INTERNET TRENDS

Internet Founder Ponders the Web's Future

An interview with Lawrence Roberts
CTO and Chair of Caspian Networks

A leading authority on packet switching and network architectures, Lawrence Roberts led the team that designed and developed Arpanet, the world's first major computer packet network and precursor to the modern Internet. Widely recognized as one of the Internet's founding fathers, Roberts launched Telenet, the world's first packet data communications carrier, in 1973. The carrier subsequently became Sprint's data division. Roberts also made major contributions to asynchronous transfer mode (ATM) technology, which underlies much of today's Internet core. He has received numerous awards, including the L.M. Ericsson prize for research in data communications. To discuss the continuing evolution of the Internet, *IT Professional* interviewed Roberts at his Caspian Networks office in San Jose, California.

> ***The developer of Arpanet, the Internet's precursor, sees intelligent, IP-based networks as the wave of the future.***

LESSONS LEARNED

IT Pro: What did you learn from your pioneering work with Telenet, the first commercial packet network?

Roberts: Most of the basic experimentation on the Arpanet had been done by 1973. I realized then that we had to make the network commercially viable. Everybody said it was great technology, but it depended on government funding. Without that funding it would probably never grow or even survive.

So I went to AT&T and asked if they'd take over Arpanet, but they declined. I felt a strong need to prove packet switching commercially viable, so I started Telenet, which worked out well. The economics worked for long-haul access. They didn't work yet for computer-to-computer operation, but worked effectively for terminal-to-computer operation.

IT Pro: What lessons did you learn from your pioneering work with ATM?

Roberts: I started with the issue of flow control, because network flow control had always been based on TCP (transmission control protocol), which was extremely slow coming from the user end. It takes 10 to 20 seconds for TCP to get up to speed because of its slow-start behavior and because the network does not cooperate intelligently.

Dumb networks were a reasonable assumption in the early days when you did not know where the network was, who was building it, and so on. But the network could do more than just support flow control because it knew when it was congested. If it could tell someone it was congested before it dropped packets, then it might operate more intelligently.

So I started working with ATM. I thought the technology offered a viable approach, one that worked extremely well for flow control. All the research I did showed that ATM offered 100 times more effective flow control than did TCP, in terms of delay, and that ATM kept the delay stable instead of letting it grow with the size of the network, as TCP's delay does. But ATM wasn't something that could integrate easily with IP (Internet protocol) networks because the network didn't use the technology end-to-end, just at its core.

In the 1980s, when we decided on ATM's packet size, we lacked a viable way to do variable-length packets in hardware, so clearly we had to do

1520-9202/00/$10.00

something to make high-speed switching possible. Fixed-size packets seemed a reasonable solution.

The problem proved to be settling on a fixed size because the packets had to be small enough for voice transmission. If you could make the packets variable in size, so that they could get small for voice and big for data, you had a much more efficient system. That's what IP does. So, ultimately, I found that IP provides a better vehicle for network flow control than ATM, despite IP's shortcomings.

IT Pro: Will IP be the best choice for tomorrow's networks?

Roberts: It's a better vehicle for integrating everything because users use it now and it does what they need. Given those realities, we must attack quality of service so that IP can support low delay and low-delay variation. ATM does those tasks well, but requires an immense change to the IP concept because if the network is dumb, it won't support ATM. IP also offers the advantage of flexibility without requiring that you change standards or the external interface.

TODAY'S INTERNET

IT Pro: What shortcomings do you see in today's IP-based Internet?

Roberts: First, we need to do something about flow control. TCP's flow control is extremely slow, which causes enormous loading times for Web pages. We can fix such

Roberts Recalls Designing the Internet

***IT Pro:* How did you become involved with the idea of networking computers?**

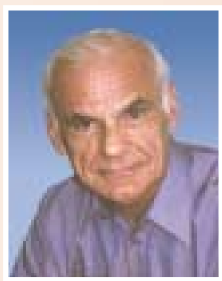

***Roberts:* I talked with ARPA (Advanced Research Projects Agency) Director J.C.R. Licklider while I was at MIT in 1964. He convinced me to attack the problem of connecting computers together and building the information network so that we could get at all the information in all the computers.**

That capability struck me as something tremendously valuable, new, and different. I did experiments at MIT, connecting two computers together to uncover the problems, which all proved to be communications related. So I incorporated the packet switching Len Kleinrock had been working on at MIT.

