LOS ANGELES – A Conversation with Steve Crocker and Leonard Kleinrock
Monday, October 13, 2014 – 18:30 to 19:30
ICANN – Los Angeles, USA

UNIDENTIFIED MALE: A conversation with Steve Crocker and Leonard Kleinrock, October 13, 2014 in the Los Angeles ballroom.

JIM TRENGROVE: Good evening. Thank you for being here. It's been a long day, but this ought to be the fun part of the day if you haven't had enough already.

My name is Jim Trengrove. I am Senior Director of Communications at ICANN and it's a privilege to have these two gentlemen here on stage. I think it's going to be a fascinating hour.

This morning Secretary of Commerce Pritzker referred to the book by Walter Isaacson, The Innovators. She talked about collaborative creativity and these two gentlemen here are certainly examples of that. I'm reading it and especially the part about the ARPANET, it struck me and again, when Steve in his opening session this morning talked about the importance of personal relationships in creating this technical marvel we've all come to rely on.

Personal relationships – so Larry Roberts who was the project manager at the Advance Research Project agency for the Defense Department at the time decided that the first node of the ARPANET should go at UCLA. It struck me, well, the reason for that was because UCLA was the best suited for that, or that Leonard Kleinrock was his office mate at the MIT Lincoln Labs earlier than that and that two of the people who got tapped to be on Dr. Kleinrock's team were Vint Cerf and Steve Crocker, who incidentally, had gone to high school together at in Van Nuys over in the San Fernando Valley.

So with that, I'm just going to introduce them – Dr. Leonard Kleinrock, Dr. Steve Crocker – and get out of their way. We're going to take some questions. They're going to speak for a while. If you have some questions, we'll have a microphone here to get going.

I'd just like to start off with one question, and for both of you, Dr. Kleinrock, if you could tell us because probably more than anyone here you've known Steve the longest; and Steve, question for you as well. Dr. Kleinrock, who was Steve Crocker back in October 1969. And Dr. Crocker who was Leonard Kleinrock at the time?

Thank you.

LEONARD KLEINROCK: Thanks. Well, I'll take that on. Who was Steve Crocker?

Steve Crocker was this brilliant, aggressive, unruly, but enormously creative leader of my software team at UCLA. You didn't have to spend much time with him, because he'd go away and do everything that needed to get done and come back with a great result and take on more.

It's amazing the way in which the authority of this entire project we're talking about was delegated down the hierarchy from ARPA's thinking to the principle investigators goals, to the real workers of the group, which were the graduate students who listened a little bit to what we said – not much – and off they ran and created the thing we now call the internet. Steve, that's a fair characterization I think.

STEVE CROCKER: I'm trying to recall back in this business of listening a little bit, I don't recall that you said anything. So maybe…

LEONARD KLEINROCK: See what I mean?

STEVE CROCKER: The circumstances that led to the development of the ARPANET and set the direction that persisted throughout the Internet really didn't seem as unusual at the time as they do in retrospect, but it was a very unusual arrangement.

There had been some unsuccessful networking attempts. There had been some successful networking attempts and some several unsuccessful networking attempts, usually on small scale in both directions. But in the ARPA office in Washington and in the senior community of the computer scientists who were the principle investigators of the major ARPA computer science projects, there was a consensus that it was time to go push forward and try to connect up computers and to do so in a somewhat ambitious way, to do so in a way that was independent of the technologies of any of the particular computers.

Not IBM computers. IBM was the dominant computer vendor of the day. Not UNIVAC computers, not digital equipment as one vendor, but something that was neutral across all of them.

That made the project hard in a way and it also was a beautiful separation of interests and could focus on the communications aspect on one side and the use of it in another aspect.

So, a formal project in the usual way, a carefully constructed RFP, and a selection of a competent group at Bolt, Beranek and Newman to design what became the first routers, the procurement of long-distance lines from AT&T to connect these together, and then all of this was given as a present, in a way – one of the kinds of presents that you can't refuse – to the existing research places of which UCLA under Len's direction was one of them.

And then there was this absence, this unbelievable vacuum of direction as to what to do with it. Instead, it was a situation in which each of these research environments were populated with people who were deeply involved in computer science research, and the research varied over the landscape from artificial intelligence, to multi-computers, super computers of the day, to graphics, to big databases, to fancy interactions, all sorts of things.

But the common theme was that we were all pressing the barrier of trying to develop new research and, to a certain extent, shared all of the basic knowledge about what computer science technology was of the day.

