Regulating Cyberspace

The explosive growth of the Internet is leading to a collision with carrying capacity—and with the legal community. Internet directors are moving to avoid the first impact, but not the second

Adorning *People* magazine's list of the "25 most intriguing people" on Earth last December was the amiable face of Vinton Cerf, a senior vice president at MCI Telecommunications Corp. Cerf won his place in the *People* honor roll for his role as "the father of the Internet." He helped design the initial architecture of a system that in 1983 had only 200 host computers, but today has more than 5 million hosts in 94 countries; it is growing at an estimated rate of 9% to 12% a month. "The fact of the matter," Cerf says, "is that [the Internet] is an unstoppable explosion."

Yet that explosion is now threatening to consume the Internet itself. Its central computers are groaning beneath the task of tracking the routes between its millions of hosts. The rampant growth of the "routing tables," as the database of interconnections is known, is taxing both the memory capacity of the hardware that stores them and the human capacity of the people who maintain them. A study backed by the Internet Society—the closest thing the global network has to an administration—predicted last year that the routing table system could break down as early as 1997. If that occurred, millions of users, including an appreciable fraction of the world's scientists and engineers, would be stuck in electronic traffic jams, unable to contact each other reliably.

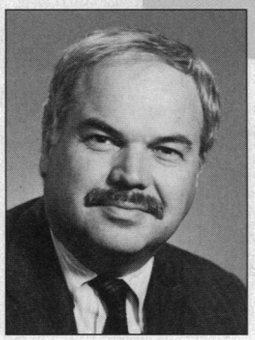

At the 1995 Internet Society Summit, held in San Diego from 11 to 12 April, the Society unveiled a possible solution: a suite of innovations known as Internet Protocol version 6. Although the changes will mostly be invisible to users, IPv6—which distributes the routing responsibility more widely, easing the burden on central computers, all the while aiming to enhance message security—may be the most fundamental change in the Internet for more than a decade. "It doesn't answer every problem," says Anthony M. Rutkowski, executive director of the Internet Society. "But it answers a lot of them. It will take the Internet to the next century."

Yet even as the Internet Society was announcing IPv6, the summit provided evidence that the Internet faces still greater challenges—cultural, legal, and social issues that Cerf and the other technophiles who created the network sometimes appear unprepared to tackle. The society's seemingly cavalier treatment of intellectual property—for instance, giving away "mci.com" as an Internet address for MCI's telecommunications rival US Sprint—has already sparked litigation. Growing demands that the now-ubiquitous Internet traffic obey the laws of every aspect of society fundamentally contradict the libertarian ethos of the network's founders. At the Internet summit, the clash was made official: For the first time, this proudly anarchic group invited a panel of lawyers to speak.

"Imagine what Coca-Cola thinks when they come to get a network address and these guys have given 'coke.com' to somebody else."

—David W. Maher

Lost in the mail

Based in Reston, Virginia, the Internet Society was established in 1992 to integrate groups of researchers that had administered the Internet in an ad hoc way when it was a government-sponsored entity. The society is a nonprofit organization with no statutory authority or enforcement power; staffed mostly by volunteers, it has just four full-time employees. Nonetheless, it creates standards for Internet transmissions, helps maintain routing tables, allocates electronic addresses, and promotes the growth of the Internet in developing countries. At its helm are 18 board members, many of them the computer scientists who helped develop the network software. As board member Jon Postel jokingly told the summit, "The secret is out—the Internet does have a government, at least a little bit of one. And we are it."

One of the most pressing tasks facing that "little bit of government" is delivering the mail. When one computer user on the Internet sends information to another, the message is broken up into small chunks, each numbered like a model plane part, and then sent independently to the destination, where the pieces are reassembled in their original order. The method is standardized by an Internet Protocol—the system now uses IP version 4.

Accomplishing this task requires adorning each packet with a "header"—about 600 bits of information that includes the unique addresses of both source and destination computers. The addresses have two independent parts: a text of several words or abbreviations separated by periods, which is used by people ("aaas.org," for instance, reaches the publisher of this journal); and a string of four numbers, separated by periods, which is recognized by the machines.

