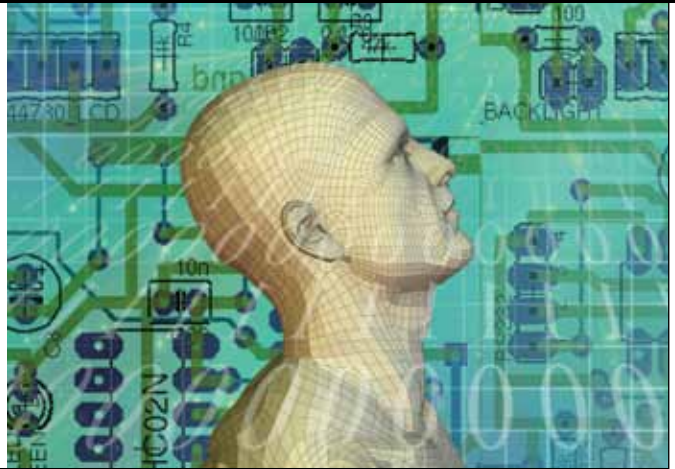# Vint Cerf: A Brief History of Packets

Charles Severance

TCP/IP evolved from 20 years' research that sought a way to move from a telephone-style circuit-switched infrastructure to a packet-switched infrastructure.

In the late 1980s and early 1990s, academics, governments, and companies around the world had built and deployed our current shared Internet infrastructure using the TCP/IP protocol. Having a well-developed open protocol implemented on a wide range of computers was an essential prerequisite for the Internet's rapid success and growth.

The simple answer to the question, "Where did TCP/IP come from?" is that its basis was ARPANET, an earlier small-scale research network. I recently spoke about this technology's emergence with Vint Cerf, who is recognized as one of the "fathers of the Internet" and was a cofounder of the Internet Society, which was established in 1992 to provide leadership in establishing Internet-related standards, education, and policy.

To view the full video of this interview, visit www.computer.org/computingconversations.

FROM THE GROUND UP

Of course, the real answer to the question of TCP/IP's origins is that it ultimately emerged from 20 years' research into a wide array of topics that explored moving from a telephone-style, circuit-switched infrastructure that required dedicated resources for every active pair of users to a packet-switched infrastructure where all active users dynamically shared all resources:

> Leonard Kleinrock specifically studied the packet-switching concept at MIT in 1961. Kleinrock focused on message switching and did a brilliant dissertation on the use of queuing theory to analyze what networks of queues would look like with a message-switching approach. Although he never used the word "packet," Kleinrock's analysis is as applicable to packet switching as it is to message switching.

One advantage of a packet-switched network is that it can dynamically route data around partial network outages as might be experienced during wartime or perhaps due to a severe large-scale weather event, such as a hurricane or typhoon. While working at RAND in 1962, Paul Baran created an extensive design for a resilient packet-switched voice network. As Cerf describes it,

> Before the existence of integrated circuits or anything else, [Paul] is saying, "We really should be digitizing and packetizing voice and using pole-mounted radios to create a highly connected environment so that if holes are knocked out by nuclear explosions, information can still get from one end to the other." You chop up the speech into little 20-millisecond pieces and dynamically route it like a "hot potato"—if you get something, you get rid of it as fast as you can. That's around 1962 and is documented in an 11-volume series called "On Distributed Communications," but he can't sell it to anybody. The traditional telcos—AT&T in particular—and the people who were in the Defense Communications Agency laughed him out of the room and say it's a silly idea that can't possibly work, so he should just go away. He never gets anywhere with the idea in spite of all his documentation.

But academics were starting to do research on how to break data into packets and send those packets over the traditional telephone network or "local area" wired networks:

> In the 1964-1965 timeframe, Donald W. Davies of the National Physical Laboratory in London also gets the packet-switching bug and tries to get money from the science research commission in England, but he only gets enough to build one node. He

Published by the IEEE Computer Society

builds a network and invents the term "packet" to describe what these objects are—and it works. He has a bunch of terminals and other things hanging off this one node so in a funny way, he builds a local-area network, but it's based on physical wires. In 1966, Larry Roberts along with Thomas Merrill does a point-to-point experiment to test packet switching. It's between the ANS FQ-7 machine at System Development Corporation in Santa Monica and the TX-2 machine at MIT Lincoln Laboratories, which is where Larry is. They demonstrate on a 2,400-bit per second modem that can move packets back and forth.

