MUSINGS ON THE INTERNET

By Vinton G. Cerf

I would like to try in this article to offer an idea of what's going on, as I see it, in the Internet world—particularly its technology, but also its policy, economics, and philosophy. Let me start by giving some figures. The number of users on the Internet has surpassed a half-billion. The most recent number I have is 544 million users, and that was in February 2002. We're well past that now. The second interesting statistic is that as of January 2002, there were about 147 million hosts with domain name system (DNS) names on the Internet. These are usually servers—could be e-mail or Web servers and so on. What's important is that they all have domain names, and they have dedicated IP (Internet Protocol) addresses. The number of clients on the Net—the laptops, the PDAs, and things like that—likely exceeds 300 million. By any reasonable accounting, that's a fairly large collection of systems. On the other hand, the telephone system has about 1.3 billion terminations—about 800 million wire-line terminations and about 500 million or so mobile phones—so plainly, it's still much larger than the Internet. However, the two are beginning to come together in interesting ways.

Vinton G. Cerf is the recipient of the 2002 Paul Evans Peters Award, presented by the Association of Research Libraries and EDUCAUSE, sponsoring organizations of the Coalition for Networked Information. This article is drawn from his speech delivered at the award presentation. Known as a "father of the Internet" for his work in co-designing TCP/IP protocols and the architecture of the Internet, Cerf is Senior Vice President of Internet Architecture and Technology for Worldcom. He has also played key leadership roles in the Internet Society, the Internet Corporation for Assigned Names and Numbers (ICANN), and the Internet Societal Task Force. Besides advancing fundamental Internet technologies and applications, he has eloquently addressed the social and policy issues raised by global networking. Cerf is widely appreciated for his ability to envision, engineer, and express the intricacies and vast potential of networked communications and technologies.

Current Developments

Regarding the technology of the Internet, a couple of things are very clear. First is the movement toward much higher speeds. At the edges of the Net, the most desirable high-speed service is gigabit Ethernet. It's not widespread yet, but my expectation is that this will become the most popular way of getting access to the Net. What's interesting about gigabit speeds is that you can do things at that rate that you wouldn't consider doing otherwise. For example, an hour's worth of an HDTV program, when it's properly encoded, takes up about 16 gigabits (2 gigabytes) of data. You could transmit that in 16 seconds, at a gigabit per second. So the idea of video-on-demand changes its character now. It's not video-on-demand when it has to be transmitted to you for an hour while you watch; now, in 16 seconds, you can grab a copy and off you go. You can look at it another time, at your leisure. So that changes the way we think about moving information around. Of course, in 16 seconds, you could also have received 2 gigabytes of software, or 2 gigabytes of a subset of a database that you're interested in working with locally.

In the core of the Net, we're already operating at 10 gigabits per second. The fibers that carry those optical signals are capable of carrying anywhere from 160 to 320 10-gigabit-per-second channels per fiber. We've already tested a 1.6-terabit-per-second system, carrying 160 colors on a single optical strand. The next technology up is 40 gigabits per second per color, and there is some debate about which of the two—10 or 40—is better. The reason there is a debate at all is that the higher-speed stuff requires repeaters to be put into the Net at closer intervals; that costs money, takes up space and power, and may not scale very well, so today 10 gigabits per second per color is the preferred way of supplying capacity to the high-bandwidth core of the Internet.

Technically speaking (I'm an engineer, so I tend to focus on a lot of the little details), one of the most important aspects of Internet technology is that it's packet-switched. The reason this is important is that the devices we carry around now, such as laptops, typically have more than one program running concurrently. And when you're in a network environment, those programs may be interacting with any number of other programs on the Net. You need a packet-switched interface to multiplex the traffic smoothly over one high-bandwidth pipe, and then of course the packets have to be switched through the network to get to their destinations. Unfortunately, this notion of packet-switching runs into a serious scaling problem as we move down closer to the core of the network, where the speeds are going up.

