



the internet reborn

BY WADE ROUSH

The Internet has transformed the way we find information, shop, and do business. But it's a dumb network built for a bygone age. A university-industry coalition is designing a vastly smarter and more secure Internet: PlanetLab.

Photograph by Beth Perkins

If you're like most cyber-citizens, you use the Internet for e-mail, Web searching, chatting with friends, music downloads, and buying books and gifts. More than 600 million people use these services worldwide—far more than anyone could have predicted in the 1970s, when the Internet's key components were conceived. An estimated \$3.9 trillion in business transactions will take place over the Internet in 2003, and the medium's reach is increasingly global: an astonishing 24 percent of Brazilians, 30 percent of Chinese, and 72 percent of Americans now go online at least once per month.

Still, despite its enormous impact, today's Internet is like a 1973 Buick refitted with air bags and emissions controls. Its decades-old infrastructure has been rigged out with the Web and all it enables (like e-commerce), plus technologies such as streaming media, peer-to-peer file sharing, and video-conferencing; but it's still a 1973 Buick. Now, a grass-roots group of nearly 100 leading computer scientists, backed by



Global ambitions: Princeton University's Larry Peterson wants to make the Internet's infrastructure more intelligent.

heavyweight industrial sponsors, is working on replacing it with a new, vastly smarter model.

The project is called PlanetLab, and within the next three years, researchers say, it will help revitalize the Internet, eventually enabling you to

- forget about hauling your laptop around. No matter where you go, you'll be able to instantly recreate your entire private computer workspace, program

for program and document for document, on any Internet terminal;

- escape the disruption caused by Internet worms and viruses—which inflicted an average of \$81,000 in repair costs per company per incident in 2002—because the network itself will detect and crush rogue data packets before they get a chance to spread to your office or home;
- instantly retrieve video and other

bandwidth-hogging data, no matter how many other users are competing for the same resources;

- archive your tax returns, digital photographs, family videos, and all your other data across the Internet itself, securely and indestructibly, for decades, making hard disks and recordable CDs seem as quaint as 78 RPM records.

These predicted PlanetLab innovations—with the potential to revolution-

ize home computing, e-commerce, and corporate information technology practices—can't be incorporated into the existing Net; that would be too disruptive. Instead, the PlanetLab researchers, who hail from Princeton, MIT, the University of California, Berkeley, and more than 50 other institutions, are building their network *on top of* the Internet. But their new machines—called smart nodes—will vastly increase its processing power and data storage capability, an idea that has quickly gained support from the National Science Foundation and industry players such as Intel, Hewlett-Packard, and Google.

Since starting out in March 2002, PlanetLab has linked 175 smart nodes at 79 sites in 13 countries, with plans to reach 1,000 nodes by 2006. It's the newest and hottest of several large-scale research efforts that have sought to address the Internet's limitations (see *"The Internet's Reinventions,"* p. 36). "The Internet has reached a plateau in terms of what it can do," says Larry Peterson, a Princeton computer scientist and the effort's leader. "The right thing to do is to start over at another level. That's the idea behind PlanetLab."

The Network /s the Computer, Finally

Like many revolutions, PlanetLab is based on a startlingly simple idea that has been around for a long time, advanced most notably by Sun Microsystems: move data and computation from desktop computers and individual mainframes into the network itself.

But this can't be done with today's Internet, which consists of basic machines, called routers, following 1970s-era procedures for breaking e-mail attachments, Web pages, and other electronic files into individually addressed packets and forwarding them to other machines. Beyond this function, the routers are dumb and inflexible: they weren't designed to handle the level of computing needed to, say, recognize and respond to virus attacks or bottlenecks elsewhere in the network.

PlanetLab's smart nodes, on the other hand, are standard PCs capable of running custom software uploaded by users. Copies of a single program can run simultaneously on many nodes around

the world. Each node is plugged directly into a traditional router, so it can exchange data with other nodes over the existing Net. (For that reason, computer scientists call PlanetLab an "overlay" network.) To manage all this, each node runs software that divides the machine's resources—such as hard-drive space and processing power—among PlanetLab's many users (see *"Planetary Pie,"* this page). If the Internet is a global, electronic nervous system, then PlanetLab is finally giving it brains.

