portability and Accountability Act, 86, 92, 185, 205, 208
 payment mechanisms, 2, 10, 24, 28, 32, 44, 219-221
 spot market purchases, 191-192
 standards for claims and payments, 92-93, 186, 219-221
 see also Health maintenance organizations; Managed care; Medicaid; Medicare

Integrated Advanced Information Management Systems program, 300
Integrated delivery networks (IDNs), 35-36, 44, 79, 179
Integrated services digital networks (ISDN), 74, 287, 288, 289, 290, 304
Integrated services model (int-serv), 12, 17, 18, 45, 137, 140-141, 142, 251, 253
Intellectual property, 14, 24, 99, 168-169, 188-189, 196, 202, 204, 215-218
 access controls, 159
 biomedical research databases, 110-111, 249, 307
 educational materials, 14, 217-218, 249, 307
 pharmaceutical companies, 168-169
Internet Engineering Task Force, 12, 18, 45-46, 138, 141, 143, 153, 174, 252-253, 254
Internet Protocol (IP), 42, 49, 140, 161-162, 260-261, 291, 294, 305
Internet Protocol Security (IPSec), 17, 150, 153-154, 251-252
Internet 2, *ix*, 1, 50-51, 52

J

Johns Hopkins University, 60
Joint Photographic Experts Group (JPEG), 286
Joint Working Group of Telemedicine/Telehealth, 224

K

Kaiser-Permanente of Northern California, 184, 196, 274-280

L

Language factors, human, 119-120, 266-267
Latency, 9, 39-40, 123, 124, 134, 136, 240
 administrative and financial applications, 93
 biomedical research, 122
 clinical care, 81-83, 87-88, 124
 consumer health applications, 62, 69-70, 279
 defined, 10, 39

emergency care, 133
guaranteed, 134-135
imaging, 115-116
professional education, 107, 124
public health, 100
satellite telecommunications, 167
Legislation, 58, 91, 205-208
 Balanced Budget Acts, 92, 219, 220-221, 224
 Collections of Information Antipiracy Act, 218
 Communications Act of 1934, 211
 Comprehensive Telehealth Act, 224
 Digital Millennium Copyright Act, 217
 European Union, 208-209, 218
 FDA Modernization Act, 120
 Health Care Personal Information Nondisclosure Act, 206
 Health Insurance Portability and Accountability Act, 86, 92, 185, 205, 208
 Medical Information Privacy and Security Act, 206-207
 Medical Information Protection Act, 206
 Privacy Act, 207
 Public Health Service Act, 224
 Telecommunications Act of 1996, 211
 Triple-A Rural Health Improvement Act, 224
Licensing of health care providers, 14, 24, 105, 186, 189, 221-223, 224, 267
Litigation, malpractice, 14, 221-223, 263
Local area networks (LANs), 33, 67, 79, 80, 104, 137, 153, 154, 284
 firewalls, 147, 148-150
 images, 80, 301
 Internet Protocol Security, 153
Local communities, 38, 312
 bioterrorist response, 99-100
 health care claims and payments, 90
 professional education, 104
 public health surveillance, 96
Local multipoint distribution system technology, 166-167
Louisiana State University Medical Center, 317-318

M

Magnetic resonance imaging, 4, 72, 77, 80, 113, 134, 238, 248, 272, 274, 284, 328

Malpractice, 14, 221-223, 263
Mammography, 48, 78, 79, 142-143, 236, 244, 272-273, 324, 332-333
Managed care, 11, 35, 95, 179, 182, 183-184, 186, 190, 191, 195, 244, 247
 efficiency and cost-effectiveness efforts, 59, 89
 integrated delivery networks (IDNs), 35
 see also Health maintenance organizations
Medicaid, 13, 21, 89, 264
Medical Information Privacy and Security Act, 206-207
Medical Information Protection Act, 206
Medical records, *see* Patient records
Medicare, 13, 21, 89-90, 219-220, 223, 224, 225, 264
Medications, *see* Drugs
MEDLARS, 33, 55 (n. 5), 135
MEDLINE, 3, 33, 37, 55 (n. 5), 103-104, 109, 110, 215, 295, 301, 307
Medtronic Inc., 66
Microsoft Corp., 66, 210, 291
Microwave technology, 74, 166, 289, 290
MINDSCAPE, 300, 301-302, 306
MRI, *see* Magnetic resonance imaging
Multicast protocols, 7, 49, 50, 119, 143-144, 176, 241, 283, 284

