

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

***The co-founder of the Internet recalls the non-commercial early days and looks at today's issues of fair use, privacy and the need for security.***

Robert E. Kahn is chairman, CEO and president of the Corporation for National Research Initiatives (CNRI), which he founded in 1986 after a 13-year term at the U.S. Defense Advanced Research Projects Agency (DARPA). Kahn conceived the idea of open-architecture networking, and is a co-inventor (with Vinton Cerf) of the TCP/IP Protocols; he was responsible for originating DARPA's Internet Program, which he led for the first three years.

**UBIQUITY:** When you were putting the TCP/IP protocol in place to set the stage for the Internet, what did you think the future would be? In other words, how much vision does a visionary have?

**KAHN:** You have to put this in the context of where we were at the time. I started my career at Bell Labs, finished a Ph.D. degree at Princeton University and then joined the faculty at MIT. I started my career as an applied mathematician, but was advised early on to get some practical experience building real things -- to focus on real problems as opposed to just mathematically interesting ones. So I took a leave of absence from MIT and went to Bolt Beranek and Newman (BBN), a small consulting firm in Cambridge, Mass that started in the architectural acoustics business. At BBN, I ended up both designing and helping to build the Arpanet, which was the first packet switched network and was based on leased *high bandwidth* land lines (50 kbps at the time).

**UBIQUITY:** When did you go to DARPA?

**KAHN:** I joined DARPA in late October of 1972. DARPA already had a computer network with some first-generation host computer protocols and a small research community that was just beginning to make use of it. I started to branch out into multiple networks, such as packet radio and packet satellite systems. The packet satellite network created an Ethernet in the sky so that we had an economic way for computer communications involving researchers in the US and various countries in Europe. Packet radio was a ground-based mobile packet network.

**UBIQUITY:** What was the first approximation like?

**KAHN:** The first two systems to be interconnected were the Arpanet and the packet radio network. Packet radio used embedded microcomputers and spread-spectrum technology in a low-cost portable configuration. Spread-spectrum had been known for many years before that, but the only *systems* employing spread-spectrum were very expensive and intended for use in the intelligence community or for special-purpose military applications. We were developing little portable radios that people could carry around with them. I called them "luggable" because they were small, approximately a cubic foot in size and weighed about forty pounds, if I recall.

**UBIQUITY:** So what other kinds communications networks did you use?

**KAHN:** We had the three different kinds described above. Specifically, we had the ground packet radio net, the packet satellite network, and we also had the Arpanet. The question that I was dealing with was how to put them all together to make a network of networks where machines on any one network could talk to machines on any other without the users having to know about their respective network, their location on the network, or how to route their data. That's what I would say was the posing of the Internet problem.

**UBIQUITY:** How did you think the problem might be solved?

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

**KAHN:** From an architectural point of view, I knew how I wanted to approach the problem, but I did not have a good idea about how to embed it into the machines that would be on the Net. That was one of the hard problems because there were several different time-shared operating systems. Vint Cerf had been through that process while helping to create the original Arpanet host protocol, which was called NCP. I was seeking to develop a new kind of protocol for use on fundamentally the same machines. I figured he was an ideal person to team up with because he had already been through that process once before.

**UBIQUITY:** Where was Vint working at the time?

**KAHN:** He'd just gotten a Ph.D. at UCLA and taken a position as an assistant professor at Stanford University, and was looking for things to work on. I approached him with the general idea of working with me on linking these three nets together. I had a view of how to go about doing the communications portion. He started thinking about how to bring his experience with NCP to bear on the problem. A few months later, we wrote the paper that described the use of TCP/IP for packet network intercommunications.

**UBIQUITY:** When was the paper published?

**KAHN:** It was published in the IEEE Transactions on Communications in May, 1974, but the paper was actually written by the two of us in the summer of 1973. We worked on the ideas partly on the East Coast and partly on the West Coast. We physically sat down to write the paper in a few days at the old Cabana Hyatt in Palo Alto. The first draft was handed out at a September 1973 meeting in Sussex, England. That's how it got started. We were working on a really interesting and challenging R&D problem, namely, "How do you link together dissimilar packet nets so that the computers on them can all interface with each other and work together?"

**UBIQUITY:** What were some of the specific issues you were dealing with?

**KAHN:** One issue was how to deal with global addressing. The existing nets viewed addresses as being local; and they all used different addressing schemes. We had to develop a global way to identify each computer as they will not generally be on the net to which the user is connected. That's where the notion of an IP address came from. We also had to deal with the issue of end-to-end reliability because we knew packets could get lost along the way. The Arpanet was designed to be reliable, so if some packets were lost, it was assumed that something in the net had broken. (This would be as if you were printing to a line printer and the line printer went down, nothing would print.) If packets didn't show up at the destination, then the idea was you just hit the equivalent of a reset button and start over again. Well, that wouldn't work in the Internet environment and so we had to develop an error-detection and retransmission scheme to get around that.