Now, Len was probably a bit of an exception in that most of the other research projects were focused on things that were not really closely related to the ARPANET itself. Len's research had been in queueing theory and development of theory for packet switching and was quite eager to – and even better if he describes all this – but at UCLA there was a very strong anticipation of this technology coming.

But even at UCLA, there was a phenomenon which I observed almost uniformly across the landscape. When the network – the piece parts, if you will, the IMPs which were the routers that arrived, and there was a lot of heavy lifting to do – the very simple thing of connecting a computer together.

You couldn't buy any parts anywhere. You had to fashion them from scratch and design them. There wasn't any software. You had to go figure out how to make incisions into the operating systems.

That was not the kind of work that attracted serious research people who had careers and status like Len and counterparts around the country, and so it was delegated off to the graduate students, which were generally called the second lowest form of life in the university. But we were eager to get involved.

There was a little bit of a separation layer there, which Len translates as we went off and did it without needing much help, and from the other side, it could just come natural. I don't take any umbrage by it, but it just observed that the main stream of the funded research took place at each of these places and the ARPANET was an extra piece that was bolted on, and so it was a secondary project virtually everywhere. There were one or two exceptions.

In the case of UCLA, it was a primary activity because Len and a number of his graduate students focused very heavily on the internal dynamics, the queueing models and the performance of the network itself, and at SRI, there was a network information center that was formed. But at most of the places, it was not the main stream activity. It was just an added thing. And so it attracted a couple of graduate students typically or secondary staff people to get involved. We all formed a loose bond and cooperated and did the best we could trying to figure out what to do with all of this.

Something I always try to say is that we were very, very conscious of how little we understood about where this was going to go, and on that basis, we wanted to leave a lot of room. We did not want to design and build and deliver a finished product and say, "These are the only services there are and it's all put together." So we designed very thin layers of protocols with the explicit understanding that people would build on top of those layers, people would blow those layers away, and put something else on the side or they would slip things in between. We wanted to do some useful things, and then as I said, leave a lot of room there.

That's probably enough. I give you another shot here.

LEONARD KLEINROCK: Well, Steve referred to the network and research we were doing. It was not so much only the underlying infrastructure, but the work we had done understanding the way a network should perform to try to get at the fundamentals.

What were the key principles that would make this technology perform well and per scale and all the rest? The notion of large shared systems, of distributed control, but mostly delegating authority throughout the network both technically and administratively was very important, so we focused on deciding what happens when you scale this thing up, what happens when you change the nature of the routing procedures.

The group that I was spending a lot of time supervising, as opposed to these guys, were the students looking at the underlying performance evaluation, the design, the optimization, the principles, and the ways in which this network would scale.

You know what's interesting regarding the networks that didn't happen? In 1964, Ivan Sutherland, the second IPTO Director at ARPA, the computer science director group, came to UCLA – he was a classmate of mine at MIT along with Larry Roberts – and he said, "You know, there are three almost identical IBM mainframes on campus. Let's network them."

Simple, straight forward problem. Not the heterogeneous network that eventually became the thrust. He came on campus and he found one in the medical school, one in the campus network, another in the business school, I believe, and said, "Let's connect them," and it never happened.

It had nothing to do with technology issues at all. It was all about political infighting and jealousies at the administrative level. It's interesting that that obviously simple problem didn't get going, because it was administrative failure between administrative groups that had no authority over each other.

When DARPA came in later and said to all of the PIs around the country, "Let's put you in a network so we don't have to give you all the resources. You want to use a resource, you want to use graphics? Log on to a machine at Utah. You want to use a database? Log on to a machine at SRI."

Almost uniformly, every one of them said, "We want nothing to do with a network. How can I possibly take my mainframe, all my time-shared system, and put it on a network and you're going to steal cycles from me? It's loaded now 100% of the time." And so they said no.

ARPA said, "Well, we're funding you guys. You shall join this network."

They caved in and they did. Almost immediately, they were thrilled that this thing was working.

But part of the key to allowing that to happen in a seamless fashion was due to one of the people help designing [inaudible] Wes Clark who said, "Don't load the individual sites and those mainframes with the job of worrying about the interface to all the other sites. Don't make them worry about the communication substrate. Put that on the device that Steve referred to before – the IMPs, the routers. Let them handle all of the packet switching, all of the segmentation, all of the error control, all of the routing control. Take it off the burden of the mainframes."

That was critical because we really would have had a problem forcing people to make major surgery on their operating systems to install a network. That was a rather important part.

So there was an administrative and a political as well as a theoretical push to make this happen and it happened across the community in a rather interesting way. The personalities and the political pressures were really important as well as the technical insights that came about.