While the research into the theoretical and technical aspects of packet switching was under way, another thread of inquiry explored what people might do with a ubiquitous always-on network:

J.C.R. Licklider is a psychologist at MIT, but he's convinced in the early 1960s that computing is important to non-numeric processing: it will allow people to collaborate in ways they never could before. He starts the Information Processing Techniques Office at ARPA and encounters Douglas Engelbart at SRI International. The two bond because Engelbart's oN-Line system (NLS) is all about non-numeric computing and the ability of people to build up a superstructure of communication and documents and interact with each other. Engelbart has a World Wide Web in a box at SRI, and Licklider sends out slightly tongue-in-cheek notes to his community of people about this "intergalactic network." Licklider really gets credit for having put this meme of communication and collaboration in place at ARPA.

MAKING CONNECTIONS

In addition to the growing notions of how people could collaborate via a shared digital network, there was also the practical consideration of deploying increasing numbers of

computers and computer terminals to meet the military's information-processing needs:

Robert W. Taylor comes along to run the IPTO after Licklighter and is all hacked-off because he has three terminals in his office at the Pentagon connected to three different machines: "Why can't there be one terminal talking to all three? We need a network!" As he's pursuing this idea with Charlie Herzfeld, the head of ARPA at the time, Charlie hands him $1 million over a 20-minute conversation, and now Taylor has to figure out who will actually do this. Taylor isn't a technologist, he's a psychologist-type, so he recruits Larry Roberts from Lincoln Laboratories, who did that earlier packet experiment over the 2,400-baud modem.

> **Another thread of inquiry explored what people might do with a ubiquitous always-on network.**

With $1 million to spend on research into a shared packet-switched network infrastructure, the IPTO wrote a Request for Quotation (RFQ) to solicit research proposals. Cerf wrote one of many responses to the ARPA RFQ for network research:

I help to write one of them with my colleague Steve Crocker while we're still at UCLA as graduate students and consulting with a company called Jacobi Systems in Santa Monica, but our proposal is not funded. But the next thing we know, Len Kleinrock, who wrote the original dissertation work on packet-switched networking at MIT, has come to UCLA to teach and explore queuing theory. Len is a close compatriot of Larry Roberts because they were both at Lincoln Labs together, so Len gets the network measurement center piece of the ARPANET project.

Cerf, Crocker, and Jon Postel—all UCLA graduate students who had attended the same high school in California's San Fernando Valley—were recruited to join Kleinrock's lab to help design and build the technology that would be used in ARPANET:

I'm the principal programmer, Steve Crocker takes responsibility for managing and leading the network working group, which led to the documents describing the host-to-host protocols, and Jon Postel becomes the keeper of the documentation. He's the Request for Comments editor, the guy who becomes the numbers czar to keep track of the address spaces and allocations.

ARPANET eventually evolved to the point that, in 1972, it worked well enough for a demonstration:

The first demonstration of ARPANET happens in the basement of the Washington Hilton Hotel in October 1972. A bunch of people from the packet-switched networking community attend, not only from the US but also from France, England, Italy, Germany, and elsewhere. That group of about 25 or 30 people convenes, sees ARPANET in operation, and then forms this international network working group modeled after the working group that Steve Crocker managed. At this point, I become the chairman of that group because Steve's busy at ARPA doing artificial intelligence.

EVOLVING PROTOCOLS

At the end of 1972, Cerf graduated from UCLA and became a faculty member at Stanford. Bob Kahn moved from BB&N to ARPA and took the ARPANET project to the next level:

In 1973, Bob comes out from ARPA and says that ARPA is also working on networking capabilities beyond the original ARPANET for the military. If

we're serious about putting computers in command and control, they have to be mobile: we need mobile radio and satellite, in addition to the fixed wire systems represented by ARPANET.