**UBIQUITY:** Other issues?

**KAHN:** Networks have different packet sizes, so a big piece that couldn't fit through a given network would be chopped into little pieces that might travel on different routes. You could also have duplicates and/or missing pieces. We had to figure out how to put all the pieces together at the other end.

**UBIQUITY:** And was reliability a great concern?

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

**KAHN:** Some networks were inherently less reliable than others. The Arpanet was intended to be very reliable but a radio net could be jammed, causing disruption and packet loss. You could drive a mobile packet radio node through a tunnel, in which case you'd like it to pick up and continue without losing a packet when you come out the other end. You could be behind a mountain and lose a signal because of the topography of the land. All of that had to be factored into the overall architecture. That's what the concept was at the time: How to make that all play together.

**UBIQUITY:** How did you go about finding solutions to these issues?

**KAHN:** The initial work that we did was to describe the method, which we documented and then built pieces of it to show how it would work. It was a technological enterprise with a lot of research issues along the way that we knew we'd have to deal with. One of the things that we had to figure out was how to build something we called a gateway at the time; it was a device that could sit between networks, interface to them, interpret the incoming IP addresses and figure out where next to route the packets. Today, this kind of device is called a router.

**UBIQUITY:** How was the research funded?

**KAHN:** I was at DARPA at the time, and DARPA funded three different groups to work on this effort. Stanford University built an implementation of TCP/IP and led the implementation effort with the other two groups in demonstrating interoperability. University College in London created an independent implementation for a different machine. This was significant because we were specifically thinking of doing this internationally right from the get-go. BBN was the third participant to develop an independent implementation for yet another different machine. In addition, we also funded BBN to build the gateway component.

**UBIQUITY:** How were your ideas generally received at the time by the technical community? With great skepticism? With great enthusiasm?

**KAHN:** Vint was clearly enthused when he and I first talked about it. Bob Metcalfe, who at the time was working on the Ethernet at Xerox Parc, was also very interested; Bob even volunteered to become a collaborator if we wanted. He was working on a specific network at the time, while our main interest was in the logical architecture for interconnecting different kinds of networks. The Ethernet eventually became the most widely used local area network on the Internet. A lot of people were interested, but at the time I don't think anybody saw it in the light of what eventually happened.

**UBIQUITY:** Was there anything that shocked people? For example, was anybody scandalized by the idea that there would be no central network administrator or control?

**KAHN:** I don't really think so. It was all a straightforward experimental activity at the time. In fact, an awful lot of the early networking activities just didn't capture the fancy of the media or others in the commercial world. People often ask, "When you installed the first node on the Arpanet, who was there from the press?" and my answer always is: "You've got to be kidding." No one was thinking press coverage at the time. The kinds of questions we were asking were whether there would be anyone on the loading dock waiting for the truck to arrive?

**UBIQUITY:** What about the military? Was there a military view about the network having no central control?

**KAHN:** Again, no, not at that time, because when we started working on networking the military had

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

very little notion of what packet-switching was. They were mainly involved in circuit-switching. It was not until a number of years later that the military realized the value of what we were doing for future defense applications. In 1975, the Defense Communication Agency (DCA) agreed to take over the day-to-day responsibility for managing the Arpanet. I think that by 1978 they wanted to use it for their own purposes. If you really want to know what happened, when DCA finally decided that packet switching was something that was a good choice for Defense, the up-front assumption was that researchers wouldn't know the real military needs and requirements. So they put together a procurement for a system called AUTODIN II. This called for a packet-switched network especially designed for use in the Defense department.

**UBIQUITY:** What was the story behind that?

**KAHN:** There had been a system called Autodin I that I believe Western Union built many years ago. It was a slow system. It originally involved the use of paper tape. Delivery time of messages was measured in hours or even days. Autodin II was an attempt to upgrade this capability and make it work a lot faster. So they put out an RFP and various proposals were submitted. I was asked to chair a blue-ribbon panel to assess the conclusions of the proposal review group that was making its recommendations to the Director of DCA. I remember asking if he wanted the panel to comment on whether Autodin II was a good idea or not, or did he just want the panel to identify issues with the specific proposals that were evaluated? He said, in effect, "don't address the issue of whether the RFP should have been issued in the first place! That's not on the table. I just want this panel to tell me the pros and cons of the proposals evaluated -- and the selected one, in particular." So they went through with the procurement and undoubtedly spent \$50M or \$100M to build it. Along the way, DCA must have concluded that the Arpanet technology could actually serve their needs, since one day after Autodin II was delivered, the system was canceled and never used operationally.