The source computer sends packets of data to a router or gateway, a specialized computer that shunts traffic between different networks. The Internet has thousands of routers (no one knows the exact number) of many different types. All contain routing tables, which are, in essence, instructions for getting from here to there. The router reads the address, looks up how to reach it, and zips the packet to the next stop on its way.

Routers at the edges of the system, Postel says, can get away with small routing tables. In most cases, he points out, they "can just send their stuff to the Big Brother routers in the center of the world"—the so-called "backbone" lines of the Internet. "But those big routers have to know everything."

As the Internet expands, knowing everything is getting harder and harder. In 1992, an Internet Society task force reported that many of the routers maintained by the National Science Foundation were already unable to keep up, and that the demands on the present system were increasing 50% faster than the growth of memory capacity. At present, big routing tables are matrices with about 22,000 entries, which have to be searched thousands of times each second as message packets come in. New high-end routers can handle up to 64,000 entries—not a large margin for a system growing at an explosive rate. Even worse, Postel says, "is keeping the routing tables up to date. When a new [network] joins, it changes how people want to get from place to place."

Enter IPv6 (v5 was assigned long ago to a little-used system). Developed primarily by Steve Deering, a researcher at Xerox's Palo

Alto Research Center, the proposed new standard was published to coincide with the Internet Summit. IPv6 changes the electronic addresses and alters the header that contains them. Today, addresses have no relation to physical location. Computers whose addresses are just one digit apart can be thousands of miles away from each other. The new addresses may be more "topological," which means, crudely speaking, that computers with similar addresses will be reachable on similar routes—an obvious step toward simplifying the routing tables.

IPv6 avoids a looming "address crunch" facing the system (*Science*, 12 August 1994, p. 879) by increasing the length of addresses fourfold, turning them into strings of 16 numbers rather than the present four. Because the present system awards each network a minimum of 256 addresses—even though there may only be 10 computers in the network—the Internet may run out of addresses by 2005. Longer addresses should push the day of reckoning centuries into the future.

Other changes, IPv6 backers hope, will improve security—an increasingly critical Internet problem. IPv6 includes two subheaders for "public-key" encryption, in which freely available mathematical algorithms (public keys, as they are known) can be used to create messages that can be decoded only by certain individuals (the owners of the associated private keys). One subheader would use this technology to let the destination computer check whether a packet was actually from its purported sender. The second subheader would let users decode an entire encrypted packet. An algorithm accompanying the packet in still another subheader would tell the recipients which of their private keys can decode it.

Unfortunately, notes Allison Mankin of the Information Sciences Institute, "the [encryption] technology is basically illegal." The U.S. government bans encryption schemes that can be exported to other countries. There is also open argument among Net experts as to whether the huge computational demands associated with encryption will jam up the Net as badly as the overwhelmed routing tables.

Partly for these reasons, no one can be certain that IPv6 will be adopted; the Internet Society can't compel anyone to take it on. "It will depend on Microsoft and cable TV, who are really going to be the big players," Harvard University's Scott Bradner told the summit, referring to the software company expected to become the world's biggest Internet service provider and the cable operators expected to become major transmitters. Although IPv6 is intended to coexist with the present IP, its success will depend on whether the major players will risk changing horses in midstream. Ultimately, Bradner said, "I don't know what the prospects for IPv6 are."

Clashing cultures

Other problems created by the growth of the Internet seem less amenable to technical adjustment. In microcosm, the difficulties can be seen in domain names.

Domain names are the symbols to the right of the familiar @ sign in e-mail addresses. In the past, assigning these domains was a simple process, handled by volunteers. According to David Maher, a Chicago lawyer who is co-chair of the International Trademark Association (INTA) Task Force on the Internet, the tradition is continued in North America by Network Solutions, a private company that under the auspices of the Internet Society allocates addresses in the supposedly decentralized network. (Two other organizations do the same thing in Europe and Asia.) "All the work [for North America] is done by a small group of computer guys in T-shirts," Maher says. "So you can imagine what Coca-Cola thinks when they come to get a network address and find out these guys have given away 'coke.com' to somebody else."

