Oral History of Vinton (Vint) Cerf

Interviewed by:
Donald Nielson

Recorded: November 7, 2007
Mountain View, California

CHM Reference number: X4308.2008

This interview took place at the Computer History Museum in Mountain View CA on November 7, 2007. Vinton Cerf, a Fellow of the Museum, and currently Google's chief Internet promoter, is interviewed by Don Nielson, a retired Vice President of SRI International and a colleague of Vint's from the early days of internetworking development.

Donald Nielson: Vint.

Vinton Cerf: Hi.

Nielson: Welcome. You've been through such interviews a number of times. I've read them, and have struggled to figure out how not to go just do another replay of your life. So I'm trying to be a little bit different here, and I may or may not succeed. But with your help, we'll see if it works.

Cerf: Different is cool, let's try it.

Nielson: I've decided to divide the interview into three parts: Vint the person; Vint the practitioner, what have you done and accomplished; and then there's Vint the philosopher, in the sense of an idea source. So we'll break it into those parts, and we're going to start with the middle one, the practitioner.

Cerf: Okay, all right.

Nielson: Okay, question one goes back to the 1960s. Was there a recognizable point in time when the creation of digital networks would first capture your interest? What influenced you toward that subject and who was it that influenced you?

Cerf: So that's really interesting. If we go to the 1960s, the networking part shows up around 1968 when ARPA puts out this RFQ for its packet switch ARPANET, before that I really didn't know too much about networks at all, and I visited the Sage computer system in 1958 thanks to my father's friendship with a programmer at Systems Development Corporation who wrote simulations of the radar tracks that you would see if the Russian bombers were coming over the Pole. And that used the 2,400 bit per second communication channel from the distant early warning radar, so I had this very vague understanding of that part at that time at 15 years old. Steve Crocker captured my interest in computing though when he got permission to use the Bendix G-15 paper tape fed machine at UCLA, so then I went to Stanford and studied computing and mathematics and went to work for IBM. So by the time 1967 rolled around, I felt I needed to go back to school for a PhD and enrolled at UCLA. Steve Crocker introduced me to Jerry Estrin and to Len Kleinrock. So Len got this contract from ARPA to do the network measurement center and I just got sucked into that because they needed a programmer to write software to capture data from this planned packet switch network ARPANET. So this wasn't me or any of my ideas that drove me into this, it was kind of circumstances, but it was instantly fascinating, and I think what attracted me more than anything was the idea that you could do something here and it would cause an effect 3,000 miles away. And the idea that a little piece, a program could do that, it wasn't something that you manually pushed or did anything. It was the fact that you wrote a piece of software here and it was running and it was interacting with another piece of software someplace else that just grabbed me as something fascinating and interesting and can imagine that there could be hundreds or thousands of machines all interacting with each other. That infected me with an excitement about networking that I still have to this day.

Nielson: But when you were at IBM you worked a little bit on a timeshare machine, didn't you?

Cerf: Yes.

Nielson: And didn't that trigger a notion of remote computer access?

Cerf: It didn't have the same reaction, nor did the distant early warning system, and the reason for that as I look back now is that in the case of the time sharing machine, there was one computer and a lot of people with terminals, but they were interacting with that one machine. The distant early warning radar is the same thing. There was data sources coming from a long ways away but there was one machine, the Sage computer that was analyzing the stuff. What intrigued me about the ARPANET was that there were computer programs interacting with each other, not just terminals interacting with remote programs but the possibility of programs interacting, and you could begin to imagine these virtual environments, not the 3-D games-based environments, but these virtual communications environments that would allow computer programs to exchange information and for some reason the idea that these programs had a life of their own and they could interact with each other that I found very intriguing.

Nielson: But it seemed to me like one of the very first reasons for the ARPANET was resource sharing. Wasn't it?

Cerf: Yes, absolutely correct. That was what Larry Roberts insisted.

Nielson: Which was quite like access to a time-share host specializing in some special application. I thought only later did the processes in one machine start interacting with the processes in another, other than something like TELNET.

Cerf: Well, but actually so that's a rather pragmatic perspective, but I have to tell you when Steve Crocker was leading the design of the network control program. NCP, the predecessor to TCP/IP, we struggled not with terminals communicating with things but with hosts communicating with each other. The terminal programs were applications that had to run on top of that, and so it was the host to host protocols in which every machine was treated as equal that forced us to think about program to program interaction.

Nielson: What was your role in NCP?

Cerf: Oh, I was just one of the guys that helped fashion the protocol. Steve Crocker led that project and Jon Postel and I and Braden and others participated and were part of the debates and he and I and there was another guy named Steve but I've lost his last name now who wrote, three of us wrote a paper for the spring 1970 Spring Joint Computer Conference that described what the NCP program was about. So I was just, you know, part of the team.

Nielson: Part of the team. Do you remember when you joined the Network Working Group or did you?

Cerf: It didn't exist until Steve Crocker started it..

Nielson: I understand that, but when it did start, you eventually became a member, didn't you?

Cerf: Yes. Instantaneously, I mean.

Nielson: I remember the first meeting was at SRI. You weren't at the meeting but on the other hand I'm sure it happened pretty soon thereafter. Anyway, I wanted to know what your impressions were of that particular group and what it was about to embark upon.

Cerf: Well, let's see, I do remember a 1968 meeting. I don't know that I, you're right, I don't think I was there, but I remember that there were notes taken of that meeting, and there was a lot of speculation about what would happen when this network was built, what information would be exchanged. I'm trying to remember the names of some of the people who were here at SRI who participated in that meeting. There's somebody named Eric that's coming to mind, and I didn't do my homework for this interview.

Nielson: I can tell you later on, but I can't pull it up either.

Cerf: Yeah, well anyway, I think that our speculations about this were way off the mark as to what would actually happen. I remember a poll that was sent out asking how much capacity did you think you would need in bit rate, and people took all the terminals that were in the building and multiplied, you know, by 110 baud or something to give an estimate, and it was all known [ph?] so the network working group as Steve led it was very focused on implementing things, doing the design, getting prototype software built, trying it out, figuring out what went wrong. The hard part I think was getting the stuff to work on multiple operating systems, so in some sense there was an instantaneous need for a group of people with different perspectives from different operating system points of view. There were more operating systems then than there are now in some sense. Today it's, you know, UNIX derivatives and a few others, and back then it was 20 or 30 or whatever numbers of operating systems depending on whose computer line you were using.

Nielson: I believe there was this notion at that time that they were just a bunch of kids and were waiting for the pros to show up. A nervousness I guess that led in part to the notion of Requests for Comments (RFCs). Something that sounded rather timid…. Well, let's send this out and see what happens. Did that atmosphere, the uncertainty of it all, come to you as you joined that group?

Cerf: Well, Steve Crocker describes this very well. In fact, I hope he's been interviewed or is interviewed. He was hesitant to appear in any way aggressive about anything because he didn't think he was really in charge, but nobody else showed up. Larry Roberts basically just handed this task to Steve, so the RFC idea was one of these kind of a meek, you know, here's the idea we have, does anybody have any comments on it, and RFC 1 came out in like April of 1969. I think we were all simultaneously intrigued and excited about going to territory that nobody had been in before, and certainly we had never been in before. That was exciting. I mean the idea that you could go explore territory that didn't exist before was another one of the reasons this networking idea was so exciting. And basically the senior people, the principle investigators, left the graduate students alone to just go try stuff out, and the freedom to do that and the degree of trust that they implicitly handed these graduate students, looking back on it was pretty incredible.

Nielson: While we're on that subject, why don't you comment a minute about that freedom and why you think it was important to how digital networking unfolded?

Cerf: The freedom to try things out is so fundamental to discovery, and we were discovering, I mean we were inventing and we were discovering, and to have this dictated top down means that the ideas that would be dictated were ideas that already somebody knew, whereas we had the freedom to discover things that nobody knew. And of course we made mistakes and things didn't work and that's how you learn, but the fact is that somebody trusted you enough to let you do that was thrilling.

Nielson: I think it was, In fact, a blessing that the old pros didn't show up.

Cerf: Well, I certainly remind myself of that now being in the, I won't call myself a pro but old for sure, and I know better, you know, working at Google that there are these kids who are out of college, they have a certain background and experience. They actually don't have to unlearn anything.

Nielson: The old pros were essentially circuit-switch oriented in those days.

Cerf: Yes, although to be fair about it, they weren't the circuit switch people that drove this project. The people that drove the project were people who thought circuit switching was a mistake for computer communications. They were believers.

Nielson: That's exactly the point.

Cerf: And, you know, CLINOX [Ph?] results demonstrated how powerful packet switching could be and later, I didn't know about it at the time, but Paul Baran's work and the work of Donald Davies reinforced that whole notion that this was a much better medium, mechanism for a computer communication.

Nielson: This maybe hard to answer in retrospect, but I'll ask it anyway. When you young researchers were deliberating in these early sessions, trying to figure out where you would go regarding digital networking, did any really profound concepts arise? Let’s take the example of complexity control. Did that ever surface as something you really needed to address in principle? Were concepts such as layering part of the design?

Cerf: It really did, and I'm glad that you brought that up, because part of the philosophy behind this design, both in the ARPANET, NCP, host to host protocol and then later the internet was to keep things as simple possible, and there were two mechanisms for doing that. One we sort of implicitly agreed that if we were going to try to make a particular thing happen, we should pick one way to do that, so that there was never any ambiguity, the software didn't have to negotiate to figure out which thing it was that you wanted to do. The second thing was this layering notion, although today I have debates with people who think that layering is a mistake, I look back on that and believe that it's fundamentally critical to the success of the system. By defining well defined interfaces between these various layers and having the layers take on some fairly basic functionality, it was possible to change the way each interior layer was implemented, and so this led me to eventually to having a t-shirt that said IP on everything, because the whole idea was to try to put the internet protocol on any underlying transmission system that was invented, trying to future-proof the architecture of the internet. And I learned that from the ARPANET experience and from the layering structure that was deliberately put in. So there was very conscious effort to limit complexity and to introduce simplifying techniques that would allow people to be creative within certain frameworks.

Nielson: There was also this so-called M times N problem; the question of how all these different terminals and different hosts would interact. Out of that realization came, I believe, the rather profound notion of the network virtual terminal.

Cerf: Exactly.

Nielson: Can you speak about that and where that came from to your knowledge?

Cerf: I certainly won't and can't claim any special credit for that idea. Others saw the value in creating a virtual common standard that everyone would interact with, and once again we're back to the N versus N squared way of translating things back and forth, so if you ejected information, emitted information out of the machine in network virtual terminal form and you received it in that form and translated back and forth, that simplified your job enormously, and so it's another manifestation of trying to keep things as simple as possible.

Nielson: So time moves on and ARPA introduces a Request for Proposal (RFP) for a switch, an Interface Message Processor (IMP). I've read someplace in the distant past that you actually bid on that RFP. Is that true?

Cerf: Yeah. Yes it is. Steve Crocker and I were consultants at a company in Santa Monica and we got, I'm not even sure I remember the name of the company.

Nielson: Jacobi.

Cerf: Well, it was Jacobi Systems, Jacobi Systems, that's right. So I wrote a simulation in GPSS of how the network might behave, and the other guys wrote down what they would do to build this system, and as it turns out, we didn't win the bid. On the other hand, since Steve and I happened to be at UCLA at the time, by great good fortune, Kleinrock got the network measurement center, we got sucked into it anyway and I look back on that thinking "Boy, how lucky can you get."

Nielson: That was a boon, because otherwise Vint would have started off in hardware production versus getting the insights that only the network measurement center at that time would've provided you.

Cerf: That's a fair point.

Nielson: That's super. There is this question about computer networks. Pardon me it is a little aside, but you remember the simple ALOHA protocol. It was kind of a disaster in terms of throughput or delays whenever the channel got reasonably occupied. I was worried in the early days that that was going to be a big problem, because as users came on, the capacity would limit or the delay would go infinite. And yet over time, on the Internet, somehow the transmission channels have gotten big enough to accommodate all this growth. Wasn't that original capacity limitation on multiple access systems a problem from the beginning given that we started with 50 KB channels?

Cerf: Yeah, or well on the backbone, right, literally 50, and then on the dial-up, you know, it was nothing.

Nielson: Even as little as 2400 baud.

Cerf: Less than that. I remember running 300 baud terminals. In fact, I even had a teletype, a Model 28 teletype at home around 1968 when I was working with Jacobi Systems. Steve Crocker and I had this time sharing system that we had written and we had these old teletypes. They would go at whatever it is, less than 110 baud, 50 baud or something, <makes machine noise> and I thought that was pretty cool at the time. So yes, the speeds have gone up dramatically over the last three decades.

Nielson: And is that increased capacity just something that the technology just provided and enabled rapid Internet growth? I remember when packet speech was first introduced; we had to severely compress the speech just to get it over the network. I guess it was a boon or a blessing that the capacities grew as fast as they did.