N

National Aeronautics and Space Administration, 24, 46, 47-48, 49, 221, 250, 283-284, 299, 300
National Cancer Institute, 120-121, 273, 301
National Center for Biotechnology Information, 109, 111, 117
National Center for Microscopy and Imaging Research, 114-116
National Committee on Vital and Health Statistics, 25, 264, 265
National Emergency Medicine Information Extranet, 316
National Heart Attack Alert Program, 128 (n. 7)
National Health Alert Network, 99-100
National Institute of Standards and Technology, 46, 47
National Institutes of Health, 19, 20-21, 46, 47, 116, 120-121, 216, 223, 225, 226, 228-229, 255-257, 258, 264

National Laboratory for the Study of Rural Telemedicine, 74

National Library of Medicine (NLM), 15, 16, 18-19, 254, 255, 258, 263

- clinical transaction standards, 84
- demonstration project awards (NGI), 314-333

Integrated Advanced Information Management Systems program, 300

MEDLARS, 33, 55 (n. 5), 135

MEDLINE, 3, 33, 37, 103-104, 109, 110, 215, 295, 301, 307

Next Generation Internet, 46-47, 305-306, 314-333

- public health initiative, 95

National Museum of Health and Medicine, 330

National Network of Libraries of Medicine, 95

National Science and Technology Council, 46

National Science Foundation, 46, 47-49, 50-51, 111, 225, 226, 250

National Telecommunications and Information Administration, 212

Native Americans, Indian Health Service, 20, 21, 24, 258, 259

Naval Research Laboratory, 171

Network availability, *see* Availability; Ubiquity of access

Network Data Express, 309

Networked 3D Virtual Human Anatomy, 315

Neurology, 73, 74, 112, 131 (n. 38), 158, 272, 274, 284, 285, 289, 308-309, 319-320, 328

Next Generation Internet, *ix*, *x*, 1-2, 16-18, 28, 46-49, 52, 291, 305-306, 312, 318-333 (*passim*)

- authentication, 17-18
- bandwidth, 16, 49, 250-251
- clinical care, 48, 319
- data security, 17, 251-252
- digital images, 48-49, 305-306, 309, 317, 320-321, 323, 324, 326, 329, 331-333
- Internet 2 and, 50-51
- security, 17, 312, 317-318
- telemedicine, 48, 318-319, 321, 324, 327-328, 332-333

testbed networks, 2, 16, 17, 20-21, 28, 47, 250-252, 259, 323, 326-327

video and teleconferencing, 48-49, 306

O

Office of the Future, 293-294

Onion routing, 171-172

Open Systems Interconnection model, 146

Organizational factors, 13-14, 15-16, 25, 28-29, 37-38, 58, 178-201, 204-205, 236, 244-245, 247-248, 265, 310-311

- barriers, external, 185-187
- barriers, internal, 187-189, 194, 262
- clinical care, 182-184, 188, 189, 192, 195, 196
- consumer health, 9, 10-11, 179-182, 183, 191, 195
- costs, 59, 179-180, 184, 185, 186-187, 194-195, 196
- data security, 38, 43, 147, 159, 203-205, 253
- firewalls, 147, 148-150
- information systems staff requirements, 14, 23, 196-199, 227-230, 260, 261-263, 298-299
- integrated delivery networks (IDNs), 35-36
- Kaiser-Permanente, 184, 196, 274-280
- leadership, 197-199, 227
- patient records, 82-83, 182-183
- public health decision making, 98
- Web hosting, 162

see also Administrative support; Health care financing; Insurance; Managed care

Overprovisioning, 137, 138-139

P

Partners in Information Access, 95

PathMaster, 328-329

Patient records, 4-5, 10, 20, 34, 57, 64-65, 70, 71-72, 80-87, 94, 194-195, 197, 236, 262, 264, 301, 320-321, 325

- audit capabilities, 19, 83, 159, 207, 242, 256, 297, 302
- availability needs, 11, 30, 64-65, 124, 133
- e-mail, incorporation into, 64