The payoff should be huge. Smarter networks will foster a new generation of distributed software programs that preempt congestion, spread out critical data, and keep the Internet secure, even as they make computer communications faster and more reliable in general. By expanding the network as quickly as possible, says Peterson, the PlanetLab researchers hope to restore the sense of risk-taking and experimentation that ruled the Internet's early days. But Peterson admits that progress won't come easily. "How do you get an innovative service out across a thousand machines and test it out?"

It helps that the network is no longer just a research sandbox, as the original Internet was during its development; instead, it's a place to deploy services that any programmer can use and help improve. And one of the Internet's original architects sees this as a tremendously exciting trait. "It's 2003, 30 years after the Internet was invented," says Vinton Cerf, who codeveloped the Internet's basic communications protocols as a Stanford University researcher in the early 1970s


and is now senior vice president for architecture and technology at MCI. "We have millions of people out there who are interested in and capable of doing experimental development." Which means it shouldn't take long to replace that Buick.

Baiting Worms

The Achilles' heel of today's Internet is that it's a system built on trust. Designed into the Net is the assumption that users at the network's endpoints know and trust one another; after all, the early Internet was a tool mainly for a few hundred government and university researchers. It delivers packets whether they are legitimate or the electronic equivalent of letter bombs. Now that the Internet has exploded into the cultural mainstream, that assumption is clearly outdated: the result is a stream of worms, viruses, and inadvertent errors that can cascade into economically devastating Internet-wide slowdowns and disruptions.

Take the Code Red Internet worm, which surfaced on July 12, 2001. It quickly spread to 360,000 machines around the world, hijacking them in an attempt to flood the White House Web site with meaningless data—a so-called denial-of-service attack that chokes off legitimate communication. Cleaning up the infected machines took system administrators months and cost businesses more than \$2.6 billion, according to Computer Economics, an independent research organization in Carlsbad, CA.

Thanks to one PlanetLab project, Netbait, that kind of scenario could





Netbait

BRENT CHUN, INTEL

▶ A Defense Early Warning line for computer worms, using PlanetLab nodes around the world as its eyes and ears.

become a thing of the past. Machines infected with Code Red and other worms and viruses often send out “probe” packets as they search for more unprotected systems to infect. Dumb routers pass along these packets, and no one is the wiser until the real invasion arrives and local systems start shutting down. But in theory, the right program running on smart routers could intercept the probes, register where they’re coming from, and help administrators track—and perhaps preempt—a networkwide infection. That’s exactly what Netbait, developed by researchers at Intel and UC Berkeley, is designed to do.

This spring, the program showed how it can map a spreading epidemic. Brent Chun, Netbait’s author, is one of several senior researchers assigned to PlanetLab by Intel, which helped launch the network by donating the hardware for its first 100 nodes. Chun ran Netbait on 90 nodes for several months earlier this year. In mid-March, it detected a six-fold spike in Code Red probes, from about 200 probes per day to more than 1,200—a level of sensitivity far beyond that of a lone, standard router. The data collected by Netbait showed that a variant of Code Red had begun to displace its older cousin.

As it turned out, there was little threat. The variant turned out to be no more malignant than its predecessor, for which remedies are now well known. But the larger point had been made. Without a global platform like PlanetLab as a vantage point, the spread of a new Code Red strain could have gone undetected until much later, when the administrators of local systems compared notes. By then, any response required would have been far more costly.

Netbait means “we can detect patterns and warn the local system administrators that certain machines are infected at their site,” says Peterson. “That’s something that people hadn’t thought about before.” By issuing alerts as soon as it detects probe packets, Netbait could even act as an early-warning system for the entire Internet.

Netbait could be running full time on PlanetLab by year’s end, according to Chun. “Assuming people deem the service to be useful, eventually it will get on the radar of people at various companies,” he says. It would then be easy, says Chun, to offer commercial Internet service providers subscriptions to Netbait, or to license the software to companies with their own planetwide computing infrastructures, such as IBM, Intel, or Akamai.

