IPng
Internet Protocol Next Generation

Edited by
Scott O. Bradner
Allison Mankin
Co-Chairs of IETF/IPng Process

Foreword by David Clark
MIT Laboratory for Computer Science

IPng

Internet Protocol Next Generation

Edited by

Scott O. Bradner
Allison Mankin

Foreword by David Clark

ADDISON-WESLEY PUBLISHING COMPANY
Reading, Massachusetts Menlo Park, California New York
Don Mills, Ontario Wokingham, England Amsterdam Bonn
Sydney Singapore Tokyo Madrid San Juan
Paris Seoul Milan Mexico City Taipei

The publisher offers discounts on this book when ordered in quantity for special sales. For more information please contact:

Corporate & Professional Publishing Group
Addison-Wesley Publishing Company
One Jacob Way
Reading, Massachusetts 01867

Library of Congress Cataloging-in-Publication Data

IPng, Internet protocol next generation / edited by Scott O. Bradner and Allison Mankin.
p. cm
Includes bibliographical references and index.
ISBN 0-201-63395-7 (alk. paper)
1. Internet (computer network). 2. Computer network protocols.
I. Bradner, Scott O., 1944- . II. Mankin, Allison, 1955- .
TK5105.875.I57I67 1995
004.6'2--dc20 95-20944
CIP

ISBN 0-201-63395-7

Text printed on recycled paper
1 2 3 4 5 6 7 8 9 MA 98979695
First Printing, September 1995

IPng

Internet Protocol Next Generation

Addison-Wesley IPng Series

Scott O. Bradner, Consulting Editor

The *Addison-Wesley IPng Series* provides networking and communications practitioners with vital information on the next generation technology being developed for the growth of the Internet. Technical and corporate managers, information providers, and all those with an interest in harnessing the new technology will find these comprehensive reference books invaluable. Scott O. Bradner, a leader in the field, is the Consulting Editor for the IPng Series.

Foreword

Dave D. Clark
M.I.T. Laboratory for Computer Science

To the casual user, the Internet is characterized by its applications: for example the World Wide Web, email, teleconferencing, or remote login. But to the builders of the Internet, it is not the applications that define it, but the foundation on which they rest. What is unique about the Internet is that it provides a basis on which a large number of applications can be built, applications with very different needs for network service qualities such as bandwidth and latency.

In contrast to the telephone network, which was initially conceived to carry exactly one application, voice communications, the Internet from the beginning was intended to provide a general infrastructure on which a wide range of applications could operate. With respect to generality, the Internet is less like a traditional network and more like a programmable computer.

In fact, the design of the Internet envisions two sorts of generality at the same time. The ability to support a range of applications is critical, but so is the ability to operate over a range of network technologies. The Internet protocols can operate over local area networks, point to point links of various speeds, other packet switched networks such as frame relay or SMDS, wireless networks, and new sorts of emerging network technology such as ATM.

To see the importance of this generality, one need only ask what engine will drive the evolution of networking. It is the creation of new applications and services, which will excite users and trigger investment in those applications. Advances in applications drive demand from users, which in turn drive innovations in network technology, to extend performance, predictability and reliability, to reduce cost and to reach into new operating regimes such as wireless networking. This drive is evidenced by the development over the last few years of ATM, 100 megabit/second Ethernet, and wireless LANs. It must be a fundamental goal of any approach to networking to allow, and indeed to stimulate, advances in applications and in technology.

The spirit of innovation and creativity that brings new applications into existence needs an environment that is open, ubiquitous, and adaptable. The importance of these characteristics in network architecture is widely recognized. A report released last year by the National Research Council, titled "Realizing the Information Future: The Internet and Beyond," articulated the power of this generality: it leads to universal connectivity for users, a foundation for innovation in services and applications, and an evolving infrastructure that can grow and change as needed. In discussions about the concept of a National Information Infrastructure, there has been considerable agreement that at least some part of such an infrastructure must be organized to meet these needs of generality and openness. This model is a compelling vision of what a network must be if it is to grow, to evolve, to stimulate new applications and be stimulated in return, and to generate both economic and technological growth.

That NRC report described an approach to achieving these objectives called the "network independent bearer service", an abstract organization of protocols and interfaces not tied to any existing protocol suite. The bearer service represented an interface within the protocol layers that cross-connects the applications to the technology without letting them become directly dependent on each other. The two can evolve independently, but the application can be assured of the behavior that it can expect from the infrastructure below.