Cerf: You know, I've often wondered about that too. It's sort of like the speed of processors and the amount of digital memory you get, people keep looking for ways of getting more in the same amount of space or for the same amount of money, and I think we have benefitted from that. I don't know that I consciously assumed that things would get faster and faster. It wasn't part of my mindset. But what was pretty clear is that every time it looked like we were going to hit a limit, somebody would come up with a way to get more capacity out of the system, and that's still going on today. So it wasn't part of the conscious planning effort, but we certainly benefitted from it.

Nielson: Let's switch to the term internet and internetworking for a second. When did "internet" first become a notion? After all the ARPANET was not an internet, but an intranet. When did the notion of disparate networks somehow come to your attention?

Cerf: Okay, so first point is that certainly nobody I knew ever thought of ARPANET as intranet, that was a term that got introduced many, many years later, it's a marketing term and I hate it anyway. All right, so much for that. The idea of interconnecting networks was Bob Kahn's idea. He had written maybe just to himself some ideas in late 1972 when he left BB&N and came to ARPA, and the ideas he was writing down had to do with an open architecture, and I'm pretty sure that he may have been driven to this idea by looking at the packet satellite network, which Larry Roberts, the project Larry Roberts had started, and the plan at the time I think was that the satellite component would be embedded inside the ARPANET, and in fact there was a satellite link, as you know, going to Norway, and it just looked like another link in the net, it just had a higher latency than the others did. And what Bob was concerned about is that if we proceeded down that kind of path and every modality would be embedded into this one single network, the ARPANET, and that one organization, BB&N would be responsible for all what any innovations and everything else associated with it. And not to try to put words into Bob's mouth, but I think he was concerned that we didn't have an open standard kind of structure that allowed other people to build networks on their own independently and somehow interconnect them. So that was what he came to talk to me about at Stanford. And so we thought about it as inter-networking of many different kinds of networks, and he pushed hard to break out the packet satellite project from the interior of the ARPANET and for the same rationale, the packet radio net was segregated into a separate process, separate network, and then the question was well, how should they interact with each other, and that was the problem that we worked on. We thought of it as internetworking for quite a long time, although the term internet arose very, very early in the whole process. My recollection now is that in the December 1974 specification of TCP, that there is reference to internet in the document, and I'm pretty sure that that's the earliest reference to that particular term.

Nielson: I thought its use might have started by the time Bob came out here in 1973 or so. Packet radio was just getting underway at that time.

Cerf: Well, the project started by that time.

Nielson: It had started. I thought there were reasons those separate types of networks existed because of the functionality they were to offer to users.

Cerf: That's right.

Nielson: Before 1973In there was just a wire network, and there was no easy way to place a wireless network in the middle of that. There were just too many existing constraints. A totally different environments or technologies were needed to serve a broad community of users like the military. That suggested the need for separate networks. Is that true or not?

Cerf: Maybe not true. Here's a place where Bob would be a better person to respond. My understanding is that he had this notion of open networking, architecture, and saw that as an important choice for design, but I think also these different networks were motivated in part by the fact that if you were going to do command and control using computers, that the computers had to be available in mobile operation, you couldn't put wires connecting the tanks and things together, and similarly on board ships at sea, and again pulling cables behind the ships didn't make any sense. So that requirement that the capability be present in things that were in motion drove a radio-based design, but you could have done it by integrating the satellite and the mobile radio system as if it were just part of the ARPANET. The problem is it would've made it really hard, because now you're mixing together the tactics that make sense in the wired world and the tactics that make sense in the mobile world, and those are not necessarily at all similar.

Nielson: And you may penalize other components of that integrated network just because of the difficulties of certain other parts.

Cerf: So this gets back to simplicity again, because if you want to use this mobility idea for the ships and the land mobile, segregating the networks made it easier to design, build, test and implement, but then came this problem, okay, now that I've got these different networks, how do I make them inter-work, and one of the ground rules was you couldn't change the networks, to know that they were part of the big internet, and so that led to a number of specific challenges that had to be overcome.

Nielson: So there came a time when you needed to make these disparate networks work together, and I know TCP was well underway at that particular point. I know in 1976 we used it to bridge the Packet Radio Network and the ARPANET. But in 1977 we're going to integrate all three networks, including the satellite net, and who was the customer for that particular demonstration? Was there one or was it just an incremental unfolding, the next obvious step in the development of the technology?

Cerf: Well, let's see. This was 1977 and by this time I'd been at ARPA for just about a year and a half, a little less than a year and a half, and really eager to demonstrate to particularly George Heilmeier, who was the director of ARPA at the time that we had a technology that actually worked in the way that we said it would, namely that these multiple networks would appear to operate as if they were a single uniform internet. So from my point of view, this demonstration was to prove to ourselves and to Heilmeier that we actually had a technology that looked like it would work. We'd already done as you mentioned a few, you know, interconnections between the packet radio net and the ARPANET and my recollection is we might even have done that as early as 1975, but I don't know whether your recollection of that is?

Nielson: I believe the spring of 1976 was the first time we at SRI started using TCP, celebrating a solid gateway to the ARPANET in August of that year. .

Cerf: Seventy-six is what you remember? Okay.

Nielson: Early, very early in 19'76.

Cerf: Okay, because we had the station, which was acting as also kind of a gateway to the ARPANET, so we had done sort of pairwise kinds of demonstrations. The thing that bothered me about a pairwise demonstration is that you could almost always figure out how to do something funny that knew about both networks to translate from one thing into another, and it wasn't convincing that you had an architecture that would work, regardless of what the network was. And looming in the background of course is the ethernet that was designed in 1973. Park Lab is a mile and a half from Stanford University. I had students who were working at Park, Park had people who were taking classes at Stanford. So there was this exchange of information, although the Park universal packet stuff was considered proprietary and so they didn't talk too much about it, but they sort of alluded to experiences that they were having while they were trying these things out. And at some point if I remember correctly, we did put a packet radio at Xerox Park and they connected their PUP [ph?] system, their ethernet system to the rest of the internet through that packet radio.

Nielson: You just triggered an idea in my head, that there's a certain generalityt you associate with that time; one that went far beyond even those three networks. In other words, was there something in those early connections that portends the the way the internet has evolved; that is,… it welcomes who knows how many different kinds of .networks?

Cerf: In any technologies, that's right.

Nielson: S you felt that particular motivation,, to form a very general, switchable network. If so, that wasa very important point in the history of the internet.

Cerf: It was, I mean we didn't quite do it right the first time, I mean, when we designed the packet formats and everything, how many networks do we need to represent, we had 256, we thought it would be more than enough, and of course we were clearly wrong about that now. But the whole point was that we deliberately designed things so that we knew nothing at the internet protocol layer of the underlying networks. There was nothing about the protocols that took advantage of the characteristics of any of those nets, so in particular we didn't have any special attention paid to the fact that some of them had broadcast capability. It wasn't part of the architecture. Looking back on that, I wished that I'd explored that idea more. But the simplicity of it and the separation of the internet protocol layer from any underlying knowledge of the underlying networks or any knowledge of the applications made for an absolutely generalized networking layer. It knew nothing except connectivity and carrying bits and identifying the destinations. That's all it did.

Nielson: Well, it seems natural to you, that is still a pretty profound concept, isn't it?

Cerf: It was, I mean, it was very deliberate. We absolutely knew that we didn't want to have any accidental ties to specific technology, because we wanted to future-proof the design, and we didn't know what people were going to invent 10, 20, 30, 40 years hence, but we didn't want the basic architecture to fail because it took some special knowledge or it took special advantage of a particular networking technology.

Nielson: That involved TCP as it was originally written and IP had not yet been broken out.

Cerf: That's correct.

Nielson: I've read earlier where you thought IP came out maybe in 1978. I know I've found evidence that said that even as late as 1979 the hosts of the ARPANET were still running TCP 2.5 . So it took a while for IP to emerge.

Cerf: To be split out, so that's the piece of research I haven't done recently. I was pretty sure that by version 3 of the spec, which I thought we had written in late '77 or possibly January of '78, we'd done the split out, but that was just the design. In terms of implementation, different story, so I agree with you there.

Nielson: While on that subject, today we talk mainly about IP on the internet, IP addresses and all that sort of stuff. What's the role of TCP now, or maybe even UDP, in the Internet of today versus IP?

Cerf: Well, I IP was split out as you will remember and then layered on top of it was this user datagram protocol, it was split out to handle real time communications, which didn't necessarily benefit from the determination of TCP to retransmit, put things in order, get rid of duplicates and all this other stuff, because it introduced delay. So if you wanted to have a real-time stream where missing something was not as serious as failing to get the most recent stuff, radar tracking for example, I don't need to know where the missile was five seconds ago, I want to know where it is now. So speech and other kinds of real-time applications benefitted from not insisting on sequential delivery or delivery of every single bit. Today TCP is still a very, very important part of the application space that people use on the net. The world wide web, HDTP protocols are designed to run on top of TCP. File transfers are becoming oddly enough increasingly common now, as even though people see a lot of streaming video and audio, as the speeds go up, it becomes just as easy to do downloads or uploads using file transfer technology, which is running over TCP, which is a disciplined protocol compared to UDP, which is not disciplined, you just blast away. So UDP, however, has been a really important element in streaming kinds of applications or in real-time kinds of exchanges where sequentiality may not be as important. So the two are still both very important to application space on the net, and now we're starting to see other kinds of protocols like real-time protocols show up where time stamps are made available so you know well, when did this thing get sent.

Nielson: One of my recollections is that when we were putting compressed speech on the packet radio network, it was the delay variance, or rather the increase in delay variance, that became one of the driving points toward the use of IP. I think you attribute the need for the separation to Danny Cohen

Cerf: Danny and several others, including David Reed by the way, whom I didn't remember from MIT, I didn't remember him until he reminded me that he was party to these debates along with Jon Postel and Danny Cohen. Danny was the most eloquent, however, he's written some wonderful things about wine versus milk, you know, wine you can wait for but milk you can't because it spoils. Danny was the most eloquent in pushing the idea that segregating the IP layer out and making for real-time applications was important and so, you know, he and Jon and David Reed deserve a lot of credit for pushing that.

Nielson: I remember Jon Postel, when he was at SRI, wrote a report to DCA on TCP design and in that report he expresses a need for a datagram type of protocol. I don't think he referred to speech but he cited some other example for it.So Jon was also part of those wanting a separate IP.

Cerf: Incidentally, the term datagram comes from Louis Pouzin, who Louis was an INRIA, it was called IRIA at the time, the Institut de Recherché d'Informatique et d'Automatique and he was the guy that pushed CIGALE network, which is pure datagram, and then CYCLADES was the system of computers that went around it, but he was very, very vocal about the need to have pure datagram systems and not circuit switched or virtual circuit switched. The vois [ph?] virtual [ph?] was the thing that led to the X 25 protocols that were virtual circuit oriented. So anyhow, just a little footnote in history that Louis contributed that particular term, just like Donald Davies contributed the term packet to our vocabulary.

Nielson: Just to be clear, Davies at the National Physical Laboratory was responsible for the concept of moving information in packets and Paul Baran at Rand who designed a network that could survive a nuclear environment…where emergency action message could be sent out reliably..

Cerf: Paul Baran, right.

Nielson: That work of his at Rand I thought was more of a distributed propagation environment rather than packets, I mean because it was only a short message as I recall that they were trying to get out.

Cerf: No, actually, well you're partly correct, but I would disagree in one respect, Paul wrote this 11-volume series called On Distributed Communication. He was at the time, this is like 1962-63, at the time he postulated digital equipment that would be sitting on tops of lamp posts or telephone poles radiating in all directions using hot potato routing, which basically said transmitted to anybody who will listen, but what he was propagating was not only emergency action messages but also packet voice, because he was trying to reconstitute the voice command and control system with a non-centralized architecture that would recover or could be recovered even after a nuclear attack, where great holes had been punched into the fabric of communication, but as long as there was at last one path somewhere, you could carry these packets. He called them message blocks instead of packets in that architecture.

Nielson: Thanks for clearing that up. As TCP progressed, there arose two particular confrontations: one in the Department of Defense when AUTODIN II was being proposed, and a second one in the International Telecommunication Union regarding the OSI model. Let's take them in that order, TCP versus AUTODIN II? I think at that particular point, maybe more than anybody in the world, you were championing the integration of TCP into both the emerging DoD systems as well as others.

Cerf: So it seems like, it only seems like yesterday, AUTODIN II was an effort to redo AUTODIN, which had been around from the 1960s. That was a message switched network and we were trying to persuade them to use packet switch. One of the biggest problems we had is that Bob Kahn took the early spec of TCP, the December 1974 spec and attached it to a communication that went to the defense communications agency, DCA then, now DSAC today, and in it he said, in the cover letter if I'm remembering right, he made reference to this attachment, saying this is an example of the kind of thing that one could do or would do with this newly devised AUTODIN II. Now the problem is that they grabbed it and they tried to implement that. The other problem is that they were four years behind us, and so we'd already discovered all kinds of bugs in that particular spec, and we'd been through four iterations, we were already at IP version 4 and TCP by the time the DCA guys started doing this AUTODIN II thing. And so I felt like I was fighting myself, because they would get up and they would say, "Well, there are these bugs." And I'd say, "I know there are those bugs. We went through this, you know, we went through four iterations of this thing." But they had gotten locked in to this particular spec, and I remember thinking, "Oh my God, you know, talk about peaking too early." So we had a big battle about this over many years and finally the whole program got canceled thanks to Steve Walker's intervention because they had just gone down a not very robust path. And it's a pity because it could've been a big deal. But when we split in 1983 when we split the MILNET off from the ARPANET and did the TCP/IP rollout, then the military did in fact get their opportunity to use these protocols in the way that they have.