Let me give an example. I've taken some liberties here, since we don't have channels that are actually running at a trillion bits per second per channel, but let's pretend for a moment that we get to the point where we can send a packet at a trillion bits per second. These packets are often one thousand bits long. It takes only one nanosecond to send a thousand-bit packet at a terabit per second. That doesn't give a whole lot of time down in the core of the Net to look at the packet to decide what to do with it. In fact, today's computers run at one or two gigahertz, one or two billion cycles per second, providing perhaps one instruction to look at the packet and do something in a nanosecond. That's clearly not enough. So we really don't like the idea of having to look at packets down in the core of the Net because there isn't enough time to do anything with them. We'd like to get to the point where the packets have been moved onto a particular wavelength on an optical fiber so that we don't have to look at packets—we only need to know the color of the light that the packet is being carried on. If it's green, it's going to go this way, and if it's red, it's going to go that way. We're switching at the optical level; we're switching photons according to their color, but we don't need to know anything about packet boundaries, so the effort required is far less. Obviously, if I need packets at the edge, and I want to do photon-switching in the middle, something has to happen in between that matches the packets with the right color that they should be on in the core of the Net. The telephone world has that same problem; it's called *grooming:* moving the circuits over so that they land on the appropriate high-bandwidth transport in the middle. The solutions that the telephone network has evolved over time to groom traffic into the right high-bandwidth pipe will also apply in the Internet world. Here's a case where the two technologies actually can benefit from each other.

Another area of interesting evolution is wireless communication. We're finding more and more opportunities to transmit data over the air. Ericsson has estimated the number of mobile devices that will be Internet-enabled by 2006 to be 1.5 billion devices. I had estimated that the number of servers and other clients on the Net, not counting mobiles, would be on the order of 900 million devices by 2006. That's a total of 2.4 billion devices on the Internet by 2006, which is now just four years away. But in addition to those mobiles that are using either GSM or third-generation wireless communication technology, there is also this kind of Ethernet in the sky: the 802.11a and 802.11b, or Wi-Fi as it's sometimes called, that's also very popular. I use it at work. I

INTERNET

Technically speaking, one of the most important aspects of Internet technology is that it's packet-switched.

> **People who think peer-to-peer is a new idea should understand that when we were designing the Internet, that was a critical and core part of the design.**

also use it at home, where it has some surprising social effects that I had not anticipated. It used to be that when my wife and I would have dinner, we would sit at the dining-room table, and we would have conversations. Often we would get stuck at some point because we didn't know some fact that we were groping for in the course of the conversation. Either we had to go on to some other topic, or we had to stop and go down the hall to our library. Well, when we got the wireless connection into the house, we could be anywhere in the house, and our laptops would pick up the Internet. So we left a laptop on the dining-room table, not to read e-mail over dinner but simply to consult the Internet to find a particular fact so that we could continue our conversations. I continue to be astonished not only at the amount of information that is on the Net but also at its quality. Yes, there is a lot of junk on the Net, but there is also an amazing amount of good-quality stuff.

I remember one particular conversation we had about where the steak we were eating had come from. Yes, we know it comes from cows, but the real question was, how did it get to our table—what was the path? And we eventually got to the point where we were saying, "Remember the cattle drives of the late 1800s and the Chisholm Trail?" We remembered the Chisholm Trail. But we didn't remember where it went—we just remembered the name. So we looked up "Chisholm Trail" and "cattle drive" on the Internet, and up comes a thirty-five-page report by some student, probably a high school student, with maps and everything else describing all of the cattle-drive activity. What really surprised me is that our whole mythology of the Wild West cowboys lasted only about twenty years, because the cattle drives were replaced by trains and then, not very long thereafter, by trucks and highways. So that whole era actually occurred over a very brief period of time. In any event, the Internet has transformed and made more interesting many dinner conversations.

Another wireless element is Bluetooth (which I thought sounded like a dental condition the first time I heard it). This is a low-power, moderate-bandwidth wireless signal. It is only a million bits per second, which is a little too slow for today's technology, but the idea is to allow small devices, like printers and computers and other desktop devices, to interact at the desktop level without wires. Looking at my desk, with all the wires all over everywhere, this seems like a good idea. Eventually this stuff will be small enough that someone could probably put it into a pair of cufflinks, and that way the left hand and the right hand could coordinate with each other—a major contribution!

Another technological development involves the idea of presence. In the telephony world, there is an assumption that the telephone is always there. You call the phone number, and the phone rings. In the Internet world, we are still only intermittently connected to the Net. Some 80 percent of people get on the Net by dialing up through a telephone. So if you're not always on, presence becomes an important notion. When you are on the Internet, you may want other people to know that you're on, so that they can get in touch with you. Instant messaging has thus become an incredibly popular tool. And of course the notion of peer-to-peer computing has also become popular. I have to point out, though, that people who think peer-to-peer is a new idea should understand that when we were designing the Internet, that was a critical and core part of the design. The TCP/IP protocols are precisely designed so that all elements in the Net are essentially of equal standing. A small mini-computer and a laptop and a big supercomputer have equal standing in the Internet world, which is how it is possible for a supercomputer to talk to a laptop without overwhelming the laptop. The two have flow control between them.

