

Tuesday, December 20, 2005

The Internet Is Broken -- Part 2

We can't keep patching the Internet's security holes. Now computer scientists are proposing an entirely new architecture.

By David Talbot

This article -- the cover story in Technology Review's December-January print issue -- has been divided into three parts for presentation online. This is part 2; [part 1](#) appeared on December 19 and [part 3](#) will appear on December 21.

In part 1, TR Chief Correspondent David Talbot argued that the "Internet has no inherent security architecture -- nothing to stop viruses or spam or anything else. Protections like firewalls and antispam software are add-ons, security patches in a digital arms race." Jonathan Zittrain, cofounder of the Berkman Center for Internet and Society at Harvard Law School, told Talbot that the Internet functions as well as it does only because of "the forbearance of the virus authors themselves." Here's more about why -- and how -- we might start to fix the problem.

Patchwork Problem

The Internet's original protocols, forged in the late 1960s, were designed to do one thing very well: facilitate communication between a few hundred academic and government users. The protocols efficiently break digital data into simple units called packets and send the packets to their destinations through a series of network routers. Both the routers and PCs, also called nodes, have unique digital addresses known as Internet Protocol or IP addresses. That's basically it. The system assumed that all users on the network could be trusted and that the computers linked by the Internet were mostly fixed objects.

The Internet's design was indifferent to whether the information packets added up to a malicious virus or a love letter; it had no provisions for doing much besides getting the data to its destination. Nor did it accommodate nodes that moved -- such as PDAs that could connect to the Internet at any of myriad locations. Over the years, a slew of patches arose: firewalls, antivirus software, spam filters, and the like. One patch assigns each mobile node a new IP address every time it moves to a new point in the network.

[[Click here](#) to view graphic representations of David D. Clark's four goals for a new Internet architecture.]

Clearly, security patches aren't keeping pace. That's partly because different people use different patches and not everyone updates them religiously; some people don't have any installed. And the most common mobility patch -- the IP addresses that constantly change as you move around -- has downsides. When your mobile computer has a new identity every time it connects to the Internet, the websites you deal with regularly won't know it's you. This means, for example, that your favorite airline's Web page might not cough up a reservation form with your name and frequent-flyer number already filled out. The constantly changing address also means you can expect breaks in service if you are using the Internet to, say, listen to a streaming radio broadcast on your PDA. It also means that someone who commits a crime online using a mobile device will be harder to track down.

In the view of many experts in the field, there are even more fundamental reasons to be concerned. Patches create an ever more complicated system, one that becomes harder to manage, understand, and improve upon. "We've been on a track for 30 years of incrementally making improvements to the Internet and fixing problems that we see," says Larry Peterson, a computer scientist at Princeton University. "We see vulnerability, we try to patch it. That approach is one that has worked for 30 years. But there is reason to be concerned. Without a long-term plan, if you are just patching the next problem you see, you end up with an increasingly complex and brittle system. It makes new services difficult to employ. It makes it much harder to manage because of the added complexity of all these point solutions that have been added. At the same time, there is concern that we will hit a dead end at some point. There will be problems we can't sufficiently patch."

The patchwork approach draws complaints even from the founder of a business that is essentially an elaborate and ingenious patch for some of the Internet's shortcomings. Tom Leighton is cofounder and chief scientist of Akamai, a company that ensures that its clients' Web pages and applications are always available, even if huge numbers of customers try to log on to them or a key fiber-optic cable is severed. Akamai closely monitors network problems, strategically stores copies of a client's website at servers around the world, and accesses those servers as needed. But while his company makes its money from patching the Net, Leighton says the whole system needs fundamental architectural change. "We are in the mode of trying to plug holes in the dike," says Leighton, an MIT mathematician who is also a member of the President's Information Technology Advisory Committee and chair of its Cyber Security Subcommittee.

"There are more and more holes, and more resources are going to plugging the holes, and there are less resources being devoted to fundamentally changing the game, to changing the Internet."

When Leighton says "resources," he's talking about billions of dollars. Take Microsoft, for example. Its software mediates between the Internet and the PC. These days, of the \$6 billion that Microsoft spends annually on research and development, approximately one-third, or \$2 billion, is directly spent on security efforts. "The evolution of the Internet, the development of threats from the Internet that could attempt to intrude on systems -- whether Web servers, Web browsers, or e-mail-based threats -- really changed the equation," says Steve Lipner, Microsoft's director of security strategy and engineering strategy. "Ten years ago, I think people here in the industry were designing software for new features, new performance, ease of use, what have you. Today, we train everybody for security." Not only does this focus on security siphon resources from other research, but it can even hamper research that does get funded. Some innovations have been kept in the lab, Lipner says, because Microsoft couldn't be sure they met security standards.