Nielson: The OSI proposal was a very formal structured and layered and proposal. Did that just get overtaken by events too?

Cerf: Well, by the time the first specification for OSI was written in 1978, written by a couple of friends of mine, one of them from Louis Pouzin's lab, Hubert Zimmerman, when they put out their first architecture document, we had just standardized on TCP/IP version 4, so they're starting, as we're sort of completing their work, they're just getting started. And what happened there is that there was a huge amount of documentation, seven layers worth of protocols, a lot of debate, some ideas good ones, some ideas not so good, but not enough implementation, and in the meantime, we're blasting along, it's 1980, you know, by 1983 we're ready to roll this thing out, and as the years unfolded, there was more, especially after the Berkeley release of 4.2 units, which incorporated TCP/IP into it, once that was available, it started to propagate to the academic community, and the OSI stuff just never got implemented enough, so even if it was better, the TCP implementations overwhelmed the OSI specifications.

Nielson: It was overtaken by the acceleration in the use of TCP.

Cerf: Exactly right.

Nielson: Okay, can we break here?

Cerf: Wanna break here? Yeah.

Nielson: In the packet radio network, there was a need to connect the central controlling node to the packet radios themselves for purposes of control, and that had to be a reliable protocol. At the outset that Channel Access Protocol (CAP) didn't have that characteristic, so you contributed some of the ideas embodied in TCP and made that particular link reliable. Do you remember that instance and to help, I will tell you that Dave Retz at SRI was one of the people you work with on it?

Cerf: It's funny. You would have had to, and you just did remind me about that, I wouldn't have remembered it offhand. But I can easily imagine myself trying to push, you know, we said look, we just learned how to do all this and it works, you know, we ought to keep using these techniques because we know they work. This popped up in another context as well. The _______ director, who was at IBM, was working on the International Science Foundation Network, and it became clear as time went on that we needed a new routing protocol, what was called a border gateway protocol, between networks. What information should be conveyed back and forth between them as sort of a peering exchange to say hi, I'm this network and I'm connected to the following other networks, and by the way my autonomous system number is whatever it is. Yacov was very concerned about the reliable exchange of this routing information, because without it, you wouldn't be able to find your way through the net, especially as things keep changing as connectivity varies. So Yacov insisted on putting TCP into the protocol that was used to do IP packet routing. And a lot us got nervous because we said, wait a minute, the TCP is a protocol layer above IP and we're relying on that protocol layer to make IP work, and what we said was, are we going to get into trouble? And we didn't get into too much trouble, but there is an issue that arose because of that. If you had a large number of peers with whom you were exchanging routing information, it meant keeping a large number of TCP connections up and running all at the same time, which have their own TCP control blocks and all that other stuff, so there were just some scaling issues associated with choosing that particular path. But, I'm not surprised that if you need reliability, well use mechanisms that are known to work.

Nielson: And that's what you did here in this case, we didn't encapsulate the protocol in TCP, we just took some of its reliability features and implemented them.

Cerf: Incidentally, just for the record, the sliding window flow control for TCP came straight out of discussions with Louie Pouzin [ph?] and his people at INRIA. So, in fact, Gerard Le Lann [ph?], who was part of Louis Pouzin's lab, came to Stanford University in 1974 and participated in a lot of the design work, and I remember Bob Metcalfe and Le Lann and I sort of lying down of the living room in my house in Palo Alto on this giant piece of paper, trying to sketch what the state diagrams were for these protocols.

Nielson: Louis Pouzin arises all over the place?

Cerf: Absolutely.

Nielson: Okay. There was a time, I believe, in the packet radio network where we actually were capable of binding a terminal device name to an address. We called it Dynamic Name-Address Binding or something like that.

Cerf: Something really, you know, had a nice ring to it.

Nielson: It does have a nice ring. Today, IP addresses are associated only with particular locations. Given the mobility of people, is there any innovation that would ever cause a person to be associated with the location of the node they were using; that the network could learn where they are, or is that too much of a privacy issue?

Cerf: Well, actually, that's an interesting question because Dynamic DNS, Dynamic Domain Name Service, is in fact a reality, and some people implement rapid domain name system updates to try to track the IP address of a target that's moving from one place to another in the network. It's interesting that if you look at the mobile world today the telephone numbers are no longer identifiers of the way you connect to the net, they're simply identifiers of the device. And underneath that is another piece of information about how am I connected to the mobile network now, it's how roaming works. IP addresses are still very much bound to how are you getting access to the network, and we don't actually have the same sort of dynamics that are part of the current telephone system's mobile environment, and I would predict that we are going to need that, especially as more and more mobiles become part of the internet environment. So, my guess is that it will definitely resurface and it will be needed because you want to initiate communications with people whose location you don't know. And that's what the domain name system allows you to do.

Nielson: Exactly. Okay, I'm now going to ask you for a thread through time. The ARPANET came into existence and over time ARPA felt that it was getting beyond their capacity or interest to maintain it, and it was transitioned to another jurisdiction. One could define a thread between the origin of the ARPANET and whatever became attached to it, then to DCA, to NSF, and so forth. First, can you draw that thread for us and then, what parts of that thread did you influence?

Cerf: So, of course the ARPANET itself was a project that I didn't have any involvement in starting, but I certainly got involved in helping to fashion, particularly at the higher layers of protocol. So the ARPANET is going along and around 1980 or 1981, NSF is starting to get interested in networking, and it initiates the CS-net project, the Computer-Science Network project. Separately, there was this bit-net thing based on IBM capability because it's time network. NSF uses TCP-IP in the Computer-Science Network, CSN. Then comes 1985, NSF wants to build a real nationwide network and somehow negotiations to use the ARPANET for that purpose don't quite work out. They use some access to the ARPANET, but primarily they want to build their own separate network. So they build, around 1985, they build this network using Dave Mill's Fuzzball PDP-11 machines, which are essentially routers moving IP packets around. They were running on 56 kilobit backbone and instantaneously overloaded the system, so now they need to move to T1 one-and-a-half mg. They put out a request for proposals and they get back a number of responses, one of which comes from a consortium of IBM, MCI, and University of Michigan. And so those three essentially designed and built the NSF net backbone. IBM was responsible for the router design, so we're now moving into 1986, we're moving into a regime which is much faster in backbone speed than the ARPANET, which never got past 56 kilobits. So by the time 1986 rolls around, the NSF team forms this one-and-a-half megabit network, and the two are primary backbones for the U.S. internet. NSF figures out that they don't want to service every single university in the country because there are like 3,000 of them, and they figured the guys running the backbone should be focused on running the backbone. So they come up with this idea, I don't know whether it was Steve Wolf, or Hans Hunter Brown [ph?}, or others, that they should have intermediate level networks that connect to the backbone and service the universities, so there were like a dozen of them, and they had, you know, names like BAR Net and so on, for Bay Area Research Network. So the idea to create these multiple intermediate level networks was the NSF's notion, Steve Wolf was in charge of the networking activity at NSF at the time. So now we're into 1986-1987, and somewhere along the line, MCI, University of Michigan, and IBM decide they should create a non-profit organization to run the network, the NSF net backbone for NSF, so they create something called ANS or Advanced Networks and Systems. So we're into a period of time now where there is increased demand for getting to higher speeds. I guess I left out one important fact; the decision to use the TCP-IP protocols in the NSF net was made by a man named Dennis Jennings. Dennis was an Irish researcher who was essentially on loan to the National Science Foundation in the super computer systems group, and he takes on the responsibility for the initial preparations for a network to connect the super computers, and he decides to use TCP-IP just like Larry Lanwebber [ph?] and the others who started CS-net agreed that they should use TCP-IP, and this was even before we had done the 1983 roll-out. So, anyhow, NSF-net is now running along, we're pushing the limits, it's 45 megabits a second, 155 megabits a second, and at some point, somewhere around 1995, this is after the NSF net has been around for almost 10 years, NSF decides that it doesn't need to run the backbone anymore, that there is enough commercial networking available to replace the NSF net with commercial service. So how did that happen? Well, what happened is that in 1988, I had left ARPA and gone to MCI, did MCI mail for four years, then rejoined Bob Kohn [ph?] at CNNI, and in 1988 we tried to figure out, how could we break the current log jam which prohibited the use of commercial traffic, carriage of commercial traffic on the NSF-net, government-sponsored backbone. So I went to NSF, and actually to the Federal Networking Council, and asked them for permission to connect the MCI mail commercial system up to the NSF net. And they gave permission to do that in 1988, and it took a year to get the interconnect worked out, a guy named David Eli wrote the software at CNNI. We announced this in 1989, and as soon as we announced it, we had this interconnection. All the other email providers said well, wait a minute, MCI can't have special access, and, so, the telenets, telemail came up, Timenet Suntime came up, and CompuServe's e-mail all came up and connected to the internet through the NSF net. And the consequence of doing that interconnection was that they could all communicate with each other, whereas, before, they were completely separate, so at a higher layer of protocol, namely email, this attempt to interconnect led to interconnection of systems that heretofore had been independent. So, by 1990, ARPA decides that the ARPANET is simply no longer a useful research tool, and so they shut it down and everybody who was using it switches over to access, to the NSF net backbone by way of one or another of the intermediate level networks. Also by 1989, we discovered that three internet service providers have already gone commercial. The Union net, which started out as a UCP network around 1986 or so, starts offering commercial internet service in 1989. The NYSER net, the New York State Education & Research network, extracts from itself all its underlying hardware and everything, and creates something called PSI net, Performance Systems International network, and offers commercial internet service. And, let's see, the third one is CERFNet, it was General Atomics in San Diego, wanted to help build a network to connect the universities in the Southern California area, so they invent this thing called CERFNet, only they were going to call it S-U-R-F Net, and they discovered that Kays Neggers [ph?] and his pals in the Netherlands have already got an S-U-R-F Net, in stands for something in Dutch, and so they change their name to something to C-E-R-F-Net and call themselves the California Educational Research Foundation Network. So, by the time 1994 or 1995 rolls around, we've got commercial networking going on, NSF decides it's going to shut down the NSF net backbone, but it makes a requirement that there be interexchange points, and everyone who was connected to the NSF net backbone, has to commit to remaining interconnected with everyone by peering at these internet exchange points, so the network access points, they call them naps. This is a really important policy move on the part of NSF. They're basically saying we want to shut down the backbone network that provided connectivity to all these different nets, but we don't want to lose the connectivity, so in order to get any further assistance from NSF, you must commit to interconnecting to the network access points. And they did that, and eventually those network access points became independent commercial enterprises or, in some cases, still nonprofits, but they're scattered all around the world now. So, you very quickly enter, around the 1995, 1996 period, the commercial world of the internet which is familiar today. Incidentally, around that same time, like 1994, Netscape Communications gets formed, and the World Wide Web starts to take off like a rocket. That, too, contributed to the public interest in the internet and its utility. So over this period of about 20 years, from 1973 to 1993, 1994, you see this explosive and public interest in the internet suddenly arising, once it becomes available commercially.

Nielson: I remember there was a time in the history of the ARPANET when Digital Equipment Corporation put a message on the ARPANET advertising a new computer.

Cerf: Hah! Yes.

Nielson: --their new latest PDP whatever, and Bob, I recall, came down on them like a ton of bricks.

Cerf: Actually, the specifics are that they put an advertisement for a person, it was a personnel ad, you know, there is a job opening available. I'm pretty sure.

Nielson: It was that, too.

Cerf: We may have to go back.

Nielson: I think there's two instances.

Cerf: Maybe there were two things. But I remember the advertisement for a job, and Dave Russell, who was running the information processing techniques office at ARPA at the time was in London. We were in the middle of an internet-- at that point, I guess it was the internet working group meeting. And he saw that message or somebody forwarded it to him, and you could hear the nuclear blast from, you know, from London. And digital was told in no uncertain terms, and under not circumstances, were they allowed to use this network for any commercial practice at all. This is clearly pre-1988.

Nielson: So, the reason I reflect back on that is that there came a time when the commercial world initiated offerings on the internetwork. There came to be this big elephant in the room called advertising, the commercialization of the Internet. Today it drives all the big money there. Did that happen by decree or did it evolve over time?