The abstract structure of the bearer service is closely matched in practice by the Internet protocol family. In the Internet protocols, it is the Internet protocol itself, sometimes called IP, that provides on the one hand the connection, and on the other hand the isolation, between the applications above and the technology options below. It represents and captures that central core of functions that the technology must support, and the application may depend on.

To propose an interface does not mean that it can actually be realized, but the Internet protocol defines a service definition that has proven successful. While there are a number of reasons for this success, perhaps the most important one is that the Internet protocol is simple. That is, the service it defines is as minimal as possible. At its heart, it defines only two things: the addressing plan and the packet delivery service. And simplicity is key—the less there is defined, the less there is to argue about.

The addresses, the numbers that are used to identify the machines attached to the Internet, are a critical point about which there must be common agreement. In simple terms, one cannot send a message to a destination unless one can identify it. And while there are continuous debates about the need for universal addresses, and what the syntax and structure should be, the basic answer is that a common address space is at the core of the Internet.

The other aspect of IP is the delivery service. What is the expectation that the user should have when a packet is handed to the Internet service for delivery? Again, the Internet protocol makes an assumption that is simple and minimal. The delivery service of the Internet, sometimes called "best effort" delivery, is that any packet handed to the Internet layer will be delivered to the destination as soon as possible, but with no specific commitment as to bandwidth, delay, or absolute reliability. The success of the Internet is that this very simple delivery service, which can be implemented over a very wide range of technology, can in fact support a wide range of applications. The Internet protocol, while very simple, is thus at the heart of the Internet and its success.

If the Internet protocol is so successful, why does it need to change? Exactly because we can now see the limits of these two core services. The problem with addressing is very compelling; the existing Internet addresses are not large enough to provide room to address all the hosts that will be attached to the Internet in the next few years. We must either move to larger addresses, or abandon the idea of universal addressability that has worked so well up to now. The problem with the best effort delivery model is that while it has proved eminently successful for a broad range of applications, there are certain applications for which it is not adequate, in particular real time delivery of data such as audio and video. Better support for these services implies a change to the structure of the IP protocol header. Finally, of course, over the period of the last 20 years, there are a number of other issues with the IP protocol that have come to light, which should be rectified as a part of a revision.

The major issue in revising the Internet protocol is to balance two concerns. One is to respond to real needs that are beyond the current protocol. The other is to control the natural desire to add new features. This desire is universal, and if not checked can lead to the creation of a complex, overloaded, and unsatisfactory system. In fact, the tendency is so well known in computer systems circles that it was given a name, the Second System Syndrome. In the case of IP, this tendency could be particularly problematical, since one of the goals of Internet was to operate over almost any sort of network technology. If the design moves away from the very simple best effort service, it will begin to limit the range of technology it can utilize. This would represent a tremendous loss of power for the protocol. The design of the next IP has had to deal with this tension.

There is another tension in the process of evolving the Internet protocol. There are two uses for the same protocol suite. One is the building of that global public network called the Internet. The other is building the many private networks that are currently based on the same protocols. These private networks, usually belonging to a corporation, are where a majority of the Internet technology sold today is used, and these networks have a different set of requirements than the public Internet. They are smaller, and less concerned with issues of scale, but they are more concerned with delivering a predictable service in support of critical applications. So there is to some extent a divergence of objectives. But the same products, implemented to the same standards, must work in both contexts. And this is the other major issue that the designers of a new Internet protocol had to balance.

Replacing the Internet protocol may be the most challenging step that the Internet Engineering Task Force has undertaken in its history. It will eventually change every host and every router in the Internet, it requires the balancing of a dauntingly broad set of requirements, and it will not succeed unless it supports and sustains the current success of the Internet, a success that, not surprisingly, is to most people more concerned with flights of fancy than foundations. But to those of us who believe that the future must sit on strong foundations, this effort is the key step towards tomorrow.

Preface

The Internet is about to become a victim of its own success.

The projection in the fall of 1991 that the Internet, with the capacity to support many millions of users, was beginning to run out of available network addresses was quite a bit of a surprise. The projection was quickly followed by numerous articles in the trade press announcing the imminent demise of what had been a promising technology.

Many groups were ready with alternatives to fix the problem. Network protocols from various official standards organizations and proprietary protocols from a number of vendors were presented as solutions that could provide a foundation for a glorious future of ubiquitous networking.

This speculation caused more than a little consternation within the Internet Engineering Task Force (IETF), the organization responsible for keeping the standards of the TCP/IP protocol suite used by the Internet. In the face of all this "end of the world as we know it" talk, the IETF felt it needed to determine just what was truly happening, and, if there was in fact a problem, what should be done about it.