**UBIQUITY:** That's hilarious. What reason did they give for canceling it?

**KAHN:** I believe the reason was that it cost much more to run than would be the case with the Arpanet technology and it had lower performance. And it was probably more vulnerable, in fact. They could have saved themselves quite a bit of money and time if they had gone directly to use of existing Arpanet technology. Unfortunately, it took time for them to get up to speed. They naturally thought conventional procurement practices could do a better job for them, but that didn't turn out to be the case. Eventually, they went back to the Arpanet technology.

**UBIQUITY:** It sounds as though, all in all, your plans for creating the Internet went pretty smoothly from the time you wrote that famous paper till now.

**KAHN:** Well, only because DARPA constantly funded both the R&D and the implementation, all in parallel. We were building something experimental that turned out to actually be useable.

**UBIQUITY:** How rare is that?

**KAHN:** Well, I think it's pretty rare and still not very well understood that something the research community might do could actually turn into real infrastructure. Of course, this is a situation where, over time, pieces of the Internet were changed out and they became a little more robust and a little more commercial. Eventually commercial suppliers were providing their own networks and technology that became part of the Internet. Today, the Internet is almost entirely commercial.

**UBIQUITY:** In many projects, the technical people end up fairly bitter about the management people.

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

Was there any of that kind of problem?

**KAHN:** I don't think so. To the extent that there were management people, up until the late 1980s, most were probably in the US Government. The Government had a substantial (dominant, initially) role in the creation of the Internet. This role subsided greatly after NSF stopped their support for the NSFNET around 1995.

**UBIQUITY:** Did you consider yourself to be management or technical?

**KAHN:** I was both for a long time. At DARPA, I was certainly playing both roles and that continued for many years thereafter. I helped to start the Internet Society back in the early 1990s. That was as primarily an administrative, organizational, and managerial entity, which functioned as a professional society.

**UBIQUITY:** It almost seems hard to believe that the project developed so smoothly; typically people are pulling out their hair saying, you know, "We didn't get money for this, we didn't get permission for that," and so forth.

**KAHN:** Well, remember, the support for the Internet came from the Information Processing Techniques Office at DARPA. In addition to starting the Internet program there in the early 1970s, I was also the office director from 1979-1985. I certainly wasn't questioning the project because it had been my idea to create what I was calling the internetting effort in the first place. In addition, I could see the growing value of it as the project went along.

**UBIQUITY:** Was there any commercial interest in it at the time?

**KAHN:** I think we were fortunate in that there was absolutely no commercial Internet industry out there. There were no Ciscos or Internet service providers; there was no commercial anything. So nobody out there saw the original internetting initiative as a threat to their business. AT&T provided a lot of our facilities, but they made a conscious decision to stay out of the DARPA computer networking efforts. I suppose they thought they could make more money by selling the equivalent of picks and shovels, namely the underlying circuits. Besides which, where could the Internet have gone back then? In the 1970s, which is when most of the basic research and development was taking place, the host computers were all big (and expensive) time-sharing systems. Only a small number of organizations were using them. The PC had not yet been invented and there were no LANs to connect them. The Internet just wasn't seen as a big thing at the time.

**UBIQUITY:** And so how important was the PC industry to the commercial success of the Internet?

**KAHN:** If the PC industry hadn't started up, I believe you never would have had the proliferation of computers and millions of people on the Net. Only the companies that could afford big time-sharing systems could have justified it.

**UBIQUITY:** Do you remember what was the first commercial application on the Internet?

**KAHN:** Yes, CNRI carried it out in the 1988-89 time frame. It was the connection of the MCI mail system to the Internet. A few years later when MCI needed to talk to others in the research community, because they were building the NSFNET, they found they could use their MCI mail system to talk to the community and thought it was just a natural part of the MCI mail system. But it wasn't. It was something we had built separately to demonstrate the utility. But they picked it up and started using it

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

and are still using it to this day. Vint Cerf was the principal investigator on that effort for CNRI.