Cerf: I think it's-- this is a classic situation where things don't happen until the environment is in a condition to support them. And, so, the internet would not have happened, ARBINET wouldn't have happened were it not for mini computers, in this case, from Honeywell, that served as the packet switches of the ARBINET. You couldn't do it at a reasonable cost otherwise, that particular architecture anyway. So, two things really happened. The first thing is spam, and spam started showing up as soon as the general public gets access to email in a way that is essentially cost free. All the predecessor publicly available email services cost money. I remember we used to charge a dollar for every MCI mail message that went through the system. That price didn't last long, but people actually paid money to get their email delivered. Now people pay money to get the email not delivered because they want to filter out the spam. So, spam was one effect that showed up very quickly when email became a free commodity. The advertising on the net, of course, it was absolutely forbidden during the early period, when it was government sponsored. There were several attempts made to use the World Wide Web as an advertising medium. Banner ads were attempted by a number of different companies, double-click is one example of another one which is helping people place ads in the best possible way. And there were several others that ultimately did not succeed in their business model. And part of the reason that they didn't succeed, I think, is that their model for delivering advertising didn't match up the advertisers very well with the targets. So, along comes Google, well, Yahoo comes first with a monetizing model that was not quite as clear. Google comes along, also with no business model at all when it first gets started, somewhere around 1996, Sergey, Brandon and Larry Page were just trying things out to try to index the internet. And in 1998, they start Google, again, without much of a business model at all. It wasn't until a couple of years later that the idea of associating advertising with what people were looking for was injected into the system, and that was clearly a trigger for a very effective combination. And I think that particular model is going to continue to grow as new media become part of the internet environment.

Nielson: I want to ask a few questions about that whole evolution you just covered. There clearly arises the question of scalability. What worked and what didn't work as the Internet grew from tens of hosts to millions of hosts; that is, as the price of a host decreased to zero? Since you have been as close to this as anybody in the world, what was there about the whole scalable design that worked so very well… that perhaps, looking back, even surprised you.

Cerf: Well, one part of the design, I think, is that the routing structure is actually working better than I ever anticipated. I won't claim, and don't want to claim, much responsibility for the design of the routing systems, because other people really did that work. I remember in the days of ARBINET, we had a routing system that created congestion in the form of these giant swarming ball of packets that would move around and congest different parts of the net. And that was remedied by some significant redesigns of the routing system. A similar kind of evolution occurred in the internet routing, particularly as we moved into a regime where we distinguished between routing with internet work and routing between networks. So the border gateway protocol routing is between nets, and interior gateway protocols are inside of nets. The decision to explicitly break out what is generically called an exterior gateway protocol was in part to prevent an ossification of routing architecture on one particular architecture, on particular design, by one company. As Bober, Eck and Newman [ph?] was the company that made the first gateways, and, therefore, they made the made the first routing protocols for the internet. And when it came time for the NSF net to become part of this architecture, Dave Mills worked on what was then called an EVP, exterior gateway protocol. That was a deliberate design decision to shift and create an interface that everyone could potentially meet. That way they could run their own interior gateway protocols, and it didn't matter to anyone else what they did. It was, once again, an attempt to layer and segregate and simplify. So scaling of the routing system has been pretty impressive. The domain name system came along in order to deal with the fact that the host dot text file that was set around on every morning was not a scalable way of finding out the-- matching between a host identifier or host name and an IP address. And that has scaled dramatically well over, you know, several decades now of use. The ability to put larger and larger routing tables into the routers has been required, and new technology has cheaper storage, larger storage, has contributed to our ability to deal with that. I do have some concerns about continued scaling, however, because, as we run out-- here's an example of something where we didn't scale very well. IP version 4 address space is only 4.3 billion addresses. At the time we decided to use 32 bits, it was 1977, and I thought that you didn't need more than 4.3 billion terminations to demonstrate this technology worked. And, as it happens, I thought that there was going to be a production redesign. If we demonstrated that it worked, then we would redesign it for production. We never got to do that. So IP version 6 is, in fact, the production, in my view, the production version of internet IP layer with the 128 bit address space. But it's turning out to be really hard to introduce it because you can't just throw a switch. When we did the TCPIP conversion from NCP in January of 1983, there were only 400 computers in the network. So we could kind of force a fly game, and even then, not everybody quite made it precisely on January 1st. Now there's 500 million machines on the internet, and maybe more than that if you count the ones that are episodically connected, like a laptop that connects from time to time. So, getting everybody to switch over to IP version 6 at the same time is impossible. That means the operating systems, the routing systems, and everything else, even the application software, have to be able to work with both of them concurrently. Well, that's more than double the complexity because now you get combinations of things happening in V4 and V6 that leave you to be unsure about the state of the net. So this is going to be a hard transition, but we're going to run out in 2010, from the ICAN, INA point of view, the internet assigned numbers authority, well it handed out the last slash 8 block of V4 address somewhere around 2010. And, subsequently, the regional internet registries will have handed out their allocations from MYANA [ph?] to the internet service providers, and at that point there will no new IPV4 address space. And if you want to keep growing the network, you need more addresses, so V6 is critical. We have to get there.

Nielson: So what's your prediction about Version 6 of TCP/IP?

Cerf: Well, it was standardized over ten years ago, and it's been a constant attempt to roll this over up to the top of the hill. I think there's going to be a lot of white water and turmoil between now and 2011, because the people will begin to realize that if they don't have IPV6, they can't continue to grow their ability to reach the edges of the network. More and more devices will come on, mobiles, set top boxes, other appliances, and they're going to need IP addresses. And at some point you just have to use V6 to do it, there is no other choice right now.

Nielson: I understand that. But this whole question, of both instantaneous and backward compatibility, is a tough issue.

Cerf: It's tough and we--

Nielson: It's a global issue. Do we just shut down the Internet to make such a transition?

Cerf: Well, we don't shut it down, but the two are not very interoperable at all. I mean, they're fundamentally different packet formats, and the problem with addresses is that if you're running the old version 4, you can't refer to 128-bit address spaces, there's no place to say anything about it. You know, I don't know whether somebody could invent an addition to the V4 protocol that would let you say go look here for the destination. I don't know, maybe we should look at that. But, in any case--

Nielson: Could you encapsulate that packet in some way?

Cerf: Well, there's a lot of encapsulation stuff, there's proxies, there are other tricks that people are thinking of playing to try to make the two interwork. At the moment, the best advice is go build IPV6, get it running in a dual sack architecture, and then serve both requirements. And you don't care anymore which, whether it's a V4 or a V6 packet that comes in, you can handle both of them. But there will come the day when some device shows up that doesn't have V4, because there's no point in having V4, because it doesn't have a V4 address, and so it's going to be V6 only, and now what? So there's work still to be done.

.**Nielson:** What <u>didn't</u> work very well in this evolution? You've talked about spam a bit. As far as I'm concerned, at least at this point, that's a scourge of the Internet and even drives some people away. Security always seemed to be a stepchild along the way. What about security? What didn't work?

Cerf: This is an interesting problem because in the earliest days, and you participated in this, too, the whole question was do we make this work at all? And to overload it with holy cow, now how to I secure this system to military standards and everything just seemed like a really tough problem, besides which, graduate students at the universities, some of them from all over the world, are involved in the internet development. And it's not the first thing you think of to have military grade crypto and other stuff in the hands of graduate students. So, around 1975, when I was still at Stanford University, I started working with NSA on the design of a fully secured internet system. And it was all classified. I had people, Volbearing Enkland [ph], Steve Cant [ph?] for example, working on the, what we call BCR, or black crypto red architecture. So packet cryptography grew out of that project, and, so, in parallel, we're running, designing, a fully secured, ______________ secured internet using packetized cryptography and other things, which were all classified. But there was no way to deliver that notion to anyone in the public, and so the two architectures kind of, you know, went unparalleled. In 1977, I believe, maybe even as early as 1976, public key cryptography notions started to emerge.

Nielson: 1976, I think.

Cerf: 1976? Whit [Whitfield] Diffie and Marty Hellman published their first paper. The problem is that nobody has implemented anything. It's a concept paper. By the time RSA gets designed, we've just standardized on TCPIP version 4, and no one has any experience of RSA. We don't know whether it works or not; we don't know how reliable it is; we don't know whether its work factor will be overcome by new alga rhythms. And, so, at this point, I am busy trying to get implementations of a standard done so that we can roll it out. And it's going to take five years to do that, not that I knew it at the time. So we didn't really focus on end to ends, authentication, we didn't use public key photography, and it still has not worked out too well. Public key photography is working out very well in unauthenticated confidentiality. So when you go to a web page server, and you do an exchange, you can secure the communications, but you don't know who you're talking to. The idea of having a single central authority for public keying certificates has not worked out well in non-command environments. So, for the general public it hasn't worked out very well. It has worked out in the military; it's worked out in corporations. I'm told when I go to Google, for example, that you can't work here unless you accept this public key certificate and, you know, it becomes part of your communications tool set. So that's an area that hasn't worked out very well. I believe there's still an opportunity to make this work, but its centralized top down command control is tough. We're starting to go into that particular technology for the domain name system itself, which is now vulnerable. People can influence the way domain names are looked up and responses come back in a way that causes the answers to be wrong. And you need digital signatures in order to assure that the information that was put into the domain name system is what you get out or what someone else gets out when they do the workup, so this is called DNS sec for DNS security, and that's being pushed along more rapidly now than before. Another thing that didn't work very well is trying to allow domain names that had text that had gone from other scripts in Latin character scripts. You'll recall that domain names only use A through Z, and zero through 9 and a hyphen, and that's all they're allowed to use for years. But, now, people whose native languages are expressing Arabic characters, or in Hebrew, or in, you know, Chinese, or Hangul for, you know, Korean, and so on, or think of all the 22 languages in India and all the various character sets that are used there. None of them can introduce domain names with those character sets. Even though they can express themselves on web pages that way, because the end code, code says available through HTML and XML, so right now that's another big push, to get internationalized domain names into the system. And it's a lot harder than it sounds. It is really, really tough. It has a lot to do with the fact that Unicode wasn't decide for this kind of computer comparison. But strings, it was designed to make printed material, you know, and all these different languages, you know, readable. So that's a big struggle. Authenticity, then, and security and confidentiality are areas of weakness where the net really needs to develop.

Nielson: I was thinking more in terms other than the main components of decryption and denial of service. There is this whole question of accountability. At the moment it seems almost impossible to ascribe traffic origin. How do I know from where any given message emanated. That ability has not come to pass and, in fact, can it?

Cerf: Well, that's, the rubric for that is authenticity and the need to authenticate a correspondent, or even to authenticate where the message emerged because the architecture of the net allows you to asset what your IP address is, which means some people can asset false IDs, source addresses.

Nielson: Exactly.

Cerf: And that leads to all kinds of abuse, denial service attacks and other things. There are ways to combat that. It's possible at the edges of the net to check to see whether or not the party that's asserting their IP address is at least within the network that's associated with that IP address space. But, to go further than that, you have to do some sort of a digital signature authentication exchange as the device is connected up to the net. And, even then, you may have to do some continuous checking to validate that party. I think we're some ways away from actually having that capability available. At higher layers of protocol the same sorts of things recur. How do I know the email came from you? How do I know that this person's browser is connecting to this website? How does the guy with the browser know he's talking to the real website and not someone who is pretending to be that particular target? All of that would benefit from public key kinds of exchanges and digital authentication.

Nielson: I'm going to skip ahead here a little bit. As the Internet grew, were there attempts to monopolize it? Though in the beginning, certain large commercial organizations didn't seem at all interested, now some can't seem to live without it and some perhaps even take it over. I guess the core of my question concerns network neutrality and can strong commercial forces come into the network and take it in directions that you, Vint Cerf, don't consider to be in the users best long term interests?

Cerf: Well, the answer is yes, as possible. When this system was first designed, the whole idea was, if you follow the following rules you can build your own piece of internet and connect to it, and it should work. That's basically all the philosophy was, let anybody build anything they want to as long as it meets these requirements. What has happened is that the economics of all of this are driving towards aggregation of service providing. People for a long time, for example, when the internet was first made available to the public, most of the public got access to it by dial up modems. This had the following interesting property: You can change internet service providers by dialing a different number. So the overhead of changing was de minimmis. Then broadband comes along, and suddenly it gets harder to switch providers because getting a broadband facility then means sometimes a cable pole, or fiber pole, or something, significant amounts of time and energy go into those broadband connections. And, if you want to switch, it's not a question of dialing a different number, it's a question of picking up the phone and talking to somebody about bringing another truck rollout to provide you with a different access channel. So the competition among broadband carriers is very limited in my view. In other parts of the world where there isn't very much competition either, broadband has been very successful anyway. So, if you go to South Korea, for example, something like 70 percent of the country is on broadband capability. The UK is doing very well in this regard. People in the Netherlands are doing very well, and Hong Kong and Singapore and so on. Part of the reason for that is the regulatory regime insisted that the providers have open broadband access wholesale capabilities so that other providers could offer service. In the United States we have this belief that competition, that deregulation guarantees fair competition, and my experience is that ain't so. And, as a result, we end up with that old joke, you know, what's worse than a regulated monopoly? An unregulated monopoly. And, so, we have a situation where there are limited players offering broadband services, and, for them, some of their key revenue producing applications are eroding away. Voice over IP is free voice service, and that was the principal thing that generated revenue for a lot of the telephone companies. Internet access to video streams is, essentially, free video, but that's eroding away, potentially, the business models of the cable companies. So those two providers, the cable companies and the _______________ are trying to find ways of hanging onto their revenue or preventing it from evaporating, or introducing new revenue streams. So there is a tension, serious tension, between openness and this ability to preserve old business models. I think in the long run the tension is going to lead to some breakage. I hope that openness wins because openness is the essence of innovation. And in its absence the internet wouldn't be where it is today. We need to preserve that.