Traffic Managers

Just as the Internet’s architects didn’t anticipate the need to defend against armies of hackers, they never foresaw flash crowds. These are throngs of users visiting a Web site simultaneously, overloading the network, the site’s server, or both. (The most famous flash crowd, perhaps, formed during a 1999 Victoria’s Secret lingerie Web broadcast that had been promoted during the Super Bowl. Within hours, viewers made 1.5 million requests to the company’s servers. Most never got through.) Such events—or their more malevolent cousins, denial-of-service attacks—can knock out sites that aren’t protected by a network like Akamai’s, which caches copies of customers’ Web sites on its own, widely scattered private servers. But the question is how many copies to make. Too few, and the overloads persist; too many, and the servers are choked with surplus copies. One solution, described in papers published in 1999 by the researchers who went on to found Akamai, is simply to set a fixed number.

In the not-too-distant future, PlanetLab nodes will adjust the number of cached copies on the fly. Here’s how it works. Each node devotes a slice of its processor time and memory to a program designed by Vivek Pai, a colleague of Peterson’s in the computer science department at Princeton. The software monitors requests for page downloads and, if it detects that a page is in high demand, copies it to the node’s hard drive, which acts like the memory in a typical Web server. As demand grows, the program automatically caches the page on additional nodes to spread out the load, constantly adjusting the number of replicas according to the page’s popularity. Pai says that simulations of a denial-of-service attack on a PlanetLab-like network showed that nodes equipped with the Princeton software absorbed twice as many page requests before failing as those running the algorithms published by the Akamai founders.

This new tool, known as CoDeeN, is already running full time on PlanetLab; anyone can use it, simply by changing his or her Web browser’s settings to connect to a nearby PlanetLab node. It’s a work in progress, so service isn’t yet fully reliable. But Pai believes the software can support a network with thousands of nodes, even-



CoDeeN

VIVEK PAI, PRINCETON UNIVERSITY

► No more World Wide Wait: CoDeeN automatically moves Web pages closer in the network to the people using them.

tually creating a free “public Akamai.” With this tool, Internet users would be able to get faster and more reliable access to any Web site they chose.

But banishing flash crowds won’t, by itself, solve Internet slowdowns. Other PlanetLab software seeks to attack a subtler problem: the absence of a decent “highway map” of the network. Over the years the Internet has grown into an opaque tangle of routers and backbone links owned by thousands of competing Internet service providers, most of them private businesses. “Packets go in, they come out, and there’s very little visibility or control as to what happens in the middle,” says Thomas Anderson, a computer scientist at the University of Washington in Seattle.

One solution is software known as Scriptroute. Developed by Anderson and his colleagues at the University of Washington, it’s a distributed program that uses smart nodes to launch probes that fan out through particular regions of the Internet and send back data about their travels. The data can be combined into a map of the active links within and between Internet service providers’ networks—along with measurements of the time packets take to traverse each link. It’s like having an aerial view of an urban

freeway system. Anderson says operators at Internet service providers such as AOL and Earthlink, as well as universities, could use Scriptroute’s maps to rapidly diagnose and repair network problems in one to three years.

Sea Change

Keeping data intact can be just as tricky as transmitting it: ask anyone who has left a personal digital assistant on a train or suffered a hard-drive crash. What’s needed, says Berkeley computer scientist John Kubiawicz, is a way to spread data

around so that we don’t have to carry it physically, but so it’s always available, invulnerable to loss or destruction, and inaccessible to unauthorized people.

That’s the grand vision behind OceanStore, a distributed storage system that’s also being tested on PlanetLab. OceanStore encrypts files—whether memos or other documents, financial records, or digital photos, music, or video clips—then breaks them into overlapping fragments. The system continually moves the fragments and replicates them on nodes around the planet. The original file can be reconstituted from just a subset of the fragments, so it’s virtually indestructible, even if a number of local nodes fail. PlanetLab nodes currently have enough memory to let a few hundred people store their records on OceanStore, says Kubiawicz. Eventually, millions of nodes would be required to store everyone’s data. Kubiawicz’s goal is to produce software capable of managing 100 trillion files, or 10,000 files for each of 10 billion people.

To keep track of distributed data, OceanStore assigns the fragments of each particular file their own ID code—a very long number called the Globally Unique Identifier. When a file’s owner wants to retrieve the file, her computer tells a node running OceanStore to search for the nearest copies of fragments with the right ID and reassemble them.