***IT Pro:* Didn't the experts insist that packet switching would never work?**

***Roberts:* Yes, but Kleinrock had written a book with all the queuing theory necessary for packet switching. That became fairly critical as I got into it because if you didn't know it was going to work, it was pretty hard to argue against all the communications scientists who said my approach was crazy. They predicted that I would run out of buffers. But the computer scientists thought it was a great idea.**

So I joined ARPA in December 1966. In 1967, I designed the network to run over 50 kilobit-per-second links. I tried to integrate everything everybody had done, although most of what I needed came from Kleinrock's work in queuing theory. I did the network topology design, bought the lines, designed the basic network and all the economics of it, and worked with the participating universities.

***IT Pro:* Were the university scientists and engineers skeptical of this new idea?**

***Roberts:* Some didn't want to be involved because they thought it would take their computer time away. The goal was to do research sharing and connect computers so that everybody could share data. So they gradually found that they could collaborate, which was tremendously more effective, and could use each other's computers during off-hours.**

As I studied the economics, I realized that in the future the network would be far better for voice, video, and messaging than telephone circuit switching, and that we ought to be thinking about using packet switching for everything.

***IT Pro:* When did you come to that conclusion?**

***Roberts:* While planning the network in 1968, I saw some of the economics, and in 1969 as we built it I saw a lot more. By 1971, as the network was running, I not only saw the economics—and had published most of that information by then—but I also realized how attractive e-mail could be. We had ignored e-mail because of our focus on research sharing. Ray Thomlinson set up the file transfer programs to send messages from one computer to another, and I wrote the first e-mail handler that let us save, forward, reply, read, and perform other functions, just like today's e-mail programs. It wasn't too difficult to do that once we had the basic mechanism of moving the message on the network.**

delays easily if we change TCP's slow-start behavior and make the network cooperate. As I showed with ATM, it can be far better if you have the network indicate what speed it can support and you can jump to that speed immediately.

Second, you need quality of service for voice and video over IP. There is no reason why IP cannot support that. Developers come up with rules that may be applicable to one technology but not the next. They keep applying such rules—like "the network should be dumb"—over and over. But a dumb network doesn't do much. Equipment can be intelligent as long as you don't go beyond its scalability or that of the network.

Third, IP networks suffer from unreliable routing: Any failure results in an enormous convergence time that destroys all connectivity in the network while the convergence lasts.

Fourth, we must address scalability. Throughout the Internet's development, we've developed switches in accord with Moore's law, which asserts that computing capacity doubles every 18 months. We put our resources into tackling delay problems because historically, traffic was too light to pose a major problem. In 1997, however, traffic caught up with the Internet and became its driving force, growing fourfold each year—more than double the speed of Moore's law. Now switches must become big very quickly—you must grow the switches' capacity by a factor of 1,000 every five years. Switches must scale to that level for someone to buy them, and making that sort of progress is hard.

Fifth, we must improve management. Today, managing a classical IP network involves hourly activities. If a route becomes overcrowded because of intense Web activity, you must manually tune the network during the middle of the day. Managing today's Internet has begun to resemble what operators did for the telephone network in its early days—if you let current trends progress for just a few more years, you'll have everyone in the world managing the Internet. The network must be more intelligent so that it can manage itself. It must do routing better. Routing must balance the load across the network, handling failures quickly and effectively without wasting bandwidth.

All these challenges point to what I believe will be our most critical tasks: improving the Internet's intelligence and making the network scalable.

IT Pro: Will diffserv (IP differentiated services) improve quality of service? Will MPLS (multiprotocol label switching) solve the routing or scalability problems?

Roberts: The diffserv effort represents an attempt to stay dumb, to know a little bit about class of service, and to try making it work. If you study queuing theory, you can prove that diffserv won't work. Nobody who follows this "dumb network" philosophy really believes they can get quality of service without throwing bandwidth at it, and queuing theory tells us you can't solve the problem that way either. You must solve it by having more intelligence in the system.

Nor is MPLS a panacea—it's just another way to set up a pipe. That's not the right thing to do in the long run, although it may help us get around routing problems in the short term. The routing problem causes the need for these pipes, that and the scaling problem implicit in building large networks. We should attack those challenges directly rather than pursue solutions that add extra layers of protocol and management.

Making Networks Smart

The public telephone system exemplifies an intelligent network. It embeds intelligence in complex, software-controlled switches and an elaborate signaling system. The network establishes a connection and maintains state information for every voice call, fax transmission, and modem linkage. Reliable and congestion-free, the network lacks only agility.

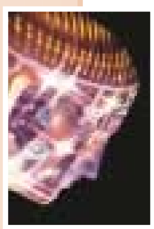

IP networks, like the Internet, don't maintain state information, so intelligence resides at the network's edge. David S. Isenberg ("The Dawn of the Stupid Network," *ACM Networker*, February/March 1998, http://www.isen.com/papers/Dawnstupid.html) writes: "In a Stupid Network, if you have congestion, you just add more connections, bandwidth, switching or processing power. If you want reliability, you add more routes or more redundancy. If you need more intelligence for features or services, you add it at the endpoints. [...] Data gets to where it must go adaptively, with no intelligence and no features."

