

Network Working Group
Request for Comments: 1462
FYI: 20

E. Krol
University of Illinois
E. Hoffman
Merit Network, Inc.
May 1993

FYI on "What is the Internet?"

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

This FYI RFC answers the question, "What is the Internet?" and is produced by the User Services Working Group of the Internet Engineering Task Force (IETF). Containing a modified chapter from Ed Krol's 1992 book, "The Whole Internet User's Guide and Catalog," the paper covers the Internet's definition, history, administration, protocols, financing, and current issues such as growth, commercialization, and privatization.

Introduction

A commonly asked question is "What is the Internet?" The reason such a question gets asked so often is because there's no agreed upon answer that neatly sums up the Internet. The Internet can be thought about in relation to its common protocols, as a physical collection of routers and circuits, as a set of shared resources, or even as an attitude about interconnecting and intercommunication. Some common definitions given in the past include:

- * a network of networks based on the TCP/IP protocols,
- * a community of people who use and develop those networks,
- * a collection of resources that can be reached from those networks.

Today's Internet is a global resource connecting millions of users that began as an experiment over 20 years ago by the U.S. Department of Defense. While the networks that make up the Internet are based on a standard set of protocols (a mutually agreed upon method of communication between parties), the Internet also has gateways to networks and services that are based on other protocols.

To help answer the question more completely, the rest of this paper contains an updated second chapter from "The Whole Internet User's Guide and Catalog" by Ed Krol (1992) that gives a more thorough explanation. (The excerpt is published through the gracious permission of the publisher, O'Reilly & Associates, Inc.)

The Internet (excerpt from "The Whole Internet User's Guide and Catalog")

The Internet was born about 20 years ago, trying to connect together a U.S. Defense Department network called the ARPAnet and various other radio and satellite networks. The ARPAnet was an experimental network designed to support military research--in particular, research about how to build networks that could withstand partial outages (like bomb attacks) and still function. (Think about this when I describe how the network works; it may give you some insight into the design of the Internet.) In the ARPAnet model, communication always occurs between a source and a destination computer. The network itself is assumed to be unreliable; any portion of the network could disappear at any moment (pick your favorite catastrophe--these days backhoes cutting cables are more of a threat than bombs). It was designed to require the minimum of information from the computer clients. To send a message on the network, a computer only had to put its data in an envelope, called an Internet Protocol (IP) packet, and "address" the packets correctly. The communicating computers--not the network itself--were also given the responsibility to ensure that the communication was accomplished. The philosophy was that every computer on the network could talk, as a peer, with any other computer.

These decisions may sound odd, like the assumption of an "unreliable" network, but history has proven that most of them were reasonably correct. Although the Organization for International Standardization (ISO) was spending years designing the ultimate standard for computer networking, people could not wait. Internet developers in the US, UK and Scandinavia, responding to market pressures, began to put their IP software on every conceivable type of computer. It became the only practical method for computers from different manufacturers to communicate. This was attractive to the government and universities, which didn't have policies saying that all computers must be bought from the same vendor. Everyone bought whichever computer they liked, and expected the computers to work together over the network.

At about the same time as the Internet was coming into being, Ethernet local area networks ("LANs") were developed. This technology matured quietly, until desktop workstations became available around 1983. Most of these workstations came with Berkeley UNIX, which included IP networking software. This created a new demand: rather

than connecting to a single large timesharing computer per site, organizations wanted to connect the ARPAnet to their entire local network. This would allow all the computers on that LAN to access ARPAnet facilities. About the same time, other organizations started building their own networks using the same communications protocols as the ARPAnet: namely, IP and its relatives. It became obvious that if these networks could talk together, users on one network could communicate with those on another; everyone would benefit.

One of the most important of these newer networks was the NSFNET, commissioned by the National Science Foundation (NSF), an agency of the U.S. government. In the late 80's the NSF created five supercomputer centers. Up to this point, the world's fastest computers had only been available to weapons developers and a few researchers from very large corporations. By creating supercomputer centers, the NSF was making these resources available for any scholarly research. Only five centers were created because they were so expensive--so they had to be shared. This created a communications problem: they needed a way to connect their centers together and to allow the clients of these centers to access them. At first, the NSF tried to use the ARPAnet for communications, but this strategy failed because of bureaucracy and staffing problems.