Nielson: I think this is evident in the connections available to us, cable verses DSL. Wireless is also coming into existence, but it is slow and slow in coming. One just gets the creepy feeling that if these offerings ever merge into the same commercial entity, then where will be the room for breakout?

Cerf: Well, I think the only answer for this is that we need to introduce a regulatory regime that inhibits the abuse of exactly that scenario. And right now we don't have an environment that is friendly to that notion of regulatory oversight. Personally, I think there isn't going to be much opportunity for serious competition in broadband, even with the wireless coming along, and that we really should be looking at how to prevent the abuses and open things up. I think we need to cut there for the next tape.

Nielson: Continuing on with this issue of whether the access providers get bigger, more powerful, and more arbitrary in how they deal with their user community, there's this question of video and large capacity users. Do you have any opinions about whether there should be cost schedules that present higher costs for those people who flood the Internet in some way? Express your feelings about the very large bandwidth offerings.

Cerf: Well, I don't know that I want to get into details about any specific providers, but I will say two things. First of all, that the technology is around for traffic shaping and limiting. So if a party believes that they have paid for access to the Net at a particular bid rate, then they ought to be able to get access up to that bid rate. Then after that, they should reasonably expect to be suppressed. If the capacity is available at the time because of idleness or something, I don't see any reason why you wouldn't let somebody go beyond their rated limit. But, there are ways to do this that are not as abusive as simply terminating the remissions or blowing somebody off the Net. So, I'd like to see a little more attention paid to fair mechanisms for controlling the data rates, but I really think truth in advertising is important here. If you're really not going to deliver to somebody what they thought they bought, which is open access to the entire Internet at a given data rate; if you can't provide that, then you should tell them what you can provide and give the consumer some understanding of what they're getting. This may require some regulatory oversight. It may be that the private sector is unable to come to grips with the need for fair kinds of access to the Net. Our big concern, and mine personal big concern is that there will be unfair discrimination among various suppliers of service of one kind, or applications of one kind of another and then demand made for payment for access to a particular consumer whereas the consumer thought that they had just paid for access to everything on the Internet, up to whatever rate it was that they had purchased. That brings up one other point. Most of the broadband services in the United States are asymmetric. While that served the surfing style for a long time where you mostly downloaded stuff, now we're seeing people generating their own content; so bloggers and YouTube and other kinds of things. We need symmetry for that. You also need it for video conferencing. You need it for applications that have a need for low latency and want high speed in both directions to reduce latency and not just to carry a lot of bids. I'm starting to see this symmetric services coming up more outside the U.S. than inside, but at least one provider has offered a 20-megabit ______________ service for Internet access, which I understand is accessible for residential, as well as business purposes. So, maybe we're slowly moving towards symmetric. But in the U.S., you don't get gigabit per second service like you can get in Tokyo or in Stockholm or some other places.

Nielson: You mentioned YouTube and the multiple sources out there for video streams. Is that the need for the symmetry you're asking for? Are you going to see the time when you pay for the amount of capacity that you use? Suppose you have a 20-megabit connection and will the e-mail user be asked to pay the same as the user that has a video source.

Cerf: Well, let's see. You'll notice that in the traditional world of telecom, for many years, it was the case, especially for residential use, that you paid for the amount of time that you use the system, and - well, and distance too, but the distance thing disappeared fairly quickly. The time-based stuff disappeared in favor of a flat rate instead of rates varying depending on how far away you called, and then finally a fixed price for virtually unlimited service. I mean how many - a thousand minutes a month or something. How long can you talk on your mobile? So, we're seeing a trend towards fixed price services and certainly, Internet service has always been a fixed price service. As far as I'm aware, it's never been a capacity limited service, although I know understand that there are some contracts that people didn't know the details of which said you can only transmit a certain number of bytes per month, after which we will cut off your ability to transmit. They would just put limits on the total. I understand all of these maneuvers to be attempts to allocate the resources, which apparently are insufficient to meet all the demand in some way which some people would consider fair. All the experience I have with this sort of situation of limited resources is that it's often cheaper to just build more resource than it is to try to figure out a bunch of different tools and mechanisms for rate limiting everybody. So there's a balance here someplace to be struck. I do believe, though, that symmetric service is going to be absolutely essential. I'd like to say one other thing, though, about video in particular. The focus of attention is constantly on streaming video. I'd like to suggest to you, especially as the speeds go up, that watching streaming video is probably not the best relevant way of dealing with that medium. If you look at what's happened with MP3 or other audio recordings, people download the audio faster than they can actually listen to it because it doesn't require a big data rate to deliver audio. So, iTunes, for example, or iPod, you just download everything and play it back whenever you want to, as many times as you'd like. Video can have the same property and in fact, YouTube is not a streaming service. My understanding is that you are literally doing a file transfer and that before the file transfer is over, you can start interpreting the file. But the actual mechanics are file transfer, not streaming video. So the streaming is coming off of your own disk, not coming from the Net. So if there's some stuttering going on there during the download, it doesn't matter because you're not actually watching that. You're watching a replay from your own local source. I think more and more people are going to go into the download the video and play it back when they want to. In some cases, it'll be faster than real-time dramatically. If you had a gigabit service, it would only take you ten or fifteen seconds to download an hour's worth of video. That's video-on-demand. That's very different from streaming.

Nielson: Yes, and there's a market for both. I mean there will be real-time events that you will want to watch.

Cerf: Exactly, or video conferencing, which requires you to be in real-time. But most of the video that people watch is prerecorded material anyway, so you might as well just download it. If you don't mind a little bit more of a rant on video, one thing that we didn't do in the Internet very well was take advantage of broadcast media - we mentioned this earlier - in the context of radio and satellite. If we started to get really smart and we started having devices that are on the Net that are programmable, you can imagine doing something kind of like TiVo or even Netflix where you go on the Net, you decide what it is you're interested in watching and the system organizes the scheduling of transfers of those videos to your local disk. If we had broadcast media available, you can imagine scheduling the broadcast, telling the receivers it's time to start recording this broadcast and then transmit the data faster than real-time. So, instead of having 500 channels of mostly junk, you might have 100 real-time channels available to people and the other 400 are aggregated together into a big fat huge pipe. If 10,000 people wanted to record *Gone With the Wind*, then you schedule all of their equipment. You do one transmission of *Gone With the Wind* at 100 billion bits a second and everybody records it. If you missed a few packets, you recover from that on a point-to-point basis. But the end result is extreme efficiency in the delivery of large video files and now, let's stop talking about video and let's start talking about digital anything. So, it could be a piece of software. What if it's a big update to somebody's operating system? You do it the same way. It's a big digital file. Maybe it's books. Maybe it's other media. It isn't just video that requires big-- What if it's medical charts and things like that? MRI images. We can use these broadcast techniques and high-speed techniques to deliver this content far more effectively than if we focus only on video streaming. I need to take a break and get a drink. Okay.

Nielson: So I guess the same is also true of the world's large libraries, for example, the Oxford Library that Google would like to digitally encode. The requests for those would be pretty much the same.

Cerf: It could be very similar. It's not just Google that wants to do that. A number of places are interested in digitizing materials, which today are not available in digital form. Of course, all of these digital forms of music, video, books and everything else lead to big debates over intellectual property and how it should be protected. My honest belief, Don, is that our intellectual property models of the past have all been oriented around physical copies of things. Because of the ease of duplicating and the inexpensive ability to store digital material, it may be that this notion of physical copy isn't the right model for compensation people for their intellectual property work. I'm not sure what the right answer is, but I think we need to back up for a bit and ask ourselves, if it isn't copies, then what it is, because I think it's going to be impossible to stop people from digitally distributing copies of material they have.

Nielson: Again, digital signatures come to mind. They might become an intrinsic part of any intellectual property such that it can't be reproduced without corrupting that digital signature.

Cerf: For example, what rights do you have? If somebody gives you a copy of something that's encrypted and you have to purchase the right to open it up, to decrypt it, that's one possibility. Then you get into somebody publishing the key. But that could be specifically declared an inimical [ph?] act just as currently duplicating material and distributing it illegally is inimical. It's a statutory violation of the copyright law. So, it's like with everything else. We say to people, "Look, there are certain things that are societally acceptable and some that are not. In our society, this is not acceptable. If we catch you doing it, there will be consequences, even though we know we can't prevent you from doing it ahead of time."

Nielson: Okay. Let's switch gears a little bit back to something that you were more personally involved in. There have been a number of regulatory bodies in the history of the ARPANET and the Internet. It starts with the Network Working Group and ultimately the International Network Working Group. You've been on those and served as chairman of lots of them. Again, I'm looking for a thread through time of these because I think most of us get confused as to whose doing what and when. Can you work a thread through time that goes from the Network Working Group clear up to ICANN? And also, what takes place under the umbrella of ICANN?

Cerf: Okay. Actually, this is an interesting question because I don't think I've ever taken that particular thread slice through all these things.

Nielson: It isn't easy.

Cerf: After I got to ARPA, I came in 1976. Somewhere around 1979, Bob Kahn says, "If you get hit by a truck, what's going to happen to this project?" He was right. That was a risk factor. So he said, "You need to do something to make sure that there's preserved knowledge of what this is about." And so, I said, "Well, let's take the people who are the leads, technical leads in the various projects that are associated with the Internet, including people involved in packet radio and packet satellite and so on, and let's form a group so that they all know what each other is doing." We called it the Internet Configuration Control Board. We picked the most boring sounding title to keep people from demanding to be in this little group of people. When I left ARPA in 1982, late 1982, just before the TCPIP cut over in January of 1983, a guy named Barry Liner [ph?] picked up my responsibilities and renamed the ICCB and restructured it to be the Internet Activities Board, made up of a number of taskforces focused on different aspects of Internet, whether it was security or the basic IP routing layer and so on. That IAB thing evolved over time and, let's see, somewhere around 1989 or so, there was a restructuring of the Internet Activities Board into an Internet Engineering Taskforce and an Internet Research Taskforce, all kind of reporting up to the Internet Activities Board. So what happened, IETF was one of the working taskforces, but it got much, much bigger because it was the place where a lot of the standardization was taking place. So, we restructured into these three bodies - IAB and the Research and the Engineering Taskforces. And then, about 1992, a lot of us thought that-- Well, actually, we were required/forced into creating another body called the Internet Society. The reason we were forced into that is that in - it was probably somewhere around 1987 or so, the IETF was starting to get so big that it needed secretariat help. So I was at CNRI at the time and I actually hired Phil Gross, who was the then chair of the IETF, to organize a secretariat to support the IETF and then CNRI went to the National Science Foundation and the Federal Networking Council and asked for contracts to support that secretariat. Later, somewhere around 1991 or so, I was told that the research guys were concerned that they were spending research money to support a secretariat for a group which was largely now commercial enterprises, and that they thought that wasn't a good use of research dollars, and I had to agree. So they said, "We'd like to cut back on the expenditures" and I thought, "Well, how do we get money? We have to have an institution that can receive contributions." So, we invented the Internet Society to do that. I became the first president and somewhere in there, I wound up serving as chair of the IAB for a couple of years. So, ISOC was started in large measure to support the Internet Engineering Taskforce. The model I had in my head was kind of a professional society like the ACM, and I also thought there was a kind of double entendre here because I thought the Internet would actually create a society, an online society of people who were using that, and that was an interesting concept. So ISOC comes into existence and it has a major focus to support the IETF. And as we come into existence, we import the Internet Activities Board and rename it the Internet Architecture Board. So that's where the change in name came was when it began operating under the auspices of the Internet Society. There was a period of turmoil where it wasn't quite clear whether the IETF wasn't happy with the decisions that the IAB were making and there was a big debate and we sort of turned things over on its head and now, the Internet Engineering Steering Group, which runs the IETF, became the principal authority for standardization and IAB became kind of advanced, looking at architectures and the IRTF is still around doing stuff that isn't ready to be standardized. So that's now the structure.

Nielson: Excuse me; are all those underneath the Internet Society at this point?

