

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

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.