# Striking at the Internet's Heart

BY STEVEN M. CHERRY
SENIOR ASSOCIATE EDITOR

E-mail worked flawlessly on 11 September; the Internet was only peripherally affected by the events of that tragic day. But how well could it have withstood some kind of frontal attack?

The Internet "was designed to route around problems," points out Internet architect John Gilmore, a San Francisco entrepreneur. "If one major point goes down, the rest of the Net doesn't know what the problem is, but the traffic will be routed around it." Indeed, the Internet originated in a military research network called the Arpanet, conceived as a distributed network so richly interconnected that if any part of it failed, all the remaining points would still be able to communicate with one another.

But one critical vulnerability, some believe, lies in the basic Internet addressing scheme: its domain name system (DNS). Although this belief is mistaken, it exists in surprising places. According to a 28 September article in *The Wall Street Journal*, "The Bush administration for months has privately expressed concerns about the security of the Internet's 13 most important computers, called root servers, which manage global Internet traffic."

Could the Internet, grandchild of the perfectly distributed Arpanet, be so centralized as to be managed by 13 computers? Is their security a cause for concern? The short answers are, no and no. First, the root servers do not manage global Internet traffic; they handle relatively few domain name queries, the ones that cannot be handled by the many thousands of lower-level name servers in the DNS system. Second, those root servers are extremely secure.

The protocols that govern the 13 root servers call for complete diversity. The servers are geographically distributed, 10 in the United States and one each in London, Stockholm, and Tokyo [see figure]. Those in the United States are almost evenly divided between the East and West coasts. Two are run by the U.S. military, but little is publicly known about them.

One server, A (they are known by letters of the alphabet), is operated by the Network Solutions subsidiary of Verisign Inc., in Mountain View, Calif. It contains the master database of top-level domains (such as .com, .org, and .uk); the other 12 are copies. Tests performed in 1999 (for possible Year 2000 problems) demonstrated that five or six servers would have to fail before the loads on the remainder would increase appreciably.

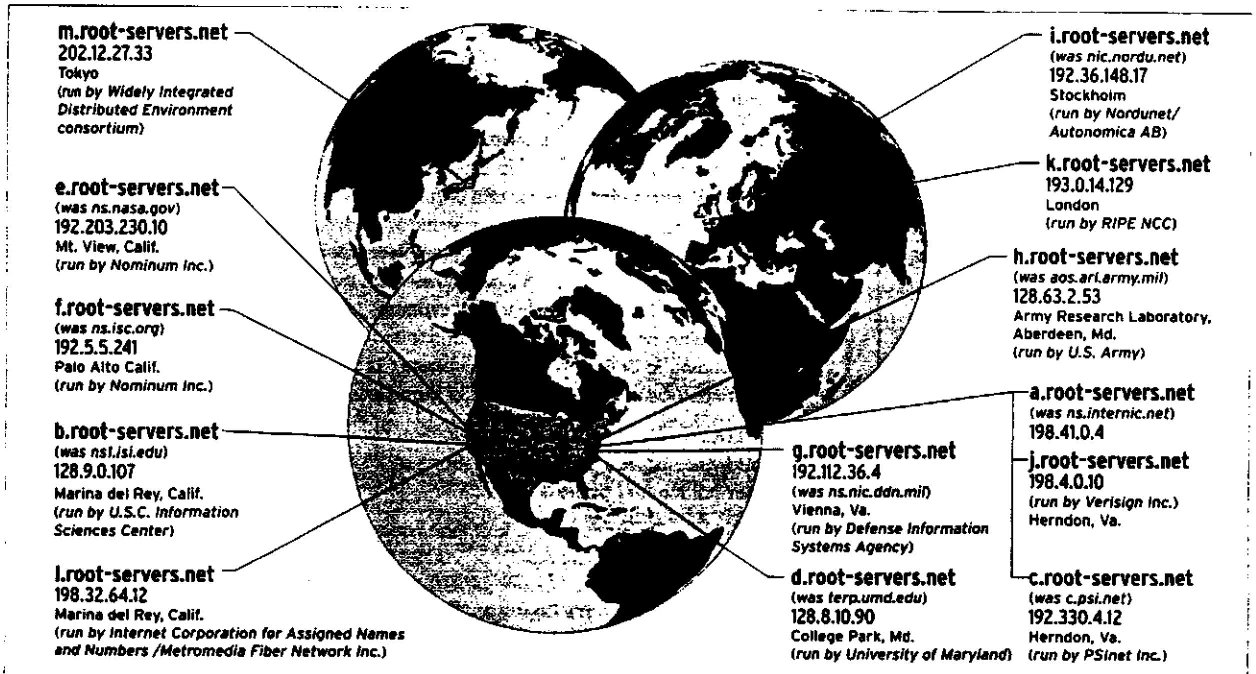

***For some time, the 13 Internet root servers around the world have been informally referred to by the letters A to M. Since 1997, those letters, prefixed to root servers.net, have been adopted as the formal names. Also listed is the former name of each root server, its numerical IP address, and the organization in charge of it.***