Another interesting idea is something that Bob Kahn and I called "knowbots": knowledge robots. We published our ideas about knowbots in 1988. Although things didn't evolve precisely according to our description, something has evolved. Java virtual machines, things that will interpret a high-level code, are a good example. Another programming language, called Python, is also a very popular high-level language. The important point is that these programs can be moved from place to place, and our notion of moving software around to the databases to do work seems to be emerging now in this Java environment.

Another very visible development is voice over IP. And of course, for us at Worldcom, this has become a very important notion. Now that it has been demonstrated that voice can be carried over the Internet, a lot of companies are putting their voice up on their virtual private networks that they run or that we run for them privately. This is very good for them because now they don't have to have two separate networks; they just have one, which carries both data and voice.

There is an initiative in the Internet Engineering Task Force called ENUM. What ENUM does is bring together the public-switched telephone net world and the Internet world by taking conventional telephone numbers, sometimes called E.164 numbers, and putting them into the domain name system. Those elements in the domain name system are mapped, as is every domain name, into something. Most domain names get mapped into Internet addresses. In the ENUM world, the mapping is not to an Internet address but to a URL or to a URI. The idea here is that mapping a telephone number into some URL on the Net gives tremendous flexibility in associating the phone number with a variety of ways that someone might be reached: a phone number, a fax number, an e-mail address, a Web page, or any program that's willing to interact using the URL that the E.164 telephone number has been mapped into. That concept is almost certain to transform the way in which we think of telephonic communication because now your telephone number becomes a path to reach you by any of a variety of different methods. And of course, not surprisingly, a lot of arm wrestling is going on over who has responsibility for that particular part of the domain name system. ICANN, where I serve as chairman, is of course square in the middle of all that hollering and yelling. Eventually it will get sorted out.

Internet-Enabled Appliances

Internet-enabled appliances will have a huge impact. We're already starting to see a great many automatic devices. In fact, one of the games I like to play as I think about what's happening in the world is to imagine that I've been transported back in time to the world of 1952. I used to think that if we were all transported back fifty years, by the end of the week many of us would have broken noses. It wouldn't be because we would get into altercations; it would be because we would run into doors that we thought were going to open on their own. Today, a number of devices do things for themselves. Toilets flush automatically. Faucets turn on and turn off. Lights go on and off when someone walks into or out of a room. People in 1952 would really think we were crazy to assume that such things were normal.

I have some favorite appliances that I am looking forward to seeing on the Net. One of them is a videocassette recorder (VCR). The reason is simple. I can't program mine. I usually have to find an eleven-year-old to do it for me, but I've run out of eleven-year-olds at home, so what I would like is to have the VCR up on the Net and be willing to accept controls through the Internet. What's important about this is that we have also reached the point where speech recognition is considered a reasonable tool. For example, we can call an airline, speak to a computer, and find out the status of a flight. Recently I was driving to Dulles Airport, and I could not remember exactly what time a plane was landing. So I called the airline on my cell phone. The computer on the other end asked for the flight number, and I said that I didn't know the flight number. The computer then asked, "What time is the plane arriving?" I said 9:40. The computer responded: "I thought you said 9:14. Is that right?" I said, "No." It said: "Oh, you must have said 9:40. What city is it coming from?" I said, "San Francisco." It said, "That's flight number 222, and it's twenty minutes early." At that point, I accelerated to try to get to the airport on time. The interaction was very natural, and the reason that it worked is what is called "speaker-independent speech recognition." I had never practiced with this system; it didn't know me from Adam. But it worked because at each point in the exchange, my responses were fairly predictable. They were from a small set of vocabulary items, making it easier for the computer to recognize the speech. We've been able to do this for a long time—it's just that we've never been able to do it in real time. Having much more powerful computers has allowed us to do speech recognition in real time on the Net. This translates into the possibility that speech will enable a variety of devices that are willing to accept control on the Net. So, for example, I should be able to pick up the phone and say, "Please record *Star Trek* at 10 o'clock on Sunday night," and the computer, which understands that, will turn around and tell my VCR what to do.