In response, NSF decided to build its own network, based on the ARPAnet's IP technology. It connected the centers with 56,000 bit per second (56k bps) telephone lines. (This is roughly the ability to transfer two full typewritten pages per second. That's slow by modern standards, but was reasonably fast in the mid 80's.) It was obvious, however, that if they tried to connect every university directly to a supercomputing center, they would go broke. You pay for these telephone lines by the mile. One line per campus with a supercomputing center at the hub, like spokes on a bike wheel, adds up to lots of miles of phone lines. Therefore, they decided to create regional networks. In each area of the country, schools would be connected to their nearest neighbor. Each chain was connected to a supercomputer center at one point and the centers were connected together. With this configuration, any computer could eventually communicate with any other by forwarding the conversation through its neighbors.

This solution was successful--and, like any successful solution, a time came when it no longer worked. Sharing supercomputers also allowed the connected sites to share a lot of other things not related to the centers. Suddenly these schools had a world of data and collaborators at their fingertips. The network's traffic increased until, eventually, the computers controlling the network and the telephone lines connecting them were overloaded. In 1987, a contract to manage and upgrade the network was awarded to Merit

Network Inc., which ran Michigan's educational network, in partnership with IBM and MCI. The old network was replaced with faster telephone lines (by a factor of 20), with faster computers to control it.

The process of running out of horsepower and getting bigger engines and better roads continues to this day. Unlike changes to the highway system, however, most of these changes aren't noticed by the people trying to use the Internet to do real work. You won't go to your office, log in to your computer, and find a message saying that the Internet will be inaccessible for the next six months because of improvements. Perhaps even more important: the process of running out of capacity and improving the network has created a technology that's extremely mature and practical. The ideas have been tested; problems have appeared, and problems have been solved.

For our purposes, the most important aspect of the NSF's networking effort is that it allowed everyone to access the network. Up to that point, Internet access had been available only to researchers in computer science, government employees, and government contractors. The NSF promoted universal educational access by funding campus connections only if the campus had a plan to spread the access around. So everyone attending a four year college could become an Internet user.

The demand keeps growing. Now that most four-year colleges are connected, people are trying to get secondary and primary schools connected. People who have graduated from college know what the Internet is good for, and talk their employers into connecting corporations. All this activity points to continued growth, networking problems to solve, evolving technologies, and job security for networkers.

What Makes Up the Internet?

What comprises the Internet is a difficult question; the answer changes over time. Five years ago the answer would have been easy: "All the networks, using the IP protocol, which cooperate to form a seamless network for their collective users." This would include various federal networks, a set of regional networks, campus networks, and some foreign networks.

More recently, some non-IP-based networks saw that the Internet was good. They wanted to provide its services to their clientele. So they developed methods of connecting these "strange" networks (e.g., Bitnet, DECnets, etc.) to the Internet. At first these connections, called "gateways", merely served to transfer electronic mail between the two networks. Some, however, have grown to translate other

services between the networks as well. Are they part of the Internet? Maybe yes and maybe no. It depends on whether, in their hearts, they want to be. If this sounds strange, read on--it gets stranger.

Who Governs the Internet?

In many ways the Internet is like a church: it has its council of elders, every member has an opinion about how things should work, and you can either take part or not. It's your choice. The Internet has no president, chief operating officer, or Pope. The constituent networks may have presidents and CEO's, but that's a different issue; there's no single authority figure for the Internet as a whole.

The ultimate authority for where the Internet is going rests with the Internet Society, or ISOC. ISOC is a voluntary membership organization whose purpose is to promote global information exchange through Internet technology. (If you'd like more information, or if you would like to join, contact information is provided in the "For More Information" section, near the end of this document.) It appoints a council of elders, which has responsibility for the technical management and direction of the Internet.

The council of elders is a group of invited volunteers called the Internet Architecture Board, or the IAB. The IAB meets regularly to "bless" standards and allocate resources, like addresses. The Internet works because there are standard ways for computers and software applications to talk to each other. This allows computers from different vendors to communicate without problems. It's not an IBM-only or Sun-only or Macintosh-only network. The IAB is responsible for these standards; it decides when a standard is necessary, and what the standard should be. When a standard is required, it considers the problem, adopts a standard, and announces it via the network. (You were expecting stone tablets?) The IAB also keeps track of various numbers (and other things) that must remain unique. For example, each computer on the Internet has a unique 32-bit address; no other computer has the same address. How does this address get assigned? The IAB worries about these kinds of problems. It doesn't actually assign the addresses, but it makes the rules about how to assign addresses.