- in emergency care, 81, 83, 144, 159, 160, 206
 - organizational factors, 82-83, 182-183
 - professional education and, 106
 - public health workers use of, 94
 - quality of service needs, 134, 164
 - standards, 84-86, 129 (n. 17), 186, 188
see also Data security; Privacy
 - Patient-Centered Access Secure Systems (PCASSO), 65
 - Pharmaceuticals, *see* Drugs
 - Physical fitness and exercise, *see* Exercise
 - Picture Archiving and Communications System (PACS), 76, 79, 80, 274
 - Pilot projects, *see* Demonstration projects
 - Platform for Internet Content Selection, 62
 - Policy issues, *see* Public policy
 - President's Information Technology Advisory Committee, 229
 - Pretty Good Privacy (PGP), 63, 150, 152, 156-157, 313
 - Preventive care, *see* Risk factors
 - Primary care physicians, 13, 30, 31, 61, 191, 220
 - certification authorities, 155-156
 - Prison telemedicine, 72, 76, 215, 286-287, 289, 290
 - Privacy, 2, 5, 15, 24, 28, 38, 40, 56, 61, 68, 153, 167-173, 186, 193, 205-206, 235, 264, 310
 - and home care, 68
 - in financial and administrative transactions, 90-91
 - patient records, 64, 65, 68, 82-84, 124, 242, 249, 256, 262, 279, 292-293, 300, 310, 316-317, 324-325, 332-333
 - encryption, 82-83, 84, 292-293
 - public policy, 203-209, 212, 214, 262
 - see also* Pretty Good Privacy; Security
 - Privacy Enhanced Mail, 152
 - Private networks, 34, 41, 43-45, 74, 154, 297
 - see also* Local area networks
 - Private sector, general, *ix*, 50-51, 264
 - biomedical research databases, 109-110
 - demonstration projects, 20, 260-261
 - protocols, 21-22
 - regulatory issues, 14, 25, 205, 207-208
 - see also* Insurance; Internet 2; Managed care
 - Professional education, 2, 3, 6-7, 14, 15, 22-23, 28, 57, 102-108, 240-241, 261-263, 315-316
 - and intellectual property, 14, 217-218, 249, 307
 - availability, 107-108, 241
 - bandwidth needs, 104, 105, 107, 123, 126, 241, 281, 282, 290-291, 297-298
 - continuing education, 105-107, 322
 - cost factors, 282-283
 - data security for, 108
 - distance education, 6-7, 214, 217-218, 238, 240-241, 249, 261, 266, 280-283, 289-291, 293-298, 306-307, 324
 - graduate education, 102-105, 322
 - information systems staffs, 14, 23, 227-230, 261-263
 - latency needs, 107, 124
 - Next Generation Internet, 48-49, 318, 319, 322
 - remotely controlled apparatus, 113-114
 - simulations, 103, 104-105, 123, 124, 250, 293-294, 320, 329
 - ubiquity of access, 108, 241
 - video and teleconferencing, 6-7, 104, 106, 240-241, 266, 273, 280-283, 290, 291, 293-294, 297-298, 315, 325
- Proprietary information, *see* Intellectual property
- Protocols, 21-22, 34, 42, 45, 55-56 (n. 15), 244, 260-261
 - authentication, 150, 151-156
 - demonstration projects, scalability, 16-17
 - hypertext transfer protocol, 84, 147, 261
 - Internet Protocol (IP), 42, 49, 140, 161-162, 260-261, 291, 294, 305
 - Internet Protocol Security (IPSec), 17, 150, 153-154, 251-252
 - latency and, 39-40
 - multicast, 7, 49, 50, 119, 143-144, 176, 241, 283, 284
 - quality of service protocols, 12, 18, 45, 132, 135-141, 142, 143, 251, 252-253
 - scalability of, 16-17, 32, 48, 72, 140, 252-253, 320-321

- security, 17, 49, 63, 150-157, 251-252; *see also* Encryption
- Transmission Control Protocol, 39-40, 42, 136-137, 260-261
- Public health, *ix*, 1, 2, 27, 36-37, 58, 94-102 and data security, 98-99, 101
- Centers for Disease Control and Prevention, 95, 96, 99-100, 225, 264, 312-313
- cost savings, 97-98, 101-102
- privacy and, 169
- standards for, 101-102, 186
- state government and, 95, 96, 97
- surveillance, 3, 4, 8, 57, 94-99, 101-102, 238-239, 258, 312-313
- technical requirements for, 4-5, 100-102, 238-239
- ubiquity of access, 101-102
- video and teleconferencing needs, 4-5, 238-239, 100
- Public Health Service Act, 224
- Public key infrastructure, 17, 63, 64, 152-156, 252, 264, 275, 276-277
- Public policy, 14-15, 23-25, 29, 202-234, 236, 248-249, 263-265
 - consumer health security, 70
 - data security, 151-152, 203-209, 212, 214, 225, 264
 - privacy, 203-209, 212, 214, 262
 - see also* Federal government; Intellectual property; Legislation; Regulatory issues; State government; *cross-references under* Demographic Factors
- PubMed Central, 216