Cerf: They're all underneath the Internet Society now. So the IAB and the Internet Engineering Taskforce and the Internet Research Taskforce are all, I wouldn't say underneath exactly, but they are supported within its umbrella. The intent is not control at all. It's a support relationship, and it's gotten much better in the recent past because the Internet Society made a bid to run the .Org comp level domain, which it won. So it has a revenue stream now, which is better than passing the hat around to industry, which is what I did when I was president of ISOC in the early days. So, ISOC persists. It's been around since 1992, and it will be coming up on its 20th anniversary in 2012. The next thing that happens is that Jon Postel, who had been the Internet Assigned Number Authority for many, many years, realized somewhere around 1996 that the Internet has now become a commercial enterprise. The worldwide Web is taking off like a rocket. The Dot boom is happening. There are liabilities associated with the money that is being made in the Internet context. My understanding is that the Information Sciences Institute and USC, the parent body, saw the IANA functions as having risks now associated with them that they didn't have before. When Jon was making decisions about who runs this top level domain or who gets this Internet address base, there's value, monetary value associated with those things and therefore, potential liability and risk. And I think Jon generally believed that this was no longer a game for two or three people. This had to be institutionalized. He didn't want it to be a body that was strictly revolving around a cult of personality. So, he started to work on something that was called the Internet Ad Hoc Group, IAHG, which or maybe IAHC - Internet Ad Hoc Committee - to try to figure out how to institutionalize the domain name system and the management of the Internet address space. For two years, there was flaring debates and lots of ad hominem attacks because people saw the Dot boom happening and thought, "Well, if I can get control over some part of the domain name system, I'll make a mint." So, there was a lot of fighting going on and finally, the White House stepped in. Ira Magaziner was asked by President Bill Clinton to go get this process sorted out so that we could privatize the management of the domain name system and the Internet address space. So he stepped in and went through a series of green paper/white paper exercises. In the end, around 1998, the white paper described an organization which would be in the private sector, but which would manage the functions that Jon had been doing on his own or with a small group for many years. Sadly and ironically, about two weeks before ICANN was formed, Jon Postel passed away. He had a heart problem that required surgery and he didn't survive the surgery. So, the ICANN organization, the Internet Corporation for Assigned Names and Numbers, came into being as Jon leaves the scene. I didn't have much direct involvement with it, in fact. I was active in and supported the Ad Hoc Committee and I did engage with Ira and others on debates with the green and white papers, but I didn't feel the need to inject myself into this ICANN activity, although I did show up at the very first meeting where the entire board, which was formed to run ICANN, was trying to decide who its executive director or president was going to be. I showed up at that particular meeting, endorsing Mike Roberts for the job, and he got that job. Ester Dyson sort of volunteered to be the first chair and so he had gone off to a rocky start, not because of the people. They were very, very talented and very qualified members of the board. Mike Roberts has proven himself many times in EDUCAUSE and EDUCOM and even at ISOC as one of its first executive directors. But the community was fighting over who had charge of what. It was the same fight that had gone on in the Ad Hoc Committee. So, ICANN has a pretty rocky start. I got put on the board of ICANN in 1999, in November and after serving for a year, Ester Dyson decided to step off the board. She had only committed to two years of work, and I was elected chairman and I served as chairman all the way up until just a couple of weeks ago when my term limits were up. The bylaws say you can't have more than whatever it is - three terms and I've had three over a period of eight years. So, this whole process of trying to govern the Internet has evolved from a casual thing because it was a private network, it was sponsored by the Defense Department as an experiment to something which is worldwide in scope, has enormous political, technical, social and commercial impact and a much more complex environment. Some governments would like more control over the Internet. Others would like to leave it completely in the private sector. There continues to be major issues that have social consequence where people abuse the network and access to it. Social norms are not uniform around the world, but there may be some things that everyone would agree are totally inappropriate like child pornography. And so, I'm anticipating, as time goes on, that there will be need for serious thinking about Internet governance concepts and the most important point I can make historically here is that what ICANN does is a very small piece of the broad question of Internet governance. That has to be repeated over and over and over again. There are some organizations that, in the past anyway, have believed that by taking over what ICANN does they will somehow control the Internet, or that the relationship between ICANN and the Department of Commerce, which inherited the Internet-related work on the domain name system from the National Science Foundation, which inherited that from the Defense Department. Some people believe, outside the U.S. especially, that the Department of Commerce controls ICANN and ICANN controls the Internet, all of which is a completely wrong syllogism. But if we are going to enjoy this gigantic global system, we are going to have to accept that there are some people who abuse it and we have to do something about that. We can't just ignore it. You may recall the law of the sea and the big question about the essentially shared resource that the oceans represent. How do we draw lines and limits? It took 20 years to come to a treaty agreement. It may be that it'll take 20 years to establish a law over the Net, but it may be necessary.

Nielson: And in a sense, one wonders where a uniform governance body could come from. Is ICANN, then, sort of the end of the road? I mean does it answers to anyone other than itself? So the question is, where, in your opinion, will that governance come from?

Cerf: Let's see. First of all, we may want to pick apart a little bit the claim that ICANN doesn't answer to anyone but itself. To first order, you may be right about that because the most recent relationship that was established with the Department of Commerce has a set of objectives in a joint project agreement between ICANN and the DOC. Those objectives were generated by the board of ICANN, not by the DOC. So the relationship has become less and less onerous at the Department of Commerce moves away from sort of directing what goes on to essentially an oversight role to see whether or not the organization has fulfilled its objectives according to its own measures. I would point out, however, if you look at the board itself and ask, "Where do the board members come from?" you'll see that they are overwhelming outside of the United States. So, it's a very international body. It's a highly distributed organization with a very bottom-up focus on things. So policy gets developed by the supporting organizations, not by the board. The board ultimately has to adopt policy, but the effort that has gone into the structure of ICANN and the restructuring of ICANN has been to generate policy from the bottom-up as much as is practicable. We tried one experiment with public elections for board members, which, in my opinion, didn't work out very well because you couldn't qualify the electorate. It could have been anybody. You don't know who they are, how old they are or anything else. So, I don't think that's a very sensible model. There is, however, an at-large advisory committee, recently created regional at-large organizations, which are made up of at-large structures like - well, what would be an example? IEEE would be an example of a potential at-large structure. These are organizations of individuals that could contribute to policymaking. So, I think we have some potential for the general public to have avenues into public policymaking for those things that the ICANN does. But that doesn't necessarily inform policymaking in the large for what will be considered internationally illegal on the Net. I don't know what the answer is to that.

Nielson: From the beginning, there's been an openness to the Internet. Now, comes ICANN and there is a certain openness in it. You may quibble as to whether or not it answers to anyone, but is there going to be a point, beyond the self-imposed constraints of, say, specific cultures, where a more global governance will occur? Do you see any end state to the whole issue of Internet governance?

Cerf: Well, I hope that it stays as open as possible, first of all. I think the likely scenario and the one I would prefer is that there are two tracks going on here. One track says you need to structure that portion of the governance that is ICANN's responsibility, which I repeat, is a very small part, having to do with domain names, Internet address allocation. Structure those rules in such a way that if there are local laws within a given national boundary that those can be applied, but they can only apply within the jurisdiction of that organization. So, the "who is" policy, "who is" database has been a big debate. Up until now, everything in that database has been open to the public. But now that the public is registering addresses, they're saying, "I should have my privacy protected. I shouldn't have to put my e-mail address or my telephone number or my postal address up in public just to register a domain name." There are big arguments over what's private and what isn't in various jurisdictions. So, I think we need to adapt whatever policy ICANN adopts to the reality that there will be different definitions of privacy in different parts of the world. By the same token, there will be things that are considered illegal differently in different parts of the world. I hope, however, that the other track says, "Let's consider, together, what things we all uniformly agree should be permitted and what things should be prevented or prohibited so as to create, in a sense, a law of the Net well beyond what ICANN is responsible for." But we've done it with other common infrastructures. We have international rules and we try to enforce them as best we can. I don't see why the Net should be any different than that.

Nielson: But as you said before, those are not easily arrived at.

Cerf: That's correct.

Nielson: Just who should deliberate over those rules?

Cerf: Well, it's not a body.

Nielson: I know there is not now a body. But given the difficulty that bodies such as the World Court have in deciding or enforcing a particular issue, isn't this going to be a difficult thing to start?

Cerf: But you're making an assumption here in this part of the debate, which I think is not the same one that I'm making. If you look at the difference between world court or the International Court of Justice _________________ or something and a multilateral treaty, it's the multilateral treaty that is the vehicle of interest here. So we have conventions that are expressed in treaties and those conventions have to do with establishing the 200-mile boundaries for fishing or hikes and things like that. That's the kind of thing that I see. I don't see a central body here. What I see is enforcement as a consequence of cooperation. That's what Interpol is all about. It's cooperating law enforcement across international boundaries, enforcing things that we all uniformly agree should be enforced.

Nielson: That's an ad hoc world then. You solve one problem in a multilateral treaty and then you go to another multilateral treaty and solve another problem and so on.

Cerf: Possibly, although you could imagine a multilateral treaty having to do with Internet that might be aggregated over time. On the other hand, there are issues like intellectual property protection, which already have a venue in which to develop agreements like the World Intellectual Property Organization. So there could be, for some of these issues, there exist bodies for establishing policy. But the aggregate of all the multilateral agreements is essentially the framework in which the Internet services are delivered.

Nielson: One more question to end this part of the interview. Since you left MCI the second time, you have become literally a man of the world. You are engaged all over the place; on advisory committees, corporate boards and other things. How would you characterize your last five years or so? Now that you're no longer chairman of ICANN, what are your jobs nowadays?

Cerf: Well, if you exclude ICANN, having departed the board literally just a few days ago, I continue to have responsibilities at Google, which do cause me to travel around the world anyway. We have engineering offices all over the place. Part of my job is to show up like an intellectual bumblebee and try to pollinate everybody with not my ideas, with other peoples' ideas that I've picked up in visiting from one place to another. It turns out to be fairly hard to maintain a coherent feeling of togetherness when people are spread over 24-hour times zones. So trying to help in that process is part of my job. Part of the job involves policymaking, or development anyhow and that's both internal and external. We care a lot about policies that are national or international in scope that might have a direct bearing on Google's ability to do its business. I'm also the recipient of a large number of proposals that come in saying, "I just invented this. Would you like to buy it?" or "Please buy my company," or "Would you give us money so we can develop this idea?" I appreciate the opportunity to see some of these interesting new ideas. Some of them turn out to be sufficiently interesting that I'll try to find the right person at Google to dive more deeply into it. I have not had time to dive as deeply into the technologies at Google as I would like. But having finished my term at ICANN, I'm going to take some of that time back and do a very deep dive into the Google technologies because I think they're really interesting, especially the ability to scale. When we talk about the numbers of processors that are all running at the same time to perform the functions that Google does, it is enormous. Then I have these little side projects that I'm doing. There's the one with the Jet Propulsion Laboratory to fashion a new set of protocols to run Internet essentially across the solar system. That's come a long way in the last ten years that we've been working on it to the point now where it's possible to credibly say to NASA and the other space agencies, "If you will adopt this set of standards, you will create an opportunity for all space missions to interact with each other and to support each other by sharing the resources." We're pushing very hard to persuade people that that would be a good thing in the long run. So that's a little side project that I've gotten excited about. I also have five books that I want to write now that I have a little bit of time. One of them I want to write is called, *I Heard That*, and it's not about me and my hearing impairment. It's about my wife and her recovery from total deafness after 50 years with cochlear implants. It's the most incredible technology and a wonderful impact on her life. The second one is, let me call it a history of the Internet, but it's too grandiose. It don't know enough of everyone's involvement to be definitive, but there are periods starting about 1970 or so, 1973 to be precise with Bob Kahn, all the way up to about 1993 where a lot of what happened has not been very well articulated and I'd like to make an attempt to do that. I know I'm going to need a lot of help with that. The third book is called, *Bindings* and it's something I committed to write when the Marconi Fellowship was conferred on me. The short story here is that I'm fascinating by the notion of bindings, which are part of programming languages. When you write a program and you say, "A+B=C" and you bind values to A and B, then C gets bound to the sum. This notion of binding value to a variable applies in our daily lives. And so, when you register at a hotel, you are bound temporarily to your hotel room. It's an important binding because if you forget what the hotel room number is, you don't have any place to sleep. There are a whole series of bindings like that that influence our lives. Some of them are really long-term bindings like DNA.

Nielson: Some are called marriage.

Cerf: Yes, marriage is one. Well, you hope it's long time. Sigrid and I are going to celebrate our 41st anniversary. You're bound to your DNA. You're bound to your relatives. You're bound to your job. These bindings have various influences on us, and I want to explore that, plus chemical bindings and other kinds of notions that really affect us as organisms. I mean if it weren't for electromagnetic bindings at the chemical level, we wouldn't exist at all. So that's going to be fun. And then I have two other books to do. One is a book of poetry that I write. Nobody knows about that, but I've written poetry over the last 50 years. Second is a book of anecdotes about sometimes and like the stories that we've been talking about. So that's my plan for the time that's been made available by stepping down from the chairmanship of ICANN.

Nielson: Well, you have enormous energy. If anybody else told me they were going to write five books, I couldn't believe them. You can possibly get it done.

Cerf: Well, actually, by telling you this, I'm sort of making it public and it'll be a forcing function because I don't want to embarrass myself by failing.

Nielson: I understand that. I wrote one book. It took me about six years to do it with a lot of interruptions as well, but it can become a tar baby.

Cerf: I can believe it.

Nielson: Let's break right now then and change tapes.