STEVE CROCKER: I don't know whether you were totally aware of it, but I actually worked quite hard on that failed UCLA project.

LEONARD KLEINROCK: Didn't know that.

STEVE CROCKER: I was deeply embedded in that then, so I can tell you from the bottom what that looked like. There was a group of nine extremely good computer scientists and me up at the business school end.

For those of you unfamiliar with UCLA, there's a north-south axis, and in those days in the mid-1960s, UCLA was endowed with three big IBM 7094 scientific computers. A 7094 was the flagship computer throughout the country. If you had one of those, you were a first-class place.

UCLA had three, which was extraordinary, and they were in three different places. There was one in the south end of campus in the medical environment. They, biomedical computing, funded out of the National Institutes of Health. There was one at the north end that was attached to the business school that IBM used as its Western Regional Data Processing Center (WDPC), and they provided that one. Then there was one in the middle of the campus to serve the basic computing needs of the campus.

The idea of connecting these three machines, as Len described, was the impulse and some money was provided, but the group that actually was formed to do this was formed all out of the business school. Instant jealousy and rejection by the other two. Inside of Western Data Processing Center, there was a recognition about the political issue and so there was a political line that was drawn.

People who were working on the project were curtained off so that they were not part of WDPC except in name, but we knew we were working on a campus-wide project. Didn't help the politics with the rest, but it was an attempt.

I got drawn in to it and worked on it from about mid-1965 through '66. We were working away furiously. I can recall vividly the back and forth about this, but the project got shut down and the residue of the funds got moved in to the engineering school. The head of WDPC suggested rather pointedly – I was an undergraduate and this was my livelihood – that I might go down to the engineering school and follow the money, which I did and the rest, from my point of view, is history.

It was quite clear that there was technical things that we could solve, and then there were some other problems.

LEONARD KLEINROCK: But the result of that failure was to implant deeply in the ARPA mentality that a network was a thing that they should consider, because the ideas were there earlier from Licklider. But when Ivan came in, the second director, that idea was now fertilized in the group and eventually, it bubbled up in to what became the ARPANET project and that came to fruition of course.

STEVE CROCKER: A lot of other things come flashing back.

LEONARD KLEINROCK: Let's talk about the arrival of the IMP and the connection there.

STEVE CROCKER: Oh yeah. So there was this extraordinary mismatch. We had a highly professional group at Bolt, Beranek, and Newman. Sophisticated engineers used to working on a schedule, very disciplined and organized, and they had won the competition for the contract. The contract, I believe, began on the 1st of January, 1969. They said they would deliver the first IMP eight months later on the ninth month of the beginning of September and they would deliver one a month after that.

Meanwhile, a rag-tag bunch of programmers tried to figure out what we were going to do. One of the challenges, as I said, was building a hardware interface and another was figuring out what software to build. It won't be a big surprise to say we were running a little late on the software. We were in touch with the group at BBN, and we heard in the summer that they had some relatively subtle hardware problem, timing bug that they were tracking down.

I said to myself, "Oh, that's good. They'll run late on this."

Well, they didn't want to run late. They tracked down their timing problem. They fixed it. I figured that between the shipping of the IMP, sending in the Honeywell engineers to check out the box because it was built on a Honeywell platform, and then sending in the hardware guys from BBN to check them out, and then sending in the software guys to make sure the software would run, that I had a couple of weeks in there. The 1st of September was on Monday (Labor Day), so I figured that was an additional benefit from that.

They put the machine on the airplane on Saturday because they didn't want it to be late. It arrived on Saturday the same day. It got wheeled from the airplane and moved over to UCLA on Saturday. It got moved into the computer room on Saturday, got plugged in on Saturday, and it was up and running on Saturday. It was still two days to go and my two weeks of latitude had turned in to minus two days.

LEONARD KLEINROCK: Steve always has a flair about him. I'm going to tell a story about you. Steve, you probably forgot.

Steve managed to get this sailboat, if you want to call it that. It was a double hull, as I recall, a CAD of some sort. A big thing. And a bunch of us were out there off Santa Monica Bay. We got on this thing, and it's sailing like mad.

It starts heading to shore. The whole beach separated as we're crashing on to the sand. We crash on to the shore and we get there and I said, "Steve, this is not quite seamanship."

Steve's reply was, "It's all about the theater."

STEVE CROCKER: That's high praise coming from you. You're an accomplished showman.

LEONARD KLEINROCK: The build-up to the arrival was really very important and there were all these parties taking place. One of the key groups was Bolt, Beranek, and Newman, as Steve said. We were, meanwhile, commissioned to become the Network Measurement Center.