Privacy and security are built in. An owner who wants to retrieve a file must first present a key that has been generated using now common encryption methods and stored in a password-protected section of her personal computer. This key contains so many digits that it’s essentially

Change of Address

Dumb routers aren’t the Internet’s only limitation. The profusion of devices with Internet connectivity—laptops, handheld devices, appliances, and soon car and aircraft systems and even remote sensors—is creating a shortage of Internet Protocol (IP) addresses. But a coming overhaul of the address system will result in an almost inconceivably vast number of unique digital addresses.

Today, any device that uses the Internet has an address made up of 32 binary digits. This scheme, laid out in the early 1980s, provided for a maximum of about 4.3 billion unique addresses. Turns out, given the proliferation of devices that’s on the way, that’s not enough. So the coming version of the Internet Protocol, which will be implemented over the next several years, will create codes that have 128 binary digits. A number that long allows so many possible combinations that there will be some 670 quadrillion (thousand trillion) addresses for every square millimeter of Earth’s surface—which ought to be enough for a while.

The Internet's Reinventions

PlanetLab aims to transform today's dumb, simple Internet communications system into a smarter and much more flexible network that can ward off worms, store huge amounts of data with perfect security, and deliver content instantly. Here's how it fits into a long tradition of academic and government research projects that developed fundamental networking, transmission, and distributed-computing technologies.

1969

ARPANET

The first major attempt to use computers for communication, and the testing ground for the standards that would come to define the Internet. Built by universities and technology firms with funding from the U.S. Defense Department's Advanced Research Projects Agency (now DARPA).

1973-1983

The Internet

A network of smaller networks in which computers exchange packets of data formatted and addressed according to, respectively, the Transport Control Protocol and the Internet Protocol (TCP/IP), which were conceived in 1973 and officially replaced the ARPANET's protocols in January 1983.

1992

MBone

The Multicast Backbone: a system that allows many people to view the same real-time information, such as video broadcasts, over the Internet. Created by members of the Internet Engineering Task Force in 1992 to overcome the limitations of standard Internet protocols, which can route a given data packet to only one destination.

1996

Internet2

A consortium of more than 200 universities that has created Abilene, a network of high-performance routers and fiber-optic links. Abilene is able to transmit an entire DVD movie in about 36 seconds, as much as 3,500 times faster than a typical home DSL or cable connection.

The Grid

A collection of public and private organizations and projects that use software developed at the U.S. Department of Energy and the University of Southern California to link scattered supercomputers, scientific instruments, and data storage facilities into a "grid" that can take on tough computational problems—like screening for new drug molecules.

2000

ABone

The Active Network Backbone: a network built to test the efficiency of "active networking," in which the network is stripped of nearly all intelligence—even the basic message-passing software that runs on today's Internet—and packets of data contain all the software and instructions needed to deliver themselves to their destinations. Funded by DARPA and created by SRI International, a private research institute in Menlo Park, CA, and the University of Southern California.

2002

PlanetLab

An effort by academic and corporate networking researchers to augment, and eventually replace, today's "dumb" Internet with a much smarter network able to monitor itself for worms and viruses, relieve bottlenecks automatically, and make personal-computing environments portable to any terminal on earth.

impossible for others to guess it and gain unauthorized access. The key provides access to OceanStore directories that map human-readable names (such as "internet.draft") to fragment ID codes. The ID codes are then used to search OceanStore for the nearest copies of the needed fragments, which are reassembled and decrypted. And there's one more layer of protection: the ID codes are themselves generated from the data's contents at the time the contents are saved using a secure cryptographic function. Like encryption keys, the codes are so long (160 binary digits) that even today's most advanced supercomputers can't guess or fake them. So if data retrieved from OceanStore has an unaltered ID, the owner can be sure the data itself hasn't been changed or corrupted.

Kubiatowicz would like to see OceanStore become a utility similar to DSL or cable Internet service, with consumers paying a monthly access fee. "Say you just got back from a trip and you have a digital camera full of pictures," he suggests. "One option is to put these pictures on your home computer or write them to CDs. Another option is that you put those pictures into OceanStore. You just copy them to a partition of your hard drive, and the data is replicated efficiently on a global scale." That option could be available within three to five years, he predicts, but in the interim, two things need to happen. First, his team needs to produce sturdier versions of the OceanStore code. Second, someone needs to provide enough nodes to enlarge the system to a useful scale. That someone is likely to be a private company looking to enter the distributed-storage business, predicts Peterson. "I could imagine OceanStore attracting the next Hotmail-like startup as its first customer," he says.