When the initial investigation confirmed the basic diagnosis, the IETF undertook a multi-pronged effort to devise a replacement for the current version of the Internet protocol, IPv4. This effort sought not just to solve the immediate address limitation and scaling problems, but to look into the Internet's future and develop a protocol that would serve its needs for many years to come.

This book offers an inside view of the process the IETF used in its successful effort to define the issues, and provides an overview of the resulting Internet Protocol next generation (IPng). Along the way, the book reveals the rationale behind the structure and features of IPng, presenting numerous explorations of applications and technologies IPng could potentially support.

Audience. This book has been written so that it can be easily understood by anyone with a basic understanding of networking and communications. Those who would benefit from this book include: managers of technical organizations, networking professionals, technology watchers, those with a stake in the growing on-line commerce industry, and anyone with an interest in the Internet prototype of the Information Superhighway.

Organization.

Part I provides the background on the issues and problems facing the Internet. Part II describes the process which the IETF used to develop the new protocol; and Part III examines the all-important time frame for developing IPng.

Next, the book turns to the outside perspective of the wider networking community, with contributions from numerous industry experts.

Part IV explores the potential role of IPng in the future of communications, and Part V, the innovative technologies IPng should consider embracing.

Part VI contains the technical criteria for judging IPng proposals - culled from all of the preceding discussions, issues, and contributive perspectives.

The IPng proposals are presented and evaluated against the technical criteria in Part VII.

All of the preceding sections culminate in Part VIII, the overview of the selected IPng proposal, the new IPv6 Internet protocol.

In Part IX, the critical issue of security is discussed, and in Part X, the ongoing process of developing the protocol in greater detail is outlined.

Sources. Much of the material in this book has been adapted from the Internet standard and documentation series known as Request for Comments (RFCs). The material has been reworked, with the authors' assistance, to make it more accessible to a general audience while retaining the technical detail inherent in the original work. Also included are a number of new pieces written specifically for this book.

Acknowledgements. Reaching this stage of the recommendation would not have been even vaguely possible without the efforts of many people. In particular, the work of the IPng Directorate, Frank Kastenholz and Craig Partridge (the authors of the Criteria document) along with Jon Crowcroft (who co-chaired the ngreq BOF) was critical. The work and cooperation of the chairs, members, and document authors of the three IPng proposal working groups, the ALE Working Group and the TACIT Working Group laid the groundwork upon which this recommendation sits.

We would also like to thank the many people who took the time to respond to RFC-1550 and who provided the broad understanding of the many requirements of data networking that any proposal for an IPng must address.

The members of the IESG, the IAB, and the always active participants in the various mailing lists provided us with many insights into the issues we faced. Many other individuals gave us sometimes spirited but always useful counsel during this process. They include (in no particular order) Radia Perlman, Noel Chiappa, Peter Ford, Dave Crocker, Tony Li, Dave Piscitello, Vint Cerf, and Dan Lynch.

Thanks to David Williams and Cheryl Chapman who, along with the very hard-working Addison-Wesley technical editors: Abigail Cooper and Kate Habib, took on the occasionally impossible task of ensuring that what is written here resembles English to some degree.

This book would have never happened without the perseverance and astonishingly good humor in the face of changing realities of Carol Long, our Executive Editor at Addison-Wesley.

To all of the many people mentioned above and those we have skipped in our forgetfulness, thank you for making this task doable.

Contributors' Biographies

R. Brian Adamson works at the Naval Research Laboratory (NRL) in Washington, DC. He is currently involved in NATO and DoD demonstration projects, which provide integrated communication services over mobile radio networks through the application of internetwork protocols and advanced service multiplexing techniques. To meet the goals of these projects, he has also written an experimental low data rate Internet voice communication software application.

Steven M. Bellovin works at AT&T Bell Laboratories, where he does research in networks, security, and why the two don't get along. While a graduate student at the University of North Carolina at Chapel Hill, he helped create netnews; for this, he and the other perpetrators were awarded the 1995 Usenix Lifetime Achievement Award. He is the co-author of the recent book *Firewalls and Internet Security: Repelling the Wily Hacker* and is currently focusing on how to write systems that are inherently more secure.