## A court of last resort

Actually, the 13 servers do very little. They are at the top of a three-tier hierarchy of servers whose job is to sort out which particular computer, identified by a series of numbers, is represented by one of the more than 36 million domain names registered on the Internet. For example, ieee.org (a domain name) is the machine 140.98.193.38, its Internet protocol address.

The domain name servers in the lowest tier are maintained by Internet service providers and organizations with internal computer networks. The middle tier consists of more than 100 authoritative computers (and many more backups) around the world; these specialize in finding protocol addresses for generic top-level domains, such as .org or .com, or for country domains, like .uk for United Kingdom or .de for Germany. They keep track of which lower-tier servers can authoritatively answer a name query.

When a user of, say, the Internet service provider Earthlink writes to someone at ieee.org, the interpretation of that domain name is handled by lower-tier DNS servers at Earthlink and elsewhere. If these fail, the query can go to other lower-tier name servers. Only when these lower-tier servers also fail is a root server queried. Its job is then just to direct the query to the authoritative middle-tier DNS server, the top-level domain name server—in this case, for .org.

In effect, as in a judicial system, the 13 root servers are the court of last resort. Although their importance cannot be overstated, they see only a small fraction of the case load.

Every imaginable form of redundancy is required at the root servers, including two or more computers at each location. One is ready to step in if the other fails. The D root server, run by the University of Maryland in College Park, is typical.

"We have multiple power supplies and network connections into the campus backbone, which is connected to the Internet in multiple ways," said Gerry Sneeringer, who belongs to the team that supports the root server plus the lower-tier name servers that handle queries for the university's own Internet traffic. "The server is also in one of several secure spaces on campus, and it's climate controlled. Physical access is limited, with layers of locked doors."

The D root server was singled out in the September *Journal* article. "When congressional auditors recently checked the security surrounding [the root servers], 'one of them was sitting in a professor's office at the University of Maryland,' " the article said, quoting a statement made by Keith Rhodes, a General Accounting Office technologist.

When *IEEE Spectrum* asked Sneeringer if he had seen the article, he laughed and said that the D root server "had always been in a machine room for as long as we've been running it." He pointed out that in 1998–1999, the Internet infrastructure was being investigated by the National Security Telecommunications Advisory Committee (NSTAC), and during the inquiry, a Network Solutions executive commented, inaccurately, that the machine "was on the desk of a grad student with a pizza box on top of it."

"NSTAC was at our doors in a flash and did an inspection," Sneeringer said. "They're the only government people to have been here—the GAO [General Accounting Office] has never even seen our root server." Rhodes confirmed the account, saying he was inaccurately quoted.

David Conrad, vice president of engineering at Nominum Inc., is also bothered by the in-a-professor's-office story, which seems to have the persistence of an urban legend. "Any suggestion that the root servers are not managed in a professional fashion is offensive," he said. The Redwood City, Calif., company operates the E and F root servers on behalf of NASA and the Internet Software Consortium, respectively. The consortium is responsible for BIND, the DNS software package most commonly used, and used by all the root servers.

## Attacks from within

To reduce possible entry points of an attack, very little software is run on the root computers, only whatever is needed to perform DNS duties. In fact, they run only an operating system (one other version of Unix), BIND, and a network time protocol program that synchronizes clocks around the Internet. Common network services, such as remote log-in, e-mail, and Web protocols, are disabled.

The one threat that root server operators consider realistic is a denial-of-service attack, which floods a server with seemingly legitimate traffic. The original Code Red virus was such an attack, mounted in July against the U.S. White House's Web server. Computers infected with it assailed the whitehouse.gov domain with page requests. Similarly, a virus could disseminate a program that, at any given time, could cause millions of computers around the Internet to repeatedly query the 13 root servers with requests for Internet protocol addresses.

"We know that the busiest root server, F, can—in fact, has—handled eight times its normal load of 8000 queries per second," said Nominum's Conrad. "And the A server can handle even more than F." Still, a clever denial-of-service attack could produce many more than 64 000 queries a second.

To reduce traffic under even ordinary circumstances, however, lower-level DNS servers are allowed, even encouraged, to cache data. This means they store the answers to queries asked in the past 24 or 48 hours (up to a week is permitted). When the same query is made again (and again, as in a denial-of-service attack), the answer is drawn from the cache. Thus, even in a successful attack, there's a 24–48-hour window during which the Internet could translate domain names without accessing the root servers.