Beyond providing distributed, secure storage, OceanStore could eventually make every computer your personal one. At its next level of development, it could store your entire computing environment—your PC desktop, plus all of the applications you're running and all the documents you have open—across the network and reconstitute it on demand, even if you popped up at an Internet termi-

nal halfway around the world. This capability would be useful to the businessperson on the road, to a doctor who suddenly needs to review a chart, or to a contractor who wants to tweak a blueprint from home. Several companies are working to realize this vision. Intel calls it Internet Suspend/Resume, and Sun researchers are testing several approaches to “desktop mobility.” But PlanetLab could provide the infrastructure that makes such technology possible, by offering a means to manage the large amounts of data—perhaps tens of gigabytes—that personal-computer users might regularly rely on.

Laundry List

Such ideas may seem radical. Then again, just a decade ago, so did e-commerce. The question now is which big idea will evolve into the Google or Amazon.com of the new, smarter Internet. By charter, PlanetLab can't be used for profit-making enterprises, but businesses may soon spring from the platform it provides. “We want it to be a place where you leave services running long-term—which brings us much closer to the point where someone commercial might want to adopt it or replicate it for profit,” Peterson says. That could happen if the experiments running now, along with the methods being developed to keep the network operating smoothly, provide a reliable model for future intelligent networks. “We don't know where that next big idea is going to come from,” says Peterson. “Our goal is just to provide the playing field.”

PlanetLab's early industry sponsors, such as Intel and Hewlett-Packard, may be among the first to jump in. HP Labs in Palo Alto, CA, for example, installed 30 PlanetLab nodes in June and plans to use the network to road-test technologies that could soon become products. One example: software developed by researcher Susie Wee that uses a CoDeeN-like distribution network to deliver high-resolution streaming video to mobile devices. The goal is to avoid wasting bandwidth, and Wee's software would do just that by streaming, say, video of a major-league baseball game to a single local node, then splitting the data into separate streams optimized for the screen resolutions of different viewers' devices—whether desktop PCs,



OceanStore

JOHN KUBIATOWICZ,
UNIVERSITY OF CALIFORNIA, BERKELEY

► Who needs disk drives when your data can be broken up, scattered across the Internet, and reassembled at will?

wireless laptops, PDAs, or cell phones. HP or its licensees could bring such a service to market within two years, Wee says. Projects like this one, says Rick McGeer, HP Labs' scientific liaison to a number of university efforts, means that PlanetLab is “not only a great experimental test bed, it's a place where you can see the demonstrable value of services you don't get on today's Internet.”

Of course, researchers' enthusiasm about smart networks doesn't keep them from pondering the new problems they could create. Until now, viruses and worms have always been launched from machines at the Internet's edges; imagine how much more damage an attack could do if it originated from a trusted node inside the network. And there's no centralized authority to force local PlanetLab machines to meet security standards, as there is with Akamai and other private networks. But researchers at Princeton and other PlanetLab member institutions say they're already working on ways to avoid these hazards.

While it's impossible to know which blockbuster new technology and business paradigms will emerge from smarter networks, projects like PlanetLab virtually ensure that the Internet will eventually fulfill some of its long-

unrealized potential in areas like broadband access, security, shared storage, and reliable video, text, and other content delivery. “There is a long laundry list of things we can and should do better on the Internet,” Internet pioneer Cerf says. “Why didn't we do it before? Well, some of it is that they are hard problems; some of it is because the technology wasn't capable enough—we needed more brute-force computing capability than we had 20 years ago. And in some cases, it's because nobody cared.”

That's now changing. Peterson expects that ultimately PlanetLab and similar networks will bring about a wholesale reinvention of the Internet. As smart nodes are installed at more of the Internet's existing hubs, these networks could multiply to the point that they cease to be add-ons at all and simply *become* the next generation's Internet. As Peterson puts it, “This is exactly the Internet all over again.” The results could be as different from e-mail and Web browsing as those technologies are from the telephone—or a 1973 Buick is from a low-emissions, fuel-efficient Toyota—with impact to match. ■

Wade Roush is a *Technology Review* senior editor based in San Francisco.