UCLA's job was to try to break the network that BBN was trying to put together. So immediately, it was not a marriage made in heaven. Realistically, our job was to stress the outer limits of this network as it came in, and so we planned to make a number of serious tests and we did.

I'll tell you some of the stories afterwards, but the antagonism between BBN's goal, hard driven, organized, commercial deadline oriented, get the job done versus the academic "We're going to make sure this thing works. We got a bunch of graduate students trying to help us put it together at our end." It was an anathema and it set the stage for some rather interesting dynamics that took place later.

But in spite of all of that, it was a tremendous success. That machine came in and it worked. As soon as you plugged it in, the operating system picked up and continued to run anywhere they shut it down, it [inaudible] on the way over.

STEVE CROCKER: You've touched on a topic that relates to a question that's very much been a resurgence of why was the network built? I think we have an opportunity here. This is completely unrehearsed and I hope that we can have an interesting dialogue and neither of us walks off the stage here.

But there's been fresh questions asked about what was the genesis? What was the reason the Defense Department built the ARPANET?

I've been listening. I was there, but I wasn't the principle person making the decisions. I worked for a number of them and knew them. I've been listening for a past few years and the stories are piecemeal, fit together, but they have quite a bit of spread.

So if you don't mind, let's dig in to that a little bit because it's not just an interesting story, it's an interesting collection of stories and even a bit of a lesson about what history looks like in reverse here.

LEONARD KLEINROCK: Sure. As a PI, receiving funding from ARPA at the time – now it's DARPA – my motivation was purely scientific, purely research, and never, ever was the notion of defense orientation application brought to any sense of recognition or observation on my part.

As a PI, we got no pressure from ARPA at all in that respect. Whatever they were thinking in talking to Congress and the Department of Defense and the Pentagon, we were not privy to. But the idea was to build a network to allow each of the sites that was being supported by ARPA to share their resources.

They would come to a new PI and say, "You're a great person, here's some money, go do some research."

And the response of the new PI would be, "Fine. Buy me a computer."

ARPA said, "Fine. We're happy to buy you a computer."

And then the same PI would say, "But I noticed that in Utah they have great graphics, and at SRI they have this terrific database, and at Illinois there's high-performance computing. I want all of that."

And ARPA said, "No. We can't afford to give everybody all those resources. However, if you want a network, as I said before, you can log on to Utah to do graphics. Log on to SRI to do some database access."

The motivation was, in a phrase, to share resources, to share the hardware, the software, the services, the applications among the other sites. That was the motivating factor from our point of view.

Whatever was told to Pentagon and touched Lacklider, now however, as Morten Bay will attest, one of the people that's going around the country with me, trying to recollect what was going on back then, we interviewed a number of people who were not only at the PI level but at the office directors at ARPA and the ARPA directors themselves. We're getting a very mixed story as to what that motivation was.

Just as you said, Steve, there were different perceptions, different recognitions, and this is over different periods of time. How much of the defense drive was pushing development of this ARPANET?

It is a mixed story. It'll never be a complete story. It'll never be a single answer. It's going to be who's view was expressed.

STEVE CROCKER: Yes. So, Steve Lukasik whom I had the privilege of working under, was the Director of DARPA for a several year period and wrote the checks that supported a lot of this research and was extremely enthusiastic. He's written a retrospective of talking about, in essence, answering the question why I wrote that, why I supported it, why I funded this effort. There's a very strong defense orientation there.

Personally, I think there's a little bit of skew involved, but I also think that there's some important subtleties that aren't generally understood across the world. The natural assumption, I think basically guessed at but taken as Gospel, is that because this is a packet switch system because it's distributed because it has natural properties of surviving individual pieces going out, that it was built specifically to survive nuclear holocaust.

Len's shaking his head, and I will shake my head, too. There are a couple of pieces that are worth talking about. Len talked about stressing the network, about trying to find what the limits are performance. By happenstance, I was responsible for finding a quite specific stress point on the network. I don't know how many of you are interested, but I'll spend a few seconds on the technical aspects.

The technology, as we've said, is packet switching, and so when one machine wants to send a message over to some other machine, it divides it up in to smaller pieces. Those pieces go in and each one of them is sent and routed individually. They may go to the same place or they may go around in other places.

BBN built this, but they were also a little concerned about overload and so they said, "We will only send one message at a time from one place to another that has the same link on it." They supplied a number of link numbers.

The idea was that if you wanted to send a series of packets, you would send one, wait for the answer to come back that it had gotten there, and then you'd send another one and so forth.