Nielson: We're going to switch into areas that are a bit more open. I call them philosophical, as in espousing new ideas. I'm not sure you will view these that way but nevertheless, let's proceed. You have on occasion regaled, I would say, at the success of the people that you have known, for example, those you became associated with at UCLA or students of yours at Stanford. Many of these people have gone on to do notable things and been successful by almost any definition. To an outsider, that seems a bit unusual. First of all, do you agree that there's been an uncommon rate of success among your contemporaries?

Cerf: Yes, I do agree, but I'm not sure that I would attribute it to me, particularly. Example: The community that was formed by the graduate students that were working on ARPA projects, ARPANET in particular-- it forged relationships which had gone on for decades. And a lot of the people who were part of that community of graduate students have gone on to extreme successes. Some have sort of faded away, but a lot of them are still working and still very active in computers and networking and related areas. There was-- something came out of that experience which said there is brilliance everywhere and that no one person has all of it. And in projects like this one and the Internet and the things that have subsequently developed, an openness is really important to allow contributions to come from any place. I mean, why would you not want a kind of philosophical base which invited everyone to contribute their best ideas. If you look at Jimmy Wales' <ph?> Wikipedia, it's as good an example of that philosophy at work as any I can think of. And for me, it would be unthinkable not to try to ingest into a program the really good ideas coming from elsewhere. Sometimes, you can't do it. I mean, the timing isn't right, it's like the _____________ photography and the question of authenticity in the Internet, we never quite got that to work in the earliest stages. But why would you not want to assure yourself that there is an avenue for creativity, no matter where it comes from, to enhance the value of a project that you're working on.

Nielson: I didn't necessarily mean to put you at the center, that you were responsible for their success, but there is still is an ingredient here that —is interesting. For example, did these years define a point in time where an important technology was ready to explode into the world? Were there, for you and your colleagues, a lot of avenues for opportunity that might not have been pursued were they to have worked inside a telephone company, for example?

Cerf: Well, there's an analogy here that might work a little bit. I think of the Internet as being similar to a road system in the following sense: The road system itself doesn't dictate absolutely what vehicles you can drive on it, although it puts some constraints on them. And Internet has a similar character to it. The idea that you could invent something, and I don't mean you, I mean a group of people, could invent something that allowed for a lot of variation within a certain set of constraints which led to interworking, interoperability, ability to work together, to share the common resource. That philosophical notion I think was essential to the success of the Internet. And to its ability to absorb and support new applications that hadn't been invented, that no one ever thought of. And I'm sure that 99 percent of all the applications for the Internet haven't even been thought of yet over time. And frankly, what excites me is being a part of that environment, which allowed this sort of invention and innovation to happen, and for people to benefit from it. Some people ask me, aren't you, you know, don't you wish you'd patented the Internet or patented TCP and IP, and of course, the answer is no, because if we'd try to do anything like that Bob and I would have never succeeded in getting the protocols accepted as international standards or adopted for use, because that would have put a barrier in the way, given an excuse to say, well, I don't want that because I have to license this patent.

Nielson: I think it enabled a rapid growth in the world of information services that didn't previously exist and it was important enough to attract people from all regions of the world. It was truly world changing. I don't think there's any question about that.

Cerf: So, let's go back for just a moment to something else that you said but which I didn't-- I don't think I fully appreciated until just this moment. You talked about the environment being ready, and I think a lot of things don't happen until conditions are ripe for them to happen. And then sometimes they happen in parallel because, you know, multiple people think of the same thing. Let's talk about software for just a second. The proliferation of computers only happens because the costs are low enough that they can be afforded by an increasingly large number of people. The fact that they can be interconnected is the consequence of the Internet's availability. When you get into the software world, you're in a world which doesn't seem to have-- it doesn't have any limits. It's an endless frontier, because the limits are only dictated by what you can figure out how to program. So this virtual space that's created by software and the network is one which is able to absorb virtually every idea that's expressible in the form of software. So if you want to invent another environment, whether it's, you know, a two or three dimensional second life or world of war craft or new applications that involve physical devices that are being controlled remotely, whatever it is that you have in mind to do, if it's expressible with software, can somehow find a place in this environment. And for me that suggests that there really isn't any limit anymore. We've entered into a time where we've invented an endless space for invention, within the context of software. And I think that is probably the most important characteristic of what the network and the computers that are on it introduces. And now, of course, it's whatever people are interested in potentially creates new monetary opportunities and feeds this virtual circle of invention and innovation.

Nielson: In the early days it was hardware that was the controlling factor. It was so expensive, and software cycled through that. Of course, now, you're painting this picture of this giant virtual tree that is implemented in terms of software and the hardware just cycles through it, but adding increasingly greater capability.

Cerf: That's actually a very good analogy; I like that.

Nielson: It's amazing, and I had not thought of this virtual space before, but it has no limit. You know, anybody in the world that has a need, if they choose, can attach a new branch to that tree to serve their individual or collective purposes.

Cerf: Grow your own branch basically.

Nielson: Grow your own branch. Right. I don't know about the roots, but the tree, anyway. When I first started reading about this rather small group, of which you were a member, I thought of it almost as a closed community; that in the beginning there was BBN, MIT, UCLA, and perhaps a few others.

Cerf: SRI's a part of it, too, as you know.

Nielson: A bit.

Cerf: -- the ARPHANET, for Pete's sake.

Nielsen: The first guy to call a Network Working Group meeting was an SRI guy by the name of Elmer Shapiro.

Cerf: That's the name I was trying to remember.

Nielson: He's the fellow that also worked for Larry Roberts and helped write the RFP for the IMP. But he and a few others like Jeff Rulifson, didn't choose to follow this kind of work. That's when SRI left the ARPANET development effort with the important exception of the NIC, part of Doug Engelbart's lab.

Cerf: But look how important that was. And remember we were using Engelbart's online system to compose an awful lot of the material. We had our own little web, it's just that it was on one computer and we had to get to it through the Net. And it didn't do everything that the World Wide Web does, but on the other hand, the World Wide Web doesn't do everything that Engelbart's system did either. So, yeah, this was-- I'm sorry, but this triggers a point that's worth making. A lot of people say, "Did you ever imagine that the Internet would explode the way it did?" That's almost always the first question people ask. And of course my first reaction is to say no, and then I realize, well, actually we did have enough experience by 1973, anyway, to see the potentials, right? Because we were using Engelbart's online system by then, email had been invented in 1971, we were doing file transfers; in a few cases, multi-computer applications. So a lot of the technology was there, that people are familiar with today. And by that time Xerox Park was around and they had the personal computers, they called them the Alto's <ph?> and they were $50,000 each. They had Ethernet; it was only running at three megabits a second, but that was a lot faster than 300 baud. So a lot of the pieces, which are familiar today, were familiar to some of us 30 years ago. And so it's not a big stretch to imagine a world in which these things are common. But I confess that I didn't really appreciate what would happen if a billion people had the opportunity to share information.

Nielson: You should be forgiven for that. That's why I attach so much significance to this event we're celebrating, because I think that was the dawning of an awareness that very dissimilar networks could interconnect. There are a lot of people that can embrace this particular technology and what it can provide. Curiously, within the telephone companies of the day, there was a resistance to this technology. I remember going back to New Brunswick or someplace on a consulting job once and I think John McQuillan was there, and some others, and you just couldn't interest them -- it was like beating your head against the wall to talk about this area. Why didn't it take, in your opinion, and why did it take those companies so long --- it almost had to become obvious before they participated. What was their reluctance? Do you know?

Cerf: It's a very interesting question why there was such dismissiveness in all of this. One of them is that they'd been around for a long time. They were one of the biggest industries in the entire world; they knew what telecommunications was. "We'd been doing it since 1876; who the heck are you?" and "by the way, this packet stuff, this is crazy, it's uncontrollable, it's not managed, and what do you mean, people colliding with each other and transmitting--" you know, "this is nuts." And so it just didn't fit with their world model at all. The idea that everything was managed and controlled and everything was reported and call e-tail records got generated and bills got produced, all of that just was completely incommensurate with the kind of chaotic anarchy with which they viewed packet switching, and yet packet switching was demonstrably one of the most effective ways of dynamically sharing capacity, which is something that Kleinholz' <ph?> analysis showed us. That aggregating higher capacities and letting people dynamically share it had statistical properties which were an improvement over dedicated circuit switching. So I think part of the problem was there. And also they didn't understand how to monetize it, and it's not even clear today how well they understand how to monetize this network environment. I would say there's lots of ways to do it. Google has been particularly successful monetizing advertising, but there are other applications that are equally solid in terms of their business models. But I think none of those business models were in anyway commensurate with the experience of the tel-co's. Right now they're trying to turn themselves into cable companies because they understand the cable television model, and I keep thinking, "Hey, that's so 20th century; you should be looking this way, not that way."

Nielson: They did use it for the signaling channel fairly early on. In other words, telephone call setup moved to a packet switched network, but they somehow couldn't get the notion that users would also benefit.

Cerf: That all the other content should be switched that way. Well, you remember what caused that. The Bluetoothers- I don't-- the Blue Boxes <ph?> were using the inband audio controls to take over the long distance network and make free telephone calls. And so the only way to combat that was to take it out of band. And when they took it out of band, they invented systems-- well, it was signaling system number 7 eventually, it was several before that, which is packet switch. But you're quite right. The content didn't flow that way, only the control.

Nielson: Right. Exactly. This goes back maybe to that question you always expect people to ask you: You have to be surprised by the depth and breadth of the Internet, don't you? I mean, I don't care whether you're controlling an astronomy experiment in a satellite or sharing resources that are terribly expensive. or other things. Plus, there is this growing use of it as a social medium. You have to be surprised, don't you?

Cerf: And I am. I tell you what surprises me. Two things. First of all, I am always stunned when I discover when I'm looking for information that it's there. I mean, sitting at the dinner table and somebody says something-- I'll tell you, I was having dinner in New York. Somebody said when was the Sony Walkman invented? And so I'm sitting at the dinner table and I, you know, get my Blackberry out, and I fire up Google, and I go and look and I find 1981. And a few hours later I get an email from somebody saying, "No, you're wrong, it was actually 1979." And we have dueling, you know, back up documents to prove one or the other. But the thing is that was a casual act. I'll give you another example of casual and stunning. I was in a hotel in Los Angeles last week, and I was on a video chat with my son, who's in Hollywood. And he says why don't we get mom on the line; she's in Washington. So we bring up a three-way video chat on the i-Chat application on the Macs. We all have the same, you know, wonderful Macs with the little television camera. So we're all three chatting away, Echo cancellation is working fine, no one's wearing headsets. And then my wife says, oh, would you like to see how the fireplace is coming? Because we're building an extra fireplace in the basement. And I said, sure. So she unplugs and she's using Wi-Fi, of course, and it's got a battery, so she just casually walks down the stairs as we're chatting away, aims the camera at the fireplace and I get to see how far the construction has gotten. And after that whole thing was over, I was thinking, I don't believe that. I mean, we used to have 17 engineers to try to figure out how to get the video set up and get the audio advance and everything else, and here are three of us, very casually, without thinking about it, had this three-way conversation, and she's wandering around with the television camera, aiming at the fireplace to tell us what's going on. And I thought, my God, have we ever come a long way from where we were in 1981 when we were trying to do video conferencing over the ARPANET and the Internet.

Nielson: It takes a gray beard to recognize that.

Cerf: That's true. It does.

Nielson: And if you talk to your son, well, your son's getting along in years, I guess, but if I talk to my grandchildren about this, it's a yawner. Just like they also can't believe that my Dad used to drive a team of horses either.

Cerf: That's in another century, and of course, the answer is yeah, you're right, it actually was in the previous century. I know, you're right, young kids today just take all this for granted. They don't see this as a big deal. For me, the fact that the damn TCP connection actually worked is amazing when you consider all the boxes it had to go through, all the pieces of software it had to go through, for all that to actually work, and I'm getting something useful out of it, is just astonishing. And for everybody else, since they don't know any of the details, it just works and what's the big deal?

Nielson: Do you ever feel just a little bit responsible?

Cerf: No --

Nielson: You could be forgiven if you do.

Cerf: Well, maybe. But, you know, I've learned a lesson in raising kids, right?

Nielson: Sure.

Cerf: What you learn is that if your kids are really successful you shouldn't take too much credit, because then when they screw up you don’t have to take too much blame, you know. You know, I played my part, but boy, tens of thousands or millions of other people have too.

Nielson: Is there anything in today's Internet, including the applications world that worries you a lot?