Q

- Quality of health care, general, 28, 191, 236
- Quality of information, 2, 14, 61-62, 246-247, 255-256, 292
- Quality of service (QOS), 10, 12, 15, 16-17, 19, 41, 45-46, 51, 132, 133-144, 173, 235-236, 237, 243-245, 254, 256-257, 259, 293, 297-298, 331
 - availability of network, 160
 - defined, 10, 41, 55 (n. 14)
 - guarantees, 18, 41, 134-135, 140, 141, 142, 253-254, 256-257
 - Next Generation Internet, 46, 49, 251
 - overprovisioning, 137, 138-139
 - patient records, 134, 164
 - private networks *vs* Internet, 43, 44, 45
 - private sector, other, 50
 - protocols, 132, 134-144
 - bandwidth, 132, 135-138
 - diff-serv/int-serv, 12, 18, 45, 137-141, 142, 143, 251, 252-253
 - QOS policy, 141-143
 - recommendations related to, 16-18, 251-253
 - research needed, 17-18, 252-253
 - video and teleconferencing, 133, 134
 - virtual overlay networks, 141
 - see also* Differentiated services model (diff-serv); Evaluation issues; Integrated services model (int-serv); Standards

R

- Realizing the Information Future: The Internet and Beyond*, 204
- Regence BlueShield, 309-313
- Regulatory issues, 186, 202, 219-223, 248-249, 263
 - insurance, 56 (n. 16)
 - licensing of health care providers, 14, 24, 105, 186, 189, 221-223, 224, 249, 267
 - malpractice, 14, 221-223, 263
 - privacy and security, 14, 25, 205, 207-208
 - public health surveillance and, 98
 - see also* Intellectual property; Legislation; Standards
- Remotely controlled equipment, 3, 4-7, 20, 36, 39, 68-69, 70, 75-76, 113-116, 121, 128 (n. 7), 133, 209, 235, 236, 240, 243, 247, 253-254, 266-267, 331
 - Next Generation Internet, 46, 48-49, 331
- Remote, Real-time Simulation for Teaching Human Anatomy and Surgery, 320
- Research, 15, 29, 174, 223-227
 - Department of Health and Human Services, 24-25
 - National Institutes of Health, 19
 - National Library of Medicine role, 18-19

network availability, 161-162
see also Biomedical research;
Demonstration projects;
Evaluation issues; Funding;
Internet 2; Next Generation
Internet
Risk factors, 3, 31, 34, 58, 59, 60, 95-96, 274,
288
Robert Wood Johnson Foundation, 156, 310
Rural areas, *ix*, 6, 48, 72, 74, 77, 96, 103, 104,
120, 195, 209, 210, 211-212, 213-
214, 215, 220, 222, 224, 236, 240,
245-246, 300, 315, 321
Rural Health Science Education, 315