As in a church, everyone has opinions about how things ought to run. Internet users express their opinions through meetings of the Internet Engineering Task Force (IETF). The IETF is another volunteer organization; it meets regularly to discuss operational and near-term technical problems of the Internet. When it considers a problem important enough to merit concern, the IETF sets up a "working group" for further investigation. (In practice, "important enough" usually means that there are enough people to volunteer for the working

group.) Anyone can attend IETF meetings and be on working groups; the important thing is that they work. Working groups have many different functions, ranging from producing documentation, to deciding how networks should cooperate when problems occur, to changing the meaning of the bits in some kind of packet. A working group usually produces a report. Depending on the kind of recommendation, it could just be documentation and made available to anyone wanting it, it could be accepted voluntarily as a good idea which people follow, or it could be sent to the IAB to be declared a standard.

If you go to a church and accept its teachings and philosophy, you are accepted by it, and receive the benefits. If you don't like it, you can leave. The church is still there, and you get none of the benefits. Such is the Internet. If a network accepts the teachings of the Internet, is connected to it, and considers itself part of it, then it is part of the Internet. It will find things it doesn't like and can address those concerns through the IETF. Some concerns may be considered valid and the Internet may change accordingly. Some of the changes may run counter to the religion, and be rejected. If the network does something that causes damage to the Internet, it could be excommunicated until it mends its evil ways.

Who Pays for It?

The old rule for when things are confusing is "follow the money." Well, this won't help you to understand the Internet. No one pays for "it"; there is no Internet, Inc. that collects fees from all Internet networks or users. Instead, everyone pays for their part. The NSF pays for NSFNET. NASA pays for the NASA Science Internet. Networks get together and decide how to connect themselves together and fund these interconnections. A college or corporation pays for their connection to some regional network, which in turn pays a national provider for its access.

What Does This Mean for Me?

The concept that the Internet is not a network, but a collection of networks, means little to the end user. You want to do something useful: run a program, or access some unique data. You shouldn't have to worry about how it's all stuck together. Consider the telephone system--it's an internet, too. Pacific Bell, AT&T, MCI, British Telephony, Telefonos de Mexico, and so on, are all separate corporations that run pieces of the telephone system. They worry about how to make it all work together; all you have to do is dial.

If you ignore cost and commercials, you shouldn't care if you are dealing with MCI, AT&T, or Sprint. Dial the number and it works.

You only care who carries your calls when a problem occurs. If something goes out of service, only one of those companies can fix it. They talk to each other about problems, but each phone carrier is responsible for fixing problems on its own part of the system. The same is true on the Internet. Each network has its own network operations center (NOC). The operation centers talk to each other and know how to resolve problems. Your site has a contract with one of the Internet's constituent networks, and its job is to keep your site happy. So if something goes wrong, they are the ones to gripe at. If it's not their problem, they'll pass it along.

What Does the Future Hold?

Finally, a question I can answer. It's not that I have a crystal ball (if I did I'd spend my time on Wall Street instead of writing a book). Rather, these are the things that the IAB and the IETF discuss at their meetings. Most people don't care about the long discussions; they only want to know how they'll be affected. So, here are highlights of the networking future.

New Standard Protocols

When I was talking about how the Internet started, I mentioned the International Standards Organization (ISO) and their set of protocol standards. Well, they finally finished designing it. Now it is an international standard, typically referred to as the ISO/OSI (Open Systems Interconnect) protocol suite. Many of the Internet's component networks allow use of OSI today. There isn't much demand, yet. The U.S. government has taken a position that government computers should be able to speak these protocols. Many have the software, but few are using it now.

It's really unclear how much demand there will be for OSI, notwithstanding the government backing. Many people feel that the current approach isn't broke, so why fix it? They are just becoming comfortable with what they have, why should they have to learn a new set of commands and terminology just because it is the standard?

Currently there are no real advantages to moving to OSI. It is more complex and less mature than IP, and hence doesn't work as efficiently. OSI does offer hope of some additional features, but it also suffers from some of the same problems which will plague IP as the network gets much bigger and faster. It's clear that some sites will convert to the OSI protocols over the next few years. The question is: how many?

International Connections

The Internet has been an international network for a long time, but it only extended to the United States' allies and overseas military bases. Now, with the less paranoid world environment, the Internet is spreading everywhere. It's currently in over 50 countries, and the number is rapidly increasing. Eastern European countries longing for western scientific ties have wanted to participate for a long time, but were excluded by government regulation. This ban has been relaxed. Third world countries that formerly didn't have the means to participate now view the Internet as a way to raise their education and technology levels.

In Europe, the development of the Internet used to be hampered by national policies mandating OSI protocols, regarding IP as a cultural threat akin to EuroDisney. These policies prevented development of large scale Internet infrastructures except for the Scandinavian countries which embraced the Internet protocols long ago and are already well-connected. In 1989, RIPE (Reseaux IP Europeens) began coordinating the operation of the Internet in Europe and presently about 25% of all hosts connected to the Internet are located in Europe.