There is an Internet-enabled refrigerator made by Electrolux in Sweden. It has a liquid-crystal touch-sensitive display and is a nice addition to today's household communication, which is made up mostly of paper and magnets on the front of the refrigerator. You can use it to send e-mail and surf the Web. Some people have suggested that the refrigerator should include a bar-code scanner as well. The refrigerator would know what you've put into it, because it can read the bar codes. So now, while you're away at work, the refrigerator-computer can look around the Net for recipes that it knows can be made with what it has inside. When you come home, you'll see a nice display of recipes it has suggested. Or you will be away on vacation, and you might get an e-mail from your refrigerator saying, "I don't know how much milk is left, but you put it in there three weeks ago, and it's going to walk out on its own." Or you might be shopping and get a page from your refrigerator saying: "Don't forget the marinara sauce. I have everything else I need for dinner."

The Japanese have made an Internet-enabled bathroom scale. The idea is that you step on the bathroom scale, and the information goes to the doctor to be part of your medical record. Of course, the obvious question is what happens if the refrigerator gets the same information. Then you may come home to find only

INTERNET

Having much more powerful computers has allowed us to do speech recognition in real time on the Net.

diet recipes on the refrigerator display. Or the refrigerator may simply refuse to open because it knows you're on (or should be on) a diet. Horrible thought!

I recently opened a bottle of wine, something I do as often as I have an opportunity. And I noticed that burned into the cork was "www.louismartini.com" (obviously a fairly recent vintage). I started thinking about the possibility of an Internet-enabled cork. Suppose that there were a passive-memory element embedded in the cork and that the passive-memory element could be activated. You could store things in the memory when you're in the proper environment. Information about the history of that particular bottle of wine could be incorporated into the cork. So, for example, you could keep track of the temperature and the humidity ranges. You could record where it was stored and bottled. You might even be able to record which merchants had transported and sold it. So when you pull the cork, if you're not happy with the wine, at least you have a clue about its history. The same process could be applied to all kinds of manufactured objects or to anything else—perhaps an object in a museum. The idea of having passive elements that can be at least enabled through interactions with a computer adds interesting possibilities to what we would normally not think of as being part of our information world at all.

The last thing I wanted to mention about Internet-enabled appliances is clothing that can detect your vital signs without any invasive requirements. Clothing could measure oxygen levels in the blood, pulse rate, heart rate—with all the information then transmitted over the air, presumably suitably encrypted, for medical observation purposes. As I started thinking about the notion of Internet-enabled clothing, I thought the following scenario was rather appealing. Let's suppose my socks are Internet-enabled. So now I can interrogate the sock drawer, and I can get back a report saying that there are seventeen matched pairs of socks in the drawer and there is one unmatched pair: sock number 144L is missing. I could then send out a multicast around the house, and a report would come back, saying: "This is sock number 144L. I'm under the sofa in the living room." We've now solved the problem of the missing sock, which I think is another major contribution to society. Of course, like any idea along these lines, this involves risk. Another scenario came to mind—one not necessarily as pleasant. Imagine that some guy calls home and says, "Hi, honey, I'm going to work late in the lab tonight." There's a brief pause, and then she says, "Well, that's really interesting, dear, because your shirt's down at 19th Street at the bar." So maybe it's not a good idea to have Internet-enabled clothes and we shouldn't go there.

Policy

I want to touch briefly on policy issues, because it's my belief that the policy problems that are arising in the context of the Internet and, more generally, networked information are as important as—and probably harder for us to deal with than—the technical ones, which are already a fairly significant challenge. One area is intellectual property, which has become quite a major conundrum in the context of networked systems. For example, there was a big debate over the one-click buying patent that Amazon.com filed, followed by an argument over whether the process was or was not a legitimate thing to patent. Another example is the unsolvable conflict between trademarks and domain names. The reason I think it's unsolvable, at least in the current form, is that trademarks are not required to be unique. Two companies can have legitimate rights in a given trademark, whereas domain names are designed to be unique. Two different parties cannot hold the exact same domain name because that's not the way domain names work. So, unfortunately, the two concepts are fundamentally in conflict with each other, and it isn't clear how that's all going to be sorted out other than by arbitration procedures and litigation. I don't see any technological solution to the problem.