Of course, some would argue that Microsoft is now scrambling to make up for years of selling insecure products. But the Microsoft example has parallels elsewhere. Eric Brewer, director of Intel's Berkeley, CA, research lab, notes that expenditures on security are like a "tax" and are "costing the nation billions and billions of dollars." This tax shows up as increased product prices, as companies' expenditures on security services and damage repair, as the portion of processor speed and storage devoted to running defensive programs, as the network capacity consumed by spam, and as the costs to the average person trying to dodge the online minefield of buying the latest firewalls. "We absolutely can leave things alone. But it has this continuous 30 percent tax, and the tax might go up," Brewer says. "The penalty for not [fixing] it isn't immediately fatal. But things will slowly get worse and might get so bad that people won't use the Internet as much as they might like."

The existing Internet architecture also stands in the way of new technologies. Networks of intelligent sensors that collectively monitor and interpret things like factory conditions, the weather, or video images could change computing as much as cheap PCs did 20 years ago. But they have entirely different communication requirements. "Future networks aren't going to be PCs docking to mainframes. It's going to be about some car contacting the car next to it. All of this is happening in an embedded context. Everything is machine to machine rather than people to people," says Dipankar Raychaudhuri, director of the Wireless Information Network Laboratory

(Winlab) at Rutgers University. With today's architecture, making such a vision reality would require more and more patches.

Architectural Digest

When Clark talks about creating a new architecture, he says the job must start with the setting of goals. First, give the medium a basic security architecture -- the ability to authenticate whom you are communicating with and prevent things like spam and viruses from ever reaching your PC. Better security is "the most important motivation for this redesign," Clark says. Second, make the new architecture practical by devising protocols that allow Internet service providers to better route traffic and collaborate to offer advanced services without compromising their businesses. Third, allow future computing devices of any size to connect to the Internet -- not just PCs but sensors and embedded processors. Fourth, add technology that makes the network easier to manage and more resilient. For example, a new design should allow all pieces of the network to detect and report emerging problems -- whether technical breakdowns, traffic jams, or replicating worms -- to network administrators.

The good news is that some of these goals are not so far off. NSF has, over the past few years, spent more than \$30 million supporting and planning such research. Academic and corporate research labs have generated a number of promising technologies: ways to authenticate who's online; ways to identify criminals while protecting the privacy of others; ways to add wireless devices and sensors. While nobody is saying that any single one of these technologies will be included in a new architecture, they provide a starting point for understanding what a "new" Internet might actually look like and how it would differ from the old one.

Some promising technologies that might figure into this new architecture are coming from PlanetLab, which Princeton's Peterson has been nurturing in recent years (see "[The Internet Reborn](#)," October 2003). In this still-growing project, researchers throughout the world have been developing software that can be grafted onto today's dumb Internet routers. One example is software that "sniffs" passing Internet traffic for worms. The software looks for telltale packets sent out by worm-infected machines searching for new hosts and can warn system administrators of infections. Other software prototypes detect the emergence of data traffic jams and come up with more efficient ways to reroute traffic around them. These kinds of algorithms could become part of a fundamental new infrastructure, Peterson says.

A second set of technologies could help authenticate Internet communications. It would be a huge boon to Internet security if you could be sure an e-mail from your

bank is really from your bank and not a scam artist, and if the bank could be sure that when someone logs in to your account, that person is really you and not someone who stole your account number.

Today, the onus of authentication is on the Internet user, who is constantly asked to present information of various kinds: passwords, social-security numbers, employee ID numbers, credit card numbers, frequent-flyer numbers, PIN numbers, and so on. But when millions of users are constantly entering these gate-opening numbers, it makes it that much easier for spyware, or a thief sniffing wireless Internet traffic, to steal, commit fraud, and do damage.

One evolving solution, developed by Internet2 -- a research consortium based in Ann Arbor, MI, that develops advanced Internet technologies for use by research laboratories and universities -- effectively creates a middleman who does the job. Called Shibboleth, the software mediates between a sender and a recipient; it transmits the appropriate ID numbers, passwords, and other identifying information to the right recipients for you, securely, through the centralized exchange of digital certificates and other means. In addition to making the dispersal of information more secure, it helps protect privacy. That's because it discloses only the "attributes" of a person pertinent to a particular transaction, rather than the person's full "identity."

Right now, Shibboleth is used by universities to mediate access to online libraries and other resources; when you log on, the university knows your "attribute" -- you are an enrolled student -- and not your name or other personal information. This basic concept can be expanded: your employment status could open the gates to your company's servers; your birth date could allow you to buy wine online. A similar scheme could give a bank confidence that online account access is legitimate and conversely give a bank customer confidence that banking communications are really from the bank.

Shibboleth and similar technologies in development can, and do, work as patches. But some of their basic elements could also be built into a replacement Internet architecture. "Most people look at the Internet as such a dominant force, they only think how they can make it a little better," Clark says. "I'm saying, 'Hey, think about the future differently. What should our communications environment of 10 to 15 years from now look like? What is your goal?'"

This is the second of a three-part article. The [last section](#), is about an effort by program managers at the National Science Foundation to launch a \$300 million research program on future Internet architectures.

Copyright Technology Review 2005.