S

Satellite telecommunications, 166-167, 300
see also National Aeronautics and Space
Administration
Scalability, 16-17, 32, 48, 72, 140, 252-253,
320-321
see also Bandwidth
Science Applications International Corp., 65
Seattle/Pacific Northwest GigaPOP, 305
Secure/Multipurpose Internet Mail
Extension (S/MIME), 157
Secure Radiologic Collaboration on the
Next Generation Internet, 318
Secure Socket Layer (SSL) encryption, 17,
63, 64, 154-155, 252
Security, *ix*, 9-11, 15, 17, 22, 28, 29, 40, 45,
46, 123, 124, 136, 144-162, 167-
173, 186, 193, 235, 237-243, 262,
310, 322-323
access controls, 64-65, 145, 157-160, 242-
243, 256, 302
attribution/nonrepudiation, 146
authentication, 17-18, 63-65, 70, 145,
253, 279, 297
confidential e-mail, 156-157
cost factors, 144, 279
defined, 10
firewalls, 147, 148-150
for digital images, 79
for patient records, 64, 65, 68, 92-84,
124, 242, 249, 256, 262, 279, 292-
293, 300, 310, 316-317, 324-325,
332-333
encryption, 82-83, 84, 292-293
in biomedical research, 110-111, 122,
159, 316
in clinical care, 65, 68, 82-84, 88; *see also*
for patient records *supra*
in clinical trials, 121
in consumer applications, 61, 63, 64-65,
68, 70-71, 238-239
in e-mail between patients and
providers, 63, 66, 300
in emergency care, 144, 159
in financial and administrative
transactions, 61, 90-92, 94, 145,
172, 190-191, 205, 241, 242, 243
in home care, 68
in private networks, 34, 41, 43-44, 74,
154, 297
in professional education, 108
in responding to bioterrorist attacks, 99-
100
information warfare, 5
Internet Protocol Security (IPSec), 17,
150, 153-154
network availability, 6, 9-11, 160-162
Next Generation Internet, 49
on consumer Web sites, 61, 64-65, 249
organizational factors, 38, 42, 147, 159,
203-205, 253; *see also* firewalls
supra
perimeter control, 146
protocols, 32, 48, 72, 140, 252-253, 320-
321
public health information, 98-99, 101
public key infrastructure, 17, 63, 64,
152-156, 252, 264, 275, 276-277
public policy, 151-152, 203-209, 212, 214,
225, 264
recommendations regarding, 16-17, 251-
253
regulatory issues, 14, 25, 205, 207-208
rural areas, 48
Secure Socket Layer (SSL) encryption,
17
Transport layer security, 154-156
see also Authentication; Encryption;
Intellectual property; Privacy
SHINE project, 130 (n. 31)
Simulations
biomedical research, 112-113, 118, 320
professional education, 103, 104-105,
123, 124, 250, 293-294, 320, 329
Smart cards, 17-18, 30

Socioeconomic status, 95
 digital divide, 14-15, 71, 125, 195, 209,
 210, 211, 245-246

Southern Governors' Association, 223

Specialists, 8, 11, 13, 76-77, 82, 182, 184, 191
 consultations, 6, 8-9, 16, 17, 18, 20, 30-
 33, 36, 57, 71, 72-76, 88, 220-221,
 228-229, 236-237, 238-239, 244,
 247, 258, 266-267, 280, 285-286,
 288-289, 292, 323
 see also specific specialties

Standards, 186, 188, 194, 196, 239, 280, 286,
 292, 310
 clinical practice standards/guidelines,
 5, 14, 22, 57, 238-239, 275, 278
 discrete multitone, 164, 166
 distance education, intellectual
 property, 218
 insurance claims and payments, 92-93,
 186, 219-221
 patient records, 84-86, 129 (n. 17), 186,
 188
 public health, 101-102, 186
 quality of service guarantees, 18, 41,
 134-135, 140, 141, 142, 252-253,
 256-257
 see also Protocols; Regulatory issues

Stanford University, 130, 131, 280-283,
 329

State government, 317-318
 insurance regulation, 56(n.16)
 licensing of health care providers, 14,
 24, 105, 186, 189, 221-223, 224,
 249, 267
 malpractice, 14, 221-223, 263
 public health surveillance, 95, 96, 97
 Washington State Laboratory Reporting
 Project, 312-313
 Washington, Wyoming, Alaska,
 Montana, and Idaho (WWAMI)
 Rural Telemedicine Network,
 300, 302-303, 306-307
 see also terms beginning University of
 SUPERNET, 49

Support groups, *ix*, 5, 60, 61, 196, 239
 see also Chat groups

Surgery, 3, 4-5, 8, 36, 71, 105, 123, 124,
 239, 244, 250, 294, 316, 320, 323,
 329