Another area where I think our concerns are particularly highly focused now is the tension between individual privacy and the need to expose some information in order to protect society from people who don't mean that society well. I recommend Amitai Etzioni's book *The Limits of Privacy,* which fairly lays out the existing tensions. Etzioni believes that it's possible to go too far in one direction, overly protecting individual rights and leaving society exposed as a consequence. Of course, it's also possible to go too far the other way—everything is revealed, and therefore all kinds of bad things are prevented from happening, but the result is a kind of society that perhaps none of us want to live in. Anyone who has been through a fairly extensive security search in an airport probably has some sense of what it would be like to live in a society that had gone pretty far in the other direction. I'm not arguing, by the way, that there is something fundamentally wrong with the current security procedures, but they are so extreme relative to the previous experiences over the last couple of hundred years in the United States that

INTERNET

> **The idea of having passive elements that can be at least enabled through interactions with a computer adds interesting possibilities.**

it's a little unnerving to think that this may be the norm for decades to come.

Economics

Some of the current market slump is the result of the market's reaction to and nervousness about the telecom industry. It's partly a problem of perception. When no one has a reason to buy, stocks tend to go down because somebody will sell and stocks tend to follow the sellers down. But if investors were to go and look at the basics of many of these companies, I think they'd find fairly healthy balance sheets and healthy cash flow. Perhaps what will happen now is a realization that there really is a persistent demand, at least for Internet service, which is still growing at a factor between 1.4 and 2 per year in terms of demand for capacity, as it has been since 1988. It's just that we went through this crazy period—the dot-com boom—when suppliers couldn't tell whether somebody who ordered their service had a good business model or not. If you're a supplier, when a company orders something, you don't usually say, "Well, I'm not going to sell this to you unless I can have a look at your business model and see whether or not you're going to survive for the next five years." No, you just take the order, and you do whatever you can. The problem is that everybody did that, and the companies—many of them, not all of them, but many of them—had huge amounts of capital to spend from their inflated initial public offerings, and the people who were running those companies somehow seemed not to understand the difference between capital and revenue. Capital is eventually used up, whereas revenue is supposed to keep coming in. And so they spent all their capital. Meanwhile the suppliers had based their prices on the presumption that they would be selling the service for a period of time, and when the companies ran out of capital, they said, "Sorry, we

INTERNET

The policy problems that are arising in the context of the Internet are probably harder for us to deal with than the technical ones.

don't need the service anymore." Poof! There went the suppliers' income. But that horrible period of mania has passed. The venture-capital folks, of course, have pulled in their horns.

All of this is, in a sense, healthy, because I think that now we're going to go back to the real basics in economics—for example, Rule Number One: make more money than you spend. I think that we are entering a period of reality check, and the health of the telecommunications business is fundamentally still sound. The one place that is really going to be interesting is voice over IP, because Internet services have always been flat-priced. And I don't see that changing, even as we put more and more voice over the Internet. It too will become flat rate.

Philosophy

We could spend many hours, if not days, discussing the philosophy of the Internet. It is populated by such a wide range of people that one would not expect to find consensus on its philosophical underpinnings. However, there are some key facts.

The design of the Internet treats all hosts as if they were equal in the eyes of the protocol. That is, PCs and supercomputers and Internet-enabled picture frames all look equal. There is no a priori client-server relationship *imposed* on the parties.

The Internet is designed around the concept of an open and layered protocol architecture so that various applications can be constructed and can use any of the exposed layers of protocol.

The IP layer is intended to rely as little as possible on the details of the underlying transport (transmission) system so as to accommodate as many ways to carry IP packets as possible. This principle is showing some signs of wear as various classes and qualities of service are becoming of more interest and rely on knowledge of the underlying system.

Although it is becoming increasingly apparent in the business use of the Internet that security is a key requirement, it is also the case that the ability to reach all parts of the Internet on an end-to-end basis is a highly valued characteristic. This end-to-end principle puts the edge devices largely in charge of their own destiny and makes them responsible for their own protection. That has not stemmed the popularity of firewalls and so-called Network Address Translation boxes, but it does remain an important philosophical design principle.

In the end, the popularity and the utility of the Internet are largely a consequence of its accessibility and openness. The influx of content on the Internet and its openness to new functions are largely a consequence of the ability of users to participate directly. Much of this can be attributed to the ease of use and the ease of creation of content to go over the World Wide Web application of the Internet. In simplifying access to the content of the Net and in creating standards for content representation, Tim Berners-Lee opened up a world of information for all the users of the Internet. There are still myriad applications waiting for the spark of animation, and these will come as access to the Internet continues to spread to the world's population. Software creativity knows no boundaries, and the Internet benefits from the unfettered richness of the human imagination. *e*