Adding bandwidth, processing power, and routes to stupid networks helps, but only to a point. Some analysts advocate a middle ground, in which equipment and protocols acquire intelligence so that they can manage the network without managing each session. The network must know when and where it is congested, what rates it can support, how to rapidly route around congestion and failure, and how to balance loads. This awareness requires more state information than that found in today's Internet, but far less than in the telephone network.

TOMORROW'S INTERNET

IT Pro: The problems you describe sound insurmountable.

Roberts: I think they can be solved. Intelligence is something that electronics can give us virtually for free. Networks can be more intelligent, but simultaneously must do it more simply

Figure 1. This poster, a stylized snapshot of the Internet on 1 January 2000, hints at the global network's staggering complexity.

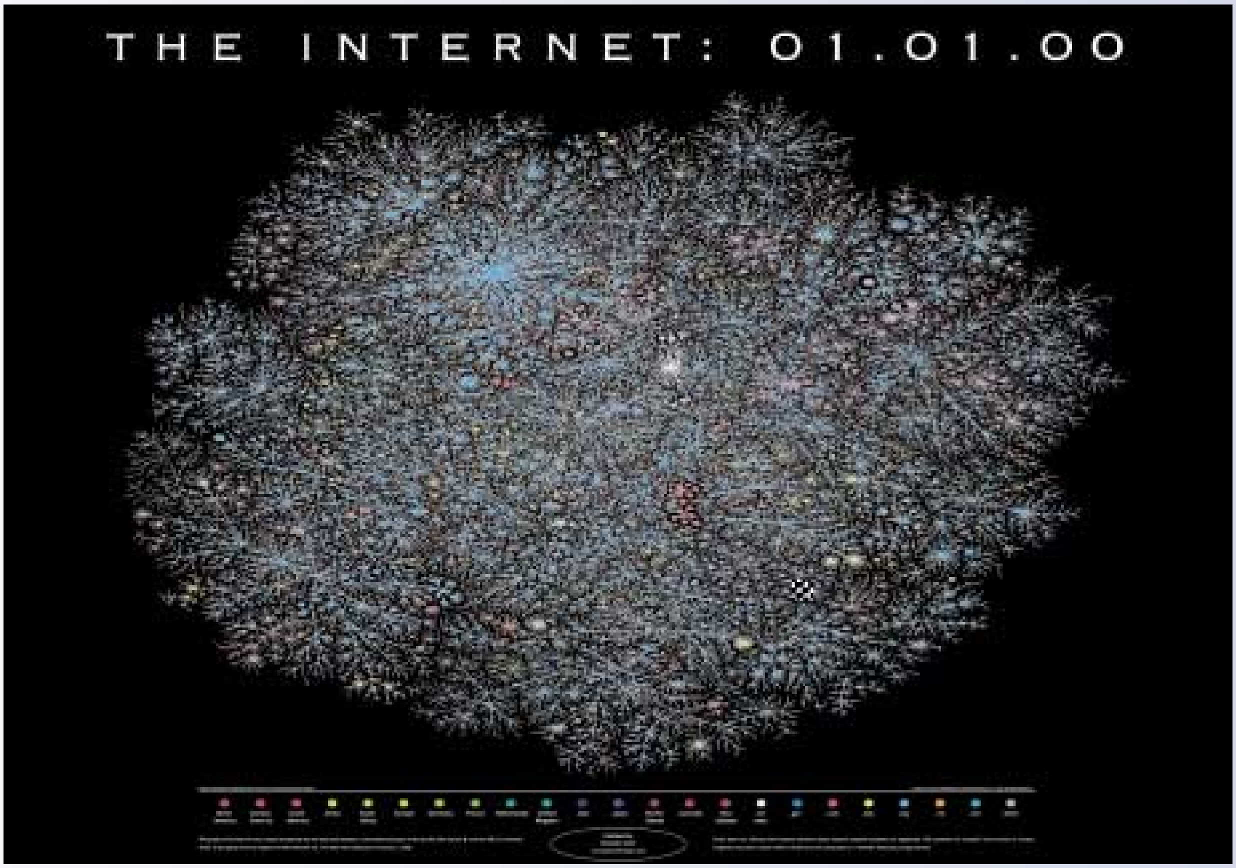

so that the reliability remains high. If you have a system with millions of lines of code, it will fail all the time. We've got to get the code down to where it's manageable, controllable, and debuggable so that we can get the reliability we really need. We must strive to build a network that is *intelligent but simple*.

Doing so requires rethinking the way we're building the Internet.