Synchronous Optical Network equipment,
 135, 146

T

Technology Opportunities Program, 212,
 213-214

Teleconferencing, *see* Video and
 teleconferencing

Telemammography, 48, 78, 79, 142-143,
 236, 244, 272-273, 324, 332-333

Telemedicine, 21, 47, 52-53, 56, 69, 71, 186,
 189, 197, 209, 224, 235, 236, 243,
 248-249, 259, 300, 321
 chat groups, *ix*, 5, 60, 191, 192-193, 196,
 239, 261, 290
 Comprehensive Telehealth Act, 224
 costs, 184, 204, 287, 302, 303
 defined, 56
 funding, 214
 licensing of health care providers, 14,
 24, 105, 186, 189, 221-223, 224,
 249, 267
 malpractice, 14, 221-223, 263
 Next Generation Internet, 48, 318-319,
 321, 324, 327-328, 332-333
 payment policies, 219-220
 prisons, 72, 76, 215, 286-287, 289, 290
 private sector (Internet 2), 50
 remote consultations, 6, 8-9, 16, 17, 18,
 20, 30-33, 36, 57, 71, 72-76, 88,
 220-221, 236-237, 238-239, 244,
 247, 258, 266-267, 280, 285-286,
 288-289, 292, 323
 site visits by committee, 280, 284-289,
 290
 support groups, *ix*, 5, 60, 61, 196, 239
 telemammography, 47, 78, 79, 142-143,
 236, 244, 272-273, 324, 332-333
 ultrasound technology, 77, 273, 284,
 293, 294, 300, 306
 vignettes, 29-32, 266-267
 Washington, Wyoming, Alaska,
 Montana, and Idaho (WWAMI)
 Rural Telemedicine Network,
 300, 302-303
 see also Home care; Remotely controlled
 equipment

*Telemedicine: A Guide to Assessing
 Telecommunications in Health Care*,
 204

Telequest, 77

Teletrauma and the NGI, 324

Terrorism, 4-5, 99, 100, 101, 225, 239

Testbed networks, 16, 17, 20-21, 28, 47, 250-252, 259, 283-284, 323, 326-327

3Com Corp., 210

Time factors, 11, 28, 184, 240

- bandwidth, impacts of increases in
 - Internet use, 136
- continuing medical education, 107
- e-mail responses, 22, 63
- encryption/decryption, 150-151
- home care, 68, 70
- image transmission, 78
- see also* Availability; Bandwidth; Latency

Token-based authentication, 17-18, 83, 84

Transmission Control Protocol, 39-40, 42, 136-137, 260-261

Transport layer security, 150, 154-156

Transport of patients, 321, 330

Triple-A Rural Health Improvement Act, 224

U

Ubiquity of access, 9, 12, 15, 38, 40-41, 71, 88, 124-126, 195, 235, 239, 245-246

- administrative and financial applications, 94
- biomedical research, 122-123
- defined, 10, 40
- digital divide, 14-15, 71, 125, 195, 209, 210, 211, 245-246
- patient records, 65
- professional education, 108, 241
- public health, 101-102
- underserved areas, *ix*, 14, 210-212; *see also* Rural areas
- see also* Bandwidth

Ultrasound technology, 77, 273, 284, 293, 294, 300, 306

UNCLE system, 295, 297

Underserved areas, *ix*, 14, 72, 210-212

Unified medical language system, 84-85

Universal Resource Locators (URLs), 170-171, 295, 298

Universal Service Administration Company, 211-212, 213

Universal Service Fund, 211, 215

University Consortium for Advanced Internet Development, 46, 50-51

University of California, 65, 76, 77, 272-274

University of Iowa, 74

University of North Carolina, 293-298

University of Washington, 298-309

Urban areas, 72, 96, 103, 120, 176, 195, 210-211, 220, 245-246

URLs, *see* Universal Resource Locators

V

Very high performance Backbone Network Service (vBNS), 49, 51, 115, 116, 284, 294, 299

Video and teleconferencing, 3, 4-7, 71, 74, 75, 129(n.10), 179, 220, 238-239, 240-241, 273, 304-305, 315, 321, 326

- bandwidth, 11, 73-74, 100, 132-133, 241, 244, 290
- Biomedical Teleimmersion, 316
- bioterrorism response, 100
- consumer use of, 30-31, 266, 288
- home care, 129(n.9)
- image transmission and, 78, 115-116
- Next Generation Internet, 48-49, 306
- professional education, 6-7, 104, 106, 240-241, 266, 273, 280-283, 290, 291, 293-294, 297-298, 315, 325
- public health, 4-5, 238-239, 100
- quality of service, 133, 134
- see also* Digital images; Simulations

Viruses, electronic, 145, 160

see also Firewalls; Hackers

W

Washington, Wyoming, Alaska, Montana, and Idaho (WWAMI) Rural Telemedicine Network, 300, 302-303, 306-307

Wavelength division multiplexing technology, 135

WebTV, 210, 291

Western Governors' Association, 222

World Wide Web Consortium, 62

World Wide Web Electronic Medical Record System, 81-83

X

X rays, 4, 36, 65, 74, 78, 80, 84, 113, 134, 164, 238, 244, 294

XML, *see* Extensible Markup Language