At present, the Internet's international expansion is hampered by the lack of a good supporting infrastructure, namely a decent telephone system. In both Eastern Europe and the third world, a state-of-the-art phone system is nonexistent. Even in major cities, connections are limited to the speeds available to the average home anywhere in the U.S., 9600 bits/second. Typically, even if one of these countries is "on the Internet," only a few sites are accessible. Usually, this is the major technical university for that country. However, as phone systems improve, you can expect this to change too; more and more, you'll see smaller sites (even individual home systems) connecting to the Internet.

Commercialization

Many big corporations have been on the Internet for years. For the most part, their participation has been limited to their research and engineering departments. The same corporations used some other network (usually a private network) for their business communications. After all, this IP stuff was only an academic toy. The IBM mainframes that handled their commercial data processing did the "real" networking using a protocol suite called System Network Architecture (SNA).

Businesses are now discovering that running multiple networks is expensive. Some are beginning to look to the Internet for "one-stop"

network shopping. They were scared away in the past by policies which excluded or restricted commercial use. Many of these policies are under review and will change. As these restrictions drop, commercial use of the Internet will become progressively more common.

This should be especially good for small businesses. Motorola or Standard Oil can afford to run nationwide networks connecting their sites, but Ace Custom Software couldn't. If Ace has a San Jose office and a Washington office, all it needs is an Internet connection on each end. For all practical purposes, they have a nationwide corporate network, just like the big boys.

Privatization

Right behind commercialization comes privatization. For years, the networking community has wanted the telephone companies and other for-profit ventures to provide "off the shelf" IP connections. That is, just like you can place an order for a telephone jack in your house for your telephone, you could do this for an Internet connection. You order, the telephone installer leaves, and you plug your computer into the Internet. Except for Bolt, Beranek and Newman, the company that ran the ARPAnet, there weren't any takers. The telephone companies have historically said, "We'll sell you phone lines, and you can do whatever you like with them." By default, the Federal government stayed in the networking business.

Now that large corporations have become interested in the Internet, the phone companies have started to change their attitude. Now they and other profit-oriented network purveyors complain that the government ought to get out of the network business. After all, who best can provide network services but the "phone companies"? They've got the ear of a lot of political people, to whom it appears to be a reasonable thing. If you talk to phone company personnel, many of them still don't really understand what the Internet is about. They ain't got religion, but they are studying the Bible furiously. (Apologies to those telephone company employees who saw the light years ago and have been trying to drag their employers into church.)

Although most people in the networking community think that privatization is a good idea, there are some obstacles in the way. Most revolve around the funding for the connections that are already in place. Many schools are connected because the government pays part of the bill. If they had to pay their own way, some schools would probably decide to spend their money elsewhere. Major research institutions would certainly stay on the net; but some smaller colleges might not, and the costs would probably be prohibitive for most secondary schools (let alone grade schools). What if the school could afford either an Internet connection or a science lab? It's

unclear which one would get funded. The Internet has not yet become a "necessity" in many people's minds. When it does, expect privatization to come quickly.

Well, enough questions about the history of the information highway system. It's time to walk to the edge of the road, try and hitch a ride, and be on your way.

Acknowledgments

We would like to thank O'Reilly & Associates for permission to reprint the chapter from their book by Ed Krol (1992), "The Whole Internet User's Guide and Catalog."

For More Information

Hoffman, E. and L. Jackson. (1993) "FYI on Introducing the Internet --A Short Bibliography of Introductory Internetworking Readings for the Network Novice," 4 p. (FYI 19, RFC 1463).

To find out how to obtain this document and other on-line introductory readings, send an e-mail message to:
nis-info@nis.merit.edu, with the following text:
send access.guide.

Krol, Ed. (1992) The Whole Internet User's Guide and Catalog, O'Reilly & Associates, Sebastopol, CA. ISBN 1-56592-025-2.

Quarterman, J. (1993) "Recent Internet Books," 15 p. (RFC 1432).

The Internet Society
Phone: (703) 620-8990
Fax: (703) 620-0913
E-mail: isoc@cnri.reston.va.us

Security Considerations

Security issues are not discussed in this memo.

Authors' Addresses

Ed Krol
Computing and Communications Service Office
Univ. of Illinois Urbana Champaign (UIUC)
1304 W Springfield
Urbana, IL 61801

Phone: (217)333-7886
EMail: e-krol@uiuc.edu

Ellen Hoffman
Merit Network, Inc.
2901 Hubbard, Pod-G
Ann Arbor, MI 48105

Phone: (313) 936-3000
EMail: ellen@merit.edu