IT Pro: Do we need IPv6?

Roberts: I believe IPv6 is an operational challenge and that the conversion to the new protocol will cost billions of dollars. What we need is extended addressing, a feature easily available by adding an extended header. We can do that incrementally, so that you can add one header or two headers or three if you need to, rather than jumping suddenly from 32 to 128 bits.

Now, 128 bits is an address size so huge we don't need it yet, and it presents extremely high operation costs. Plus, the rest of the protocol changes in IP aren't needed anymore. So we can avoid all those.

IT Pro: Does optical switching provide the intelligence and reliability we need?

Roberts: Clearly, we're going to use DWDM (dense wavelength division multiplexing), with many frequencies on each fiber, and each fiber will have much more bandwidth. But when everybody says "optical switching will take over the core, the core can be dumb, you don't need anything else, and bandwidth is free," they're wrong.

Optical switching is still in its infancy from a packet-switching perspective, and, because it's circuit switching, it does a very different job than packet switching. Packet switching fills the circuit up, and circuit switching does not. It's not efficient.

We need a better protocol for routing so that we don't have to balance the network with circuit switching. Several methods have been proposed for optical switching. One

Figure 2. How Internet paths branch at each network switch or router along the way.

© Lucent Technologies and TeleGeography Inc., 2000, Peacock Maps, http://www.peacockmaps.com

rebalances the network so frequently it updates every few seconds, which would be unnecessary if we had better routing. Another focuses on reliability to handle switchover in the case of failure. This solution could be achieved by giving packet switching multiple paths in the network instead of optical protection. Optical switching, on the other hand, does network configuration well, such as implementing city bypass when a circuit is full. Tomorrow's networks will need both packet and optical switching.

IT Pro: Will we have one unified, converged network in five to seven years?

Roberts: I think we can definitely get there. The converged network will be easy to build, but the problem will be standards. If we could go in and tune the standards, we could achieve convergence quickly. Equipment at the core of today's network gets replaced every year or two, so the equipment change-out would happen quickly. We must work with standards committees to make the standards modifiable, but that's often a pretty difficult proposition given the timeframes and politics involved.

IP can support it all. When you're done, there will not necessarily be a full 40-byte header on every IP packet. We may compress the header significantly if we don't need it. Today, header length is an IP inefficiency we don't need. For voice calls, you would like to get rid of that overhead after the first setup. Voice has very short packets, which makes the header much longer than the payload.

IMPLICATIONS AND OPPORTUNITIES

IT Pro: What implications do current developments have for readers who manage enterprise networks?

Roberts: I think the same things hold true for enterprise users as for global users. Enterprise users need equipment that can handle all the different qualities of service that must be run over a single pipe.

They need to manage their network as a single composite that can be handled without a lot of extra attention. They don't need to be tuning their MPLS pipes all day long to handle a specific fraction of traffic and similar settings. Those tasks should all be relatively automated. Enterprise users should also be able to manage delay-sensitive voice and video traffic over their network, achieving the same quality as toll-quality voice traffic.

IT Pro: What about the last-mile connection from the enterprise to the Internet, or to a private wide area network?

Roberts: I think that DSL (digital subscriber line) and fiber to the office will make the last-mile problem disappear quickly. Most big buildings get good access now, and that will continue to improve dramatically. Homes already get much better access, and wideband wireless in the megabits-per-second range will be here next year to change the playing field for people even when they travel.

These rapid advancements put a bigger demand on the network to fix its problems. Many people who subscribe to cable or DSL find that they don't have much faster page access. They can get faster file transfer, but page access suffers because TCP once again causes a bottleneck—proof the network needs more intelligence.

IT Pro: If you were twenty- or thirty-something today, what would you be focusing on?

Roberts: The evolution of the network will continue dramatically over the next decade. We'll grow from about 200 bits per second per user in the US today to about 4 million bits per second per user in seven years. We'll have TV bandwidth available for every person in the US in about seven years.

Applications will enjoy tremendous opportunities for feeding people anything they want, anytime they want—content that people never thought of communicating before, beyond video and voice. Virtual reality could certainly come to life relatively easily.

In e-commerce, selected areas—such as buying parts through the network—have already proven tremendously effective. People may not realize the truly revolutionary impact on business the network is having. Yet building and running the network will require enormous effort because it must evolve to solve all the problems we've discussed, and must grow by an order of magnitude every year or so. Maintaining that growth rate will be a tremendous challenge. ■

Lawrence Roberts *is chief technology officer and chairman of Caspian Networks (http://www.caspiannetworks. com), a Silicon Valley start-up founded in 1999 to meet the challenges of building the infrastructure for the next-generation Internet. Contact him at lroberts@caspiannetworks.com.*