Jim Bound is a member of the Technical Staff in the Network Integration Software Engineering Group at Digital Equipment Corporation, and was formerly Digital's TCP/IP Engineering Program Manager in Networks Engineering. One of the members of the Internet Engineering Task Force (IETF) IPng Directorate, Mr. Bound is the IPv6 Technical Leader and one of the advanced development implementors building a prototype to test these emerging IETF specifications. In addition, he is co-author of several IETF IPv6 specifications in progress. An IEEE Computer Society Member, Mr. Bound worked on the IEEE POSIX 1003 System Services (1003.1), Test Methods (1003.3), and Real Time (1003.4) standard committees.

Scott Bradner, co-director of the Internet Engineering Task Force (IETF) IPng effort, is a senior technical consultant at the Harvard Office of Information Technology, Network Service Division, where he works on the design and development of network-based applications and manages the Network Device Test Lab. In addition, he is the co-director of the Operational Requirements Area in the IETF, a member of the IESG, and an elected trustee of the Internet Society. Mr. Bradner has been involved in the design, operation and use of data networks at Harvard University since the early days of the ARPANET. He was involved in the design of the Harvard High-Speed Data Network (HSDN), the Longwood Medical Area network (LMAnet), and NEARNET. He is currently chair of the technical committees of LMAnet, NEARNET and CoREN.

Christina Brazdziunas is currently a software engineer developing Cellular Digital Packet Data (CDPD), a wireless data technology, at Evolving Systems, Inc. She has had extensive experience as a researcher in ATM technology at Bellcore, where she was responsible for consulting with major telecommunication industry clients on IP and ATM network integration.

Edward Britton became involved in TCP/IP at the U.S. Defense Communications Agency. Since joining IBM in 1981, he has contributed to SNA architecture, OSI system design, telephony, telecommunications cross-product design, and many aspects of IBM's TCP/IP products. He has expertise in IPng, security, performance, wireless and Asynchronous Transfer Mode communications, and TCP/IP's relationships with other protocol suites.

Ross Callon is a consulting engineer at Bay Networks Incorporated in Billerica, Massachusetts, and has more than 15 years' experience in data communications. He was the original proposer of a dual stack transition scheme and of TUBA, as well as co-author of the NSAP Guidelines standard in the IETF. He is co-chair of the IETF IPng Working Group, and was a member of the IPng Directorate. He is also a regular contributor to the ATM Forum, and has been involved in the ATM PNNI routing and multi-protocol over ATM efforts.

Brian E. Carpenter has been Group Leader of the Communications Systems group at CERN since 1985. His previous experience includes ten years in software for process control systems at CERN and three years teaching undergraduate computer science at Massey University in New Zealand. Mr. Carpenter is a member of the Internet Architecture Board (IAB) and an active participant in the Internet Engineering Task Force (IETF).

J. Noel Chiappa is currently an independent researcher in the area of computer networks and system software. He has been a member of the TCP/IP technical community since 1977. While a member of the Research Staff at M.I.T., he worked on packet switching and local area networks, and was responsible for the invention of the multi-protocol router; he later worked with a number of companies to bring networking products based on work done at M.I.T. to the marketplace.

David Clark has worked at the MIT Laboratory for Computer Science since 1973, where he is currently a Senior Research Scientist. His research interests include networks, network protocols, operating systems, distributed systems and computer and communications security. Since the mid 70s, he has been involved in the development of the Internet. From 1981-1989 he acted as Chief Protocol Architect for the Internet, and chaired the Internet Activities Board. His current research includes extensions to the Internet protocols for advanced service requirements and security.

John Curran is chief technical officer for BBN Planet, where he is responsible for the company's strategic technology and business development initiatives, including establishing the overall direction for products and services. He leads design activities for BBN Planet's service infrastructure with a particular emphasis on scaling, management, and security issues, as well as providing consulting for key BBN ISC clients. Mr. Curran is an active member of the Internet Engineering Task Force (IETF) and the Internet Engineering Planning Group (IEPG), working on a wide range of topics including inter-provider coordination and IPv6.

Stephen E. Deering is currently a member of the research staff at the Xerox Palo Alto Research Center (PARC), where he works on multicast routing and such other advanced internetworking topics as mobile communications, scalable addressing, and support for multimedia applications over the Internet. He has been studying, designing and implementing computer communication protocols since 1978, including work on X.25 software for connecting to public networks, on the first implementation of the X.400 protocol suite for email, and on high-performance transport protocols for distributed systems. He has chaired numerous working groups in the Internet Engineering Task Force (IETF) and co-founded the Internet Multicast Backbone (the "MBone") group.