**UBIQUITY:** Speaking of Vint, when and how did your career and Vint's diverge?

**KAHN:** As I mentioned, I had previously interacted with him at UCLA on the Arpanet, and I later approached him to work with me on the Internet project after he moved to Stanford. In 1976, he joined DARPA and took over running the Internet project from me. So we overlapped at DARPA from 1976 through roughly September or October of 1982 when he left to join MCI. I ran the Internet program at DARPA from 1973 until he joined DARPA in 1976. He ran it through the end of the summer of 1982. Then I picked it up again after that and ran it for a year while working with Barry Leiner to take over the responsibility. Vint left DARPA in 1982 to take a job at MCI as the key engineer on building the MCI mail system, which is why we picked that as an ideal first application after he later joined CNRI.

**UBIQUITY:** Was getting it approved by the government easy or hard?

**KAHN:** Getting it approved as an application was a political hurdle, because the government had never approved a commercial Internet application before. We were still dealing with government network facilities. As soon as the application was staffed and approved, AT&T and Sprint had their own systems up very quickly thereafter. However, the CNRI application set the standard for such activities and opened the floodgates for eventual large-scale use of the Internet.

**UBIQUITY:** What was the next step in the growth of the Internet?

**KAHN:** The next step involved passing of the so-called Boucher bill in early 1993, which allowed NSF to open the NSFNet to commercial use. The NSFNet, which was created in the late 1980s, eventually replaced the Arpanet as the backbone of the Internet. The Arpanet was decommissioned in 1990. It was deemed not needed by then because NSF had built a one-and-a-half megabit backbone that could take its place. The Arpanet had become a low-speed net compared to the NSFNet. NSF, under an agreement with a Michigan non-profit called MERIT, ran the NSFNet until approximately 1995 when they decommissioned it. At that point in time, there were so many commercial networks that they felt the NSFNet wasn't needed anymore.

**UBIQUITY:** When did you leave DARPA?

**KAHN:** I left DARPA in late September of 1985 and started the organization that I now run, CNRI, in January of 1986. Vint Cerf was employee Number 2 at CNRI; he was corporate vice president at CNRI for about eight years until the beginning of 1994 at which time he was recruited back to MCI. Of course, once back at MCI, he got tied up in the whole Internet wave, which was breaking wide open at the time.

**UBIQUITY:** Do you keep in touch with him?

**KAHN:** Yes indeed, I see him fairly often. For one thing, we get many joint honors: it's very hard for anybody with a history of collaboration such as ours to be able to tease apart who did what. So we are often honored together.

**UBIQUITY:** Let's turn to CNRI. Why did you decide to start that organization?

**KAHN:** The mid-80s were the days when the Gramm-Rudman bill was being considered to limit government spending. There were all kinds of budget cuts that I could anticipate. I thought staying in

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

the government wouldn't help me to advance my notion of a national information infrastructure, and instead I'd be in the mode of just trying to protect what budget we already had. So I thought it was a good time to start this organization. Very few people knew what the Internet was at that point except perhaps a small set of researchers. I thought CNRI could be helpful in nurturing the Internet, helping it to grow and to apply the Internet in related areas with the business community, the banking community, on building knowledge banks, digital libraries and a whole variety of similar infrastructure ideas.

**UBIQUITY:** What was on CNRI's initial agenda?

**KAHN:** One of the early areas that we worked on was digital libraries. We produced an interesting report on that subject entitled "The Digital Library Project; Volume 1: The World of Knowbots (DRAFT)" and later began implementing the ideas in that report. We also explored the possibility of developing a "Knowledge Bank" for science and technology, but there wasn't a sufficiently developed technology base at that time. DARPA decided to fund basic research in that area and we moved along to other more accessible areas such as high-speed networking. The first government support for CNRI came from the National Science Foundation in 1988 and that was to carry out a research initiative to demonstrate the architecture and utility of very high-speed networks that operate about a billion bits a second or more, and which were a major leap from where we were before. This initiative was big news at the time; it made the front page of the *New York Times* in June of 1990. The Sunday business section had a front-page article about what we were doing with networking in September of that same year. It really put high-speed networking on the map.

**UBIQUITY:** What then?

**KAHN:** We helped create five gigabit testbeds linking participating researchers at more than a dozen research laboratories across the country with contributions from many of the carriers including MCI, AT&T, GTE, NYNEX, Bell Atlantic, Pacific Bell, and Bellsouth plus organizations such as IBM and Bellcore. This initiative supported the first operational experiments with optical amplifiers, the first deployed ATM switch operating at OC-12 speeds in a carrier central office, and the testing of various advanced transmission, switching and computer technologies. We funded a group at the University of Illinois to build a point-and-click browser for accessing their supercomputer simulations. This became a starting point for the creation of the Mosaic browser for the Web. We have continued to do a lot of things along those lines.