But system administrators would realize an attack was under way because they would see very many, very similar requests. "We would tell providers, companies like Earthlink and AOL, 'don't send packets that look like such-and-such,' " pointed out Conrad. Added Sneeringer, a denial-of-service attack "would result in an immense amount of cooperation from the Internet service provider community before the cache ran out."

An attack on the DNS and its 13 root servers might be a prodigious waste of time. With an architecture built for security, a fundamentally redundant nature, and widely dispersed sites, the network's ability to withstand natural or man-made disasters should be the envy of computer networks everywhere. ●

What if people bought cars like they buy computers?

General Motors doesn't have a "help line" for people who don't know how to drive, because people don't buy cars like they buy computers --but imagine if they did...

*********************************************************

HELPLINE: "General Motors Helpline, how can I help you?"

CUSTOMER: "I got in my car and closed the door, and nothing happened!"

HELPLINE: "Did you put the key in the ignition and turn it?"

CUSTOMER: "What's an ignition?"

HELPLINE: "It's a starter motor that draws current from your battery and turns over the engine."

CUSTOMER: "Ignition? Motor? Battery? Engine? How come I have to know all of these technical terms just to use my car?"

*********************************************************

HELPLINE: "General Motors Helpline, how can I help you?"

CUSTOMER: "My car ran fine for a week, and now it won't go anywhere!"

HELPLINE: "Is the gas tank empty?"

CUSTOMER: "Huh? How do I know?"

HELPLINE: "There's a little gauge on the front panel, with a needle, and markings from 'E' to 'F'. Where is the needle pointing?"

CUSTOMER: "I see an 'E' but no 'F'."

HELPLINE: "You see the 'E' and just to the right is the 'F'.

CUSTOMER: "No, just to the right of the first 'E' is a 'V'.

HELPLINE: "A 'V'?!?"

CUSTOMER: "Yeah, there's a 'C', an 'H', the first 'E', then a 'V', followed by 'R', 'O', 'L' ..."

HELPLINE: "No, no, no sir! That's the front of the car. When you sit behind the steering wheel, that's the panel I'm talking about."

CUSTOMER: "That steering wheel thing -- Is that the round thing that honks the horn?"

HELPLINE: "Yes, among other things."

CUSTOMER: "The needle's pointing to 'E'. What does that mean?"

HELPLINE: "It means that you have to visit a gasoline vendor and purchase some more gasoline. You can install it yourself, or pay the vendor to install it for you."

CUSTOMER: "What? I paid $12,000 for this car! Now you tell me that I have to keep buying more components?

I want a car that comes with everything built in!"

*******************************************************

HELPLINE: "General Motors Helpline, how can I help you?"

CUSTOMER: "Your cars suck!"

HELPLINE: "What's wrong?"

CUSTOMER: "It crashed, that's what went wrong!"

HELPLINE: "What were you doing?"

CUSTOMER: "I wanted to go faster, so I pushed the accelerator pedal all the way to the floor. It worked for a while, and then it crashed -- and now it won't even start up!"

HELPLINE: "I'm sorry, sir, but it's your responsibility if you misuse the product."

CUSTOMER: "Misuse it? I was just following this damned manual of yours. It said to make the car go to put the transmission in 'D' and press the accelerator pedal. That's exactly what I did --now the damn thing's crashed."

HELPLINE: "Did you read the entire operator's manual before operating the car sir?"

CUSTOMER: "What? Of course I did! I told you I did EVERYTHING the manual said and it didn't work!"

HELPLINE: "Didn't you attempt to slow down so you wouldn't crash?"

CUSTOMER: "How do you do THAT?"

HELPLINE: "You said you read the entire manual, sir. It's on page 14. The pedal next to the accelerator."

CUSTOMER: "Well, I don't have all day to sit around and read this manual you know."

HELPLINE: "Of course not. What do you expect us to do about it?"

CUSTOMER: "I want you to send me one of the latest versions that goes fast and won't crash anymore!"

*******************************************************

HELPLINE: "General Motors Helpline, how can I help you?"

CUSTOMER: "Hi! I just bought my first car, and I chose your car because it has automatic transmission, cruise control, power steering, power brakes, and power door locks."

HELPLINE: "Thanks for buying our car. How can I help you?"

CUSTOMER: "How do I work it?"

HELPLINE: "Do you know how to drive?"

CUSTOMER: "Do I know how to what?"

HELPLINE: "Do you know how to DRIVE?"

CUSTOMER: "I'm not a technical person! I just want to go places in my car!"