Deborah Estrin is an Associate Professor of Computer Science at the University of Southern California in Los Angeles. Her research has focused on the design of network and routing protocols for very large global networks. She is a co-PI on the NSF Routing Arbiter project, co-chairs the Internet Engineering Task Force's (IETF) Source Demand Routing Working Group, and is a primary participant in the Inter-Domain Multicast Routing and RSVP Working Groups.

Eric Fleischman is a Senior Principal Scientist within Boeing's Information and Support Services organization. He is the program manager for the company's Data Communications Architecture group, which establishes Boeing's tactical and strategic data communications architecture and technical standards, and oversees the many resulting projects that seek to achieve the architecture. He is also the project manager for Boeing's Virtual Collocation Project, which integrates collaborative authoring/engineering and real-time desktop video with various asynchronous collaborative problem-solving approaches. In addition, Mr. Fleischman is Boeing's representative to the Internet Engineering Task Force (IETF).

Antonia Ghiselli is a network researcher at INFN, the Italian National Institute for Nuclear Physics. She has been working since 1980 in setting up, managing, and developing software for the INFN network. She has worked at CERN as a scientific associate and is now coordinating and planning the INFN network and the Italian Research Network.

Robert E. Gilligan is a Senior Staff Engineer at Sun Microsystems, Inc., where he is responsible for developing the TCP/IP software in Solaris 2. Prior to joining Sun, he worked on Internet and Packet Radio projects at SRI International. Mr. Gilligan is also an active participant in the Internet Engineering Task Force (IETF).

Daniel Green is a staff scientist specializing in the area of tactical networking at the Naval Surface Warfare Center-Dahlgren Division. He has 37 years of experience with the Navy in various aspects of computer information technology. Currently he is technical co-chair of the U.S. Navy's High Performance Network (HPN) project, which is concerned with advanced networks for Navy aircraft, ships, and submarines.

Phill Gross is Director of Internet Engineering at MCI, where his group built the IP backbone that underlies the new internetMCI family of services. Previously, Mr. Gross was Vice President for Network Technology at Advanced Network and Services (ANS), Inc. Before that, he worked at the Corporation for National Research Initiatives (CNRI) on various Internet projects, where he organized and chaired the Federal Engineering Planning Group (FEPG) of the U.S. Federal Networking Council (FNC) and the Intercontinental Engineering Planning Group (IEPG). One of the founders of the Internet Engineering Task Force (IETF), Mr. Gross chaired the IETF from 1987 to 1994, during which time he formed the Internet Engineering Steering Group (IESG), co-chaired the IETF Routing and

Addressing (ROAD) group, formed the IPng Area, and wrote the original IPng direction statement. He is also a member of the Internet Architecture Board (IAB).

Denise Heagerty is a member of the External Networking Section at CERN, Geneva, Switzerland. In addition to her operational duties, she has been responsible for the DECnet/OSI transition and implementing CERN's electronic mail gateway strategy. Her previous experience covers telephone switching and real-time control systems.

Robert M. Hinden, Director of Software at Ipsilon Networks, a Mountain View, California startup, was formerly responsible for the Internet Engineering group at Sun Microsystems where he won the Sun Microsystems Presidents Award. He is the document editor for IP Next Generation, working in the Internet Engineering Task Force (IETF), and is a member of the IPng Directorate. He was formerly area director for routing in the Internet Engineering Steering Group (IESG) and chair of a number of IETF working groups.

Christian Huitema has conducted research in network protocols and network applications for a number of years. Currently at INRIA in Sophia-Antipolis, he leads the RODEO project, which researches and defines communication protocols for very high-speed networks, at one Gigabit or more. Mr. Huitema has been a member of the Internet Architecture Board (IAB) since 1991 and the organization's chairman since 1993.

Phil Irey is a Computer Scientist at the Naval Surface Warfare Center-Dahlgren Division. He has specialized in the application of communications protocols (particularly Transport layer protocols) to tactical applications, and recently has become involved in the instrumentation of communications subsystems. In addition, Mr. Irey has participated in the U.S. Navy's SAFENET and High Performance Network (HPN) projects.

Frank Kastenholz is currently an engineer with FTP Software, INC., the leading vendor of TCP/IP software for PCs, where he is responsible for developing advanced applications and network protocols. He is also a long-time Internet Engineering Task Force (IETF) participant and is currently on the Internet Engineering Steering Group (IESG) as Co-Area Director for the Internet Area.

Tony Li is currently a Technical Lead for Cisco Systems, Inc. where he is a bit slinger specializing in IP exterior routing protocols, high-performance switching, and fixing really nasty bugs. Mr. Li is currently