**UBIQUITY:** Give us an example.

**KAHN:** We made a major effort in the area of digital objects and identifiers for digital objects. For many years, the system we developed as part of that project has been the predominant system used by the book publishers for identifying content on the Internet. A publisher, for example, can move a book from one place to another and that book can then move around the Internet and even be in many different places at once, yet one identifier will always get you to it. We call the system for resolving identifiers the "Handle System". It's really an interesting system with a lot of commercial spin-off based around it, especially in the publishing community, which uses handles but calls them Digital Object Identifiers or DOIs.

**UBIQUITY:** What are your views on the current status of copyright in the digital environment?

**KAHN:** We have been promoting copyright and, more generally, intellectual property protection in the network probably as much as anybody in the research world. CNRI built a system for the U.S.

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

Copyright Office to manage the registration of copyright claims and the attendant submission of copyright information and digital objects online; the system is called CORDS (cords.loc.gov). In my view, one of the problems that has not been satisfactorily dealt with in this country is the widespread lack of respect for the value of intellectual property. People think that they can do anything they want with intellectual property just because they themselves don't happen to see any cost associated with accessing it on the net and, perhaps sending it to others or otherwise using it. I think this is clearly an educational issue as much as it is a constitutional issue.

**UBIQUITY:** Most of the publicity seems to go to people who are worried that there will be some constriction in the fair use principle.

**KAHN:** You need to realize that fair use is not something that's in the law; it's something that's been determined in the courts. Fair use emerged from different lawsuits over time, where somebody thought allowing certain limited kinds of use of copyrighted material without the copyright holder's permission was just a reasonable thing to do -- something that should be allowable in certain educational contexts, for example.

**UBIQUITY:** Another controversial area is the issue of surveillance, such as in the recent controversy over a national Total Information Awareness system that's been in the news lately. What are your feelings about that?

**KAHN:** There are two parts. Clearly there's tension that always exists between the need of law enforcement to do its job and the need of individuals for privacy and the need of everybody for security. The tension was shown during the 1990s with the Clinton Administration when they were trying to get their "clipper chip" approach adopted. People worried about Big Brother watching you and the like. This issue has been around for a long time. But we have methods in this country for dealing with some these issues -- court procedures that deal with wiretaps, for instance. I think after 9/11 the world changed for most of us in pretty fundamental ways. We suddenly realized that we weren't nearly as safe as we thought we were and that the potential attacks on this country could be really serious.

**UBIQUITY:** Then what's next?

**KAHN:** The jury is still out as to what the real threat is, what we need to do about it, and therefore how fast we need to act. If you compare this situation to decisions about nuclear weapons: It's one thing to do the research on how to build a nuclear weapon and a quite different thing to actually deploy a nuclear weapon capability and still another to actually use it. I don't think we can ever feel safe by keeping our collective heads in the sand.

**UBIQUITY:** Do you feel the government is calibrating its security measures proportional to the real threat?

**KAHN:** That's probably what they're doing. Of course, I understand that skepticism often goes along with government statements, with people hearing a government assertion and thinking, "Yeah, well, that's what they're saying, but what are they really doing?"

**UBIQUITY:** What's your own opinion?

**KAHN:** My personal opinion is that, yes, they're really doing what they say they're doing. However, having said that, the fact is that once you have a technology like a Total Information Awareness

## Putting it all Together with Robert Kahn (Ubiquity interview with Robert Kahn)

system you're a lot more likely to use it just because it exists. So it seems to me the ultimate questions are would we ever deploy such a system and use it, what rules would apply to it, who would decide and what oversight would there be.

**UBIQUITY:** How should these issues be addressed?

**KAHN:** I think we have an obligation in this country to understand what our technological options are. If we can figure out how to do something technically then, presumably, so can somebody else, so the only way that we'll know how to counter a technology is to understand it well enough to be able to do that. As I indicated earlier, I don't think you can put your head in the sand and just say, "Somebody might be able to do that, but it's really very difficult to do, so we're not going even look at it." I think national security dictates that you're obliged to look at every one of those options and possibilities to determine what's viable and what is not, and to be prepared to react if one ever needed to. As to what the actual policy should be, I think you need to draw on experienced and knowledgeable individuals. I don't know what I would do if somebody said, "Okay, Bob. You decide." I think it's one of those problems that is so thorny and complex that you have to have the problem and access to all the attendant intelligence information to be forced to actually think about it in sufficient detail.

END

[This is the first of two Ubiquity interviews with Bob Kahn. In the next interview, intended for publication this Spring, he will talk to us about his digital object architecture and various other topics related to his current work and interests.]

Ubiquity, Vol. 4, Issue 3, March 11 - 17, 2003