I went off to work at the DARPA office. I was enthralled with having a front row seat on all the research that was being funded. Ten days after I was there, I was sent down to Oklahoma to Tinker Air Force Base because there was an experiment where the air Force was going to see how the ARPANET performed compared to the existing AUTODIN I system.

I went down and I talked to a couple of Air Force captains who were really computer science geeks. They were doing the programming for the Air Force computers. We talked about how we were going to get high through-put through the system.

I said, "You know, you can only send one message at a time, but you could divide up and use multiple links sort of like a Gatling gun. You could spray on links."

I gave them that advice and did not pay a lot of attention, when off in the side a few months later, there was some discussion about the network could be locked up if there were too many packets in there.

Several months later – we're talking about from summer time when I first went there until February the next year – they turned this on. Four seconds later the entire ARPANET came down across the country. Packets flooded, everything stopped, and every IMP died. There was no automatic restart mechanism. It was kind of ugly. So the Bolt, Beranek, and Newman folks just restarted their IMP and then they called up the next place and they said, "Please reload the paper tape and start that one and call the next one and so forth."

The two guys down in Oklahoma were scratching their head. "I don't know what happened. Let's try it again." Four seconds later, the entire network came down.

After it was all untangled, the folks at BBN were quite angry, quite hot. I happened to have been in a meeting discussing research things at MIT, completely unconnected to any of this, and I get a message that I've got an urgent phone call. I'm a program manager working for the

Department of Defense, I get an urgent phone call, I go out and take the call.

"Why did you do that?"

I said, "Who is this?"

"Why did you? You did it on purpose."

I said, "Oh, Hi Severo. What's up?" This was Severo Ornstein at Bolt, Beranek, and Newman.

My first task obviously was to get him to calm down long enough so I could understand what he was saying.

He said, "You smiled when you told them to do that."

I said, "Yeah, I thought it was a cool idea."

He said, "You did that. You broke my network."

I said, "No. Look Severo, it's not your network. It's really more my network than yours and I wouldn't actually do that on purpose." It was an interesting conversation for all.

But the relevance back to this big question is if this had been a deliberate project to build a nuclear hardened thing and there had been a flaw in it that permitted a casual error like that to bring it down, I believe heads would have rolled. That would be a totally unacceptable in a hard-core military oriented network. In a research network, different matter.

I've always thought that there had to be quite a bit of distance between the project as structured as ARPANET versus trying to build the holocaust oriented networks.

One more thing. There are plenty of other reasons – good, solid military reasons – why you want to connect computers together. Military has hundreds of computers, thousands of computers, they have lots of data that has to move back and forth. There's all kinds of coordination ranging from weapon system stuff all the way back to whether or not you've got enough blankets and other provisions and moving logistics around or just the payroll or personnel stuff, so there's enormous reasons why it's valuable to connect computers together and why there's good justification for the military to do it. The very last one of which is the survival in the extreme.

LEONARD KLEINROCK: So picking up, you generate a few comments. First of all, that idea of sending one message and requesting the next message, there was a control message which made that request and its name was RFNM. RFNM means Request for Next Message and it really slowed down things in our intent.

The second comment is the BBN story again. If BBN had their choice, a lot of heads would have rolled. Yours and mine were on the dock ready to be clipped at any time.

Again the mentality, in their minds we did break their network and we did it willingly every so often because we were testing it. We tell them, "Look. This happened." They hadn't yet released the software through us, so we couldn't tell them how to fix it. We said, "Would you please fix it," and it would take them six months to fix it.

Not too long thereafter, they were forced to release the software to us, so when it broke, we could see how to fix it. We said, "It broke. Here's how to fix it." It still took ix months. This was a very structured organization that took their time, was very careful about things, but were very punitive in their response when things happened.

But getting back to the military, in those early days we, the PIs, didn't have any sense there was a military drive for this thing.

However, in the early 70s, we started looking at packet radio – the Aloha network, the CSMA, the various mobile networks etc. It was clear to us that the major application driving that thing was military deployment of soldiers in the field and tanks etc.

Not that that was a driver for the way we did the work, but in our sense, we could see the applications there. They didn't have to tell us what it was. Our research wasn't driven that way, but the natural problems came out of immediately deployed, ad-hoc deployed networks, and it was a great amount of research, which was wonderful research, took place in the 70s.

We settled a lot of the issues and we couldn't build a damn thing. The reason is the radios that they needed were 25 pounds, a cubic foot, and 25 watts. We had to wait for things like we all