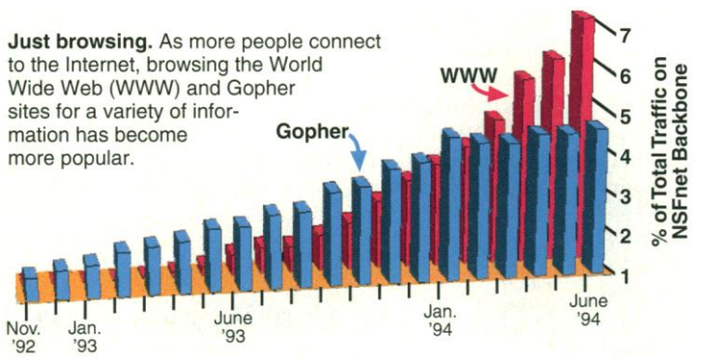

SOURCE: ISOC

Just browsing. As more people connect to the Internet, browsing the World Wide Web (WWW) and Gopher sites for a variety of information has become more popular.

Because such names are trademarked, Maher told the summit, corporations are becoming increasingly concerned at what seems to them the Internet Society's indifference to pirating. In addition to the contretemps between rivals MCI and Sprint, the music network MTV ended up in court battling a former employee who had snapped up "mtv.com" for himself. And the legal wrangles are beginning to hit home. In January, Network Solutions was sued because it had awarded the domain name "knowledge.net" to D. L. Boone & Co., a Virginia-based consulting firm—a decision that enraged KnowledgeNet, an Illinois consulting and software company.

"You can't tell these huge enterprises that the Internet will be the medium of commerce in the future and then tell them they can't use their own name," said Lance Rose, a New Jersey lawyer and author of *Netlaw*. "Anyone who thinks Net culture will survive the influx of billions of dollars doesn't know that millions of those billions will be spent on lawyers."

Speaking to the summit, Maher suggested that trademark searches need to be conducted before new domain names are awarded. Longtime Net regulars were not pleased—such searches are lengthy procedures, made all the more cumbersome by the need to examine records in every nation in which the Internet operates. "That's a ridiculous idea," snapped one summit participant. When Maher suggested that failure to grapple with the problem could lead to Congressional action, scornful muttering filled the room. "This is a real issue," Maher responded. "To say we [lawyers] aren't involved just won't go."

The lawyers had gloomier news yet. If you think domain names cause trouble, panelists said, you'd better start thinking about libel laws, obscenity statutes, and interstate and international commerce regulations, which will become applicable as the Internet grows into a vehicle for publishing and commercial transactions. "The regulators are drooling to run things and set directions" for the Internet, said Jan Constantine, a lawyer for Delphi, an Internet access provider. "One by one, the states will pick you off."

Typical of the reactions was Cerf's argument that the states "couldn't possibly do that," because the interconnected structure of the Internet makes it impossible for the system to "obey a patchwork of laws." Network operators would have to satisfy the legal codes of every nation in the Internet, not to mention countless states and provinces. With those operators not even able to monitor what is on their own systems, he said, the idea is "impossible on its face." This type of regulation, Cerf argued, would lead to countless situations like the recent case in which the operators of a bulletin board in California were indicted because their files could be used in Tennessee.

Afterward, Maher shook his head at the response. The Internet, he said, was "colliding with America—it's ignoring an enormous legal, cultural, and political culture." Working with that culture, he suggested, "will be a lot harder than anything they've done before. The technical problems [like developing IPv6] are nothing." And unlike adjusting to Internet protocols, compliance with those standards will not be voluntary.

–Charles Mann

Charles Mann is the co-author of Noah's Choice: The Future of Endangered Species.