Bob's brilliant idea is not to build one network with all those technologies embedded in it—instead, he breaks them apart and says let's build a packet satellite network that takes into account that it has a half-second of round-trip time. Let's build a packet-radio network that optimizes a system whose connectivity is changing with time as things move around and you get variable delay and interference.

Expanding from the original telephone-line based ARPANET to a network architecture that included many cooperating networks with points of interconnection required a new design:

We decide to build a gateway that today we call a "router" and also introduce other things like how to refer to another network. Each network thinks it's the only network in the universe. At that time, you didn't have a vocabulary that said, "Take this packet and move it to another computer on another network somewhere else that you might not even be connected to." So we have to invent an "Internet" address space to solve that problem. We have to find a way to allow packet losses in this path to be recovered. TCP now becomes a manager of reliability on an end-to-end basis instead of relying on each underlying network to be reliable. ARPANET was built on the assumption that you could build a reliable underlying network. The Internet was based on the assumption that no network was necessarily reliable and you had to do end-to-end retransmissions to recover.

This work on a "network of networks" began in earnest as a research project funded by ARPA. Cerf and his colleagues did all of their design work in the open and shared it within the academic and commercial communities:

Bob and I get the first paper written and published in IEEE Transactions on Communications in May 1974. Nobody pays much attention to it. Meanwhile, ARPA is funding us to make this actually work. At Stanford, I'm working with my graduate students, some at Xerox PARC and some at Stanford, on detailed specifications of TCP/IP. We publish it in December 1974, and it's the first time the word "Internet" shows up in print anywhere.

> **From its earliest days, ARPANET focused on connecting people, information, and technology.**

From 1973 to 1978, the research team designed and implemented four complete iterations of the Internet protocols as they found and solved new challenges. One innovation was to separate the TCP layer into the TCP and IP layers, which let real-time applications use the Internet without the error correction added by TCP:

For the next five years, we do everything we can to get TCP/IP implemented on every operating system we can find. It goes onto IBM machines, Digital machines, HP, Unix. We sent a Unix version built by BB&N out to Berkeley to the BSD release guys, and Bill Joy says, "I don't like that code," so he writes his own and puts it into BSD 4.2, the version of Unix that carries TCP/IP to the academic world. Around the same time, Sun Microsystems comes along and builds these fantastic workstations; it wants to use open protocols, so it adopts Unix, and TCP/IP comes along with it. Ethernet connects the workstations together. Sun drives the academic community to TCP/IP.

As the number of workstations and mainframe computers that could support TCP/IP and be connected to ARPANET grew and an increasingly rich set of networked applications were developed, more universities and research labs wanted to be connected to ARPANET, and its infrastructure started to groan under the network traffic load:

All of this places huge demands on the ARPANET backbone, which is only running at 50 kilobits per second, and eventually leads to the need for higher speed. NSF jumps into the fray, seeing how valuable all this is for the academic community, and concludes that it should build a network that runs even faster. It does, and the result is NSFNet.

What's amazing about the story of ARPANET is that from its earliest days, it focused on connecting people, information, and technology. It's also a story of having patience and being willing to take the time to build it right technically, even if doing it right meant starting over from time to time. A relatively small group of well-funded researchers worked closely together for more than 20 years, starting over and redesigning their systems multiple times as new use cases presented themselves. It's fortunate that by the mid-1980s, TCP/IP was a well-developed technology that provided a solid basis for the Internet revolution of the 1980s and the Web revolution of the 1990s. ⬛

Charles Severance, Computing Conversations column editor and Computer's *multimedia editor, is a clinical associate professor and teaches in the School of Information at the University of Michigan. Follow him on Twitter @drchuck or contact him at csev@umich.edu.*

cn Selected CS articles and columns are available for free at http://ComputingNow.computer.org.