Cerf: Let's see, there is some opportunities not taken which we alluded to a little bit earlier. The harnessing of broadcast media to do a better job of delivering large amounts of content, identical content, to people. I want to be able to do broadcast IP, basically, and take advantage of that in a sensible way. What worries me in part is that the openness of this medium may be shut down by a commercial interest, the broadband carriers trying to figure out how to control this more to their advantage, without understanding that by not controlling it, they have a bigger opportunity to grow remunerative applications. They're too fixated on control. Governments that want to control discourse on the network, that want to limit access of information. You know, in the United States we talk about freedom of speech a lot. It's built into our Bill of Rights, and we don't talk about the freedom to hear, so to speak. And I think the two have to come together. It does you no good to speak if no one can hear you, and so if there is government sponsored censorship, that is an attempt to limit either the ability to speak or to hear. "I will refuse to let you have access to information that you want to attain." I worry about that a lot. And I think generally speaking we don't fully understand what the future of this capability is going to be. And like we said earlier, it's a software environment, and that puts very little limits on what's potentially possible. There are real abuses on the net, too, though. Spam is an example; identity theft is an example; content, which is really, you know, harmful, like child pornography, and fraud and other kinds of abuse. It's all part of the Net, and as someone recently captured this for me, when he said that once the network became available to the public all of the ills of our human nature were injected into this network as well, so we shouldn't be surprised that people abuse it, just like drinking and getting in the car and driving, which is not a good combination. People do it; we can't stop them. All we can say is, is "This isn't acceptable in our society and if we catch you, there will be consequences."

Nielson: The Internet certainly reflects the human condition or the human culture or our propensities, good and bad. Maybe the important question is whether it does so preferentially in any direction, good or bad. My guess is that it just mirrors us.

Cerf: I think it is-- what is the right word for it. "Anisotropic." Or no, it's isotropic in the sense that it is not favoring any particular application or content or anything else. So I'd prefer it to be isotropic in that sense, and therefore, we have to introduce other mechanisms than those within the network itself to deal with the social and cultural issues that arise.

Nielson: And back again to the nature of this virtual tree. If you want a particular branch to mirror some aspect of our culture, then you may have to come up with restrictions or enablements to do that.

Cerf: You know, if you want to use this tree analogy, as beautiful as-- what are the little Japanese trees that are so carefully manicured?

Nielson: Bonsai.

Cerf: Bonsai. Bonsai trees. As beautiful as they are, it's not clear to me that we would want to apply bonsai philosophy to the Internet, because cutting off branches to satisfy one person's aesthetic view is not necessarily the most beneficial way to harness everybody's knowledge and experience.

Nielson: No, I know, but that isn't what I meant--

Cerf: I know you didn't mean that, but the image of trimming off pieces of the tree, for me, immediately suggested that somebody has to decide which pieces to trim off.

Nielson: But that is exactly the prerogative of a caring parent, to take care of their child's interface so that they are trimming the tree in some respects.

Cerf: Ah. But here's where--

Nielson: But it's locally imposed--

Cerf: This gets to be a very interesting deal, analogies are always a problem. I don't want to cut the branch off necessarily as much as I might to say, "You can't go on that branch."

Nielson: That's what I mean. Control the access.

Cerf: And there we would agree. I will control your access to this part of the tree. But I think we would do well to adopt some kinds of mechanisms that would permit that kind of imposition as long as it's local.

Nielson: Well, I'm talking about a virtual trimming of it so it applies to you and to the domain that you wish to restrict; that's all I meant. So, I guess I already assume by this that while you may have your own personal feelings about practices on the Internet, you're not one to rush in to prohibit them from existing on the network?

Cerf: No, I'm not. I really think that this open environment is too valuable in almost every, by every metric I can think of to arbitrarily inhibit. But I would abide by agreements that are global in scope that are beneficial to everyone, which from the technical point of view, that's been important to me. The allocation of Internet address space has to be done in a way which preserves the network's stability and security. So you want to adopt rules for that. And you want to adopt rules that will allow for innovative, new uses of the net to come about. You wouldn't want to have a rule which says that the tel-cos can decide what applications you're allowed to put up on the net. I like this permission-less network idea, which has allowed companies like Amazon and eBay and Google and Yahoo and Skype and others, to exist, as opposed to somebody trying to decide which ones should be allowed on.

Nielson: We've already talked a bit about network neutrality and the question of private internets. Maybe it's useful to revisit it for a second. Let's say someone starts a private network with Internet like capabilities with, perhaps, gateways into the Internet. Are the risks or benefits from establishing world wide private networks?

Cerf: We've already gone through this once. I would say that the walled garden notions of AOL, for example, are clear instances of attempts to control content and access and to bring people into a place where they only see what you want them to see. And the general public has said, "We don't want that." The same problem happened with computer companies that have their own proprietary networking software, S&A at IBM and Dec-Net <ph?> at Digital, and so on. Eventually they discovered that their customers said, we want something that allows any computer to interact with any other computer. "I don't want to be trapped into one vendor's offerings." So I think there's huge pressure to keep things open. At the same time, when the Internet was being designed, I made an assumption, which turned out to be wrong, that we could keep a completely open network; every computer could interact with every other computer, and if you didn't want to talk to somebody, you didn't have to. If you insisted on authenticating at the other end, you could do that. You could resist any communications unless you'd had an authenticated handshake, and that, you know, every computer had to defend itself. Well, that wasn't practical, and so you end up with perimeter defenses in the form of firewalls in an attempt to create virtual private networks by isolating some resources that you either never want the public to have access to, or only under certain controlled conditions. So I didn't quite get that right, and I think if I were to go back and do a clean sheet design, authenticity at the edges of the net and the notion of virtual private networking or some kind of enterprise networking that deliberately wanted to isolate all but a certain portion of itself to the rest of the public net, should have been built into the architecture, and it wasn't. So we ended up with firewalls as a substitute for that. And of course there are lots and lots of internets that don't connect to the public network. They just use IP. And I don't have any problem with that. There might be a question about what address space they should use, and in fact we've oscillated on that a little bit, especially as we've started to run out of the IP Version 4 address space. But the idea that a network could be using Internet privately is perfectly okay. It's just using Internet technology privately is perfectly okay; I don’t see anything bad about that.

Nielson: Do you see a business model that would propose the building of a private internet, one that would draw people off of the Internet into this other world with other characteristics or do you think that's been preempted by the size and growth of the Internet?

Cerf: I think an attempt to do that would either have to be highly compelling because of some application that everybody wants or eventually will disintegrate, for the same reason that the walled garden notion has disintegrated. The trouble here is that a privatized network can't create-- it cannot keep up with the creativity of the open net. That's really the issue.

Nielson: Let's see. —Now that you are no longer leading as many Internet organizations, are you in danger of becoming an elder statesman?

Cerf: Well, every once in a while --

Nielson: --and would that be good or bad?

Cerf: I occasionally get the grandfather of the Internet as opposed to the father of the Internet, you know, label. Well, perhaps. I don't regret, you know, sort of sliding into that position. It was a lot of work to be chairman of ICANN, and I had a huge amount of help, too. If I hadn't had that help I couldn't do the job at all. So I don't regret spending eight years on it but I'm now happy to turn this over to Peter Dengate Thrust, who is the New Zealand barrister who's picked up the job.

Nielson: But do you have any sense that you want to be above the fray. It doesn't seem like Vint wants to be above it, he wants to be in it.

Cerf: So actually that's an interesting question, because I have tended, even in my business incarnations at MCI and even at Google, to stay a little bit out of the business itself, and to try to speak about Internet as Internet independent of any particular business that we happen to be in. And I thought that that was an important role to play, and although I'm sure I haven't done it perfectly, I think I've managed to remain capable of speaking on behalf of the Internet as a whole, and not just my company's interest in it. And for me, that's always been important. So in that sense, I'm a little above the commercial fray. But when it comes to policy issues, I want to be in there, and when it comes to new technical ideas, I want to be in there. I want to understand what new things people have come up with that you can make this system do that I never would have thought of. And it's exciting to be around people who are young and, you know, are too young to know you can't do that. So they just do it, whatever it was.

Nielson: I don't know how much time we have left but I'd like to test you a bit in this regard. You have talked about the inner planetary, intergalactic Internet or whatever. Okay, what about the microscopic nanonet? I want to see whether your eyes light up.

Cerf: Ah. Interesting idea. Interesting idea.

Nielson: I just wanted to see because it's a test, to see if your wheels are going. We have these nanomachines and nanotubes all forming very small systems. How they talk to each other and………

Cerf: I think that's a very reasonable kind of proposition. One thing I want to make sure is clear is that J. C.R. Lickliter is the one that wrote the memo in nineteen sixty something, about the intergalactic network, which was a pun reference, or sort of a tongue in cheek reference to what became the ARPANET. I've never proposed to do an intergalactic net because the time delays are too big to contemplate.

Nielson: I'm sorry about that.

Cerf: That's okay. Some people say that, and I say, no, you have to think about this. It's hard enough to do interplanetary. I do want to-- I'd love to be able to take a step towards one interstellar mission. The problem is I don't have any idea what kind of a power source would be sufficient to deliver a signal all the way back to, you know, four or five light years away. But it would be fun to think about that. The nanonet idea is really fascinating. I'm not quite sure how much we can cram into these little things, but when we start talking about molecular computers, it seems perfectly possible to create something as complex as today's personal computer in the small. Then we get into this question of well, what about the signal levels and you know, is it a mesh network, and do you have to find another nano friend over here in order to propagate the stuff? I think it would be really interesting.

Nielson: You've bitten already.

Cerf: Oh, of course. This is cool. This is a space that nobody's been in yet.

Nielson: In regard to this interplanetary network; what comes to my mind is that during the time of signal transit, things that you thought were true could have changed.

Cerf: Yes, and some people really don't get it. We've had these arguments with people who think you can do IP everywhere. And it's actually true that you can use the packet format and, you know, do it everywhere. The problem is that the assumptions that get built in, about delays and so on, the dynamic range is so big that to make it work consistently, you would have to impose horrible delay time-outs on the loadlay <ph?> terrestrial environment or the loadlay planetary environment, to say nothing about being worried about running out of IP addresses again. And so the interplanetary design actually says every planet has its own IP address base, and we don't use IP addresses as the identifiers of the end points of an interplanetary communication. Instead we use a higher level construct, which is an end point identifier.

Nielson: Right. Well, I think we're about out of this tape. Let me ask one more far-out question. When Version 6 comes into existence, will you, Vint Cerf, have a unique IP address?

Cerf: I will have one, but it won't be dedicated to me. In other words, IP addresses are just the definitions, even IPV6, are the definition of an end point access.

Nielson: That's the way it is now. Maybe the way I should say this is: Will you have a name that's recognizable anywhere on the network and that can be bound to any IP address that you approach?

Cerf: Okay. Will I have an identity? Well, I already have a domain name, several of them actually. But one of them's "cerf's up dot com." So I can easily imagine taking that particular domain name and binding it to any IP address that happens to be appropriate. But I'm also fascinated by the possibility of other identifier spaces being bound dynamically to IP addresses. We do this all the time in Instant Messaging. You have an Instant Messaging handle, and it eventually binds to a particular IP address of the moment. And there's no reason why IPV6 couldn't do that as well. Some people have this model, I know you don't, but others think that somehow IP addresses are fixed identifiers that you'd be assigned on birth and never change. That's not what IPV6 is right now. It really is just the termination point of the physical network.

Nielson: But the reason I mention Version 6 is because the address base gets huge--

Cerf: Huge, yeah.

Nielson: And now can I as an entity have some identifier associated with me? Maybe an IP address is not the right label, because we associate that with network places. But somehow, an identifier that will enable anyone to have, wherever they go in the world or in space, an instant recognition of their location, be that logical or physical?

Cerf: Oh, of your location?

Nielson: Well, and I know there are some privacy issues associated with it. My location can be expressed, either from GPS or a network connectivity point of view? Should these be bound together? This is just one question of where we are headed when we have a much richer addressing environment?

Cerf: Okay. So let me try to respond here. First of all, IP address, both V6 and V4, were not designed to be geographically identifiable. They were intended to just be topologically significant. But in the meantime, there were allocations of IP address space to various networks and those networks are known to be in certain locations. And so it's now very common for somebody to do a look up of an IP address to figure out where it likely is, where that termination point is. And then people, including Google people, applications, try to figure out well, what should I present of what I know and Google will pull up a French Web page if it thinks your IP address is in France. And that's beneficial to somebody who happens to be there, because they speak French as opposed to you're there because you're traveling and you don't speak French, and it didn't help a bit. I think that there are ideas about identifiers that are very worth pursuing. Bob Kahn <ph?> and I did some work, some years back, on what he now calls the handle system. And what's interesting about it is that it's an identifier that's unique but it is not necessarily bound to any particular kind of object. It could be bound to a virtual object like a book or a piece of music. It's a digital object of some kind. And it could be anywhere in the network. The problem is you have to find it and you need to do a look-up, just like you do in the domain name system. You look up the object identifier handle and you go to a system which keeps track of where those objects are. If the object is moving around, and the identifier that comes back tells you something about where it actually is, then there is an issue about privacy. So right now I don't think the handle system has that problem, but this whole question of how do you know where people are and how do you protect against that and how can you interpret the identifiers that are termination points in the net, creates a tension, and the only way I can see to resolve that is to use various encapsulation methods so that-- the ones that are common in virtual private networking cause you to get an identifier which is not the same as is associated with where you physically are because you tunneled through the net. There are lots of ways, I think, of protecting people's location privacy. So I'm not exercised or worried about that protection, because I believe it's easily accomplished.

Nielson: Okay. The rest that I have are more of a personal nature. So let's stop here.

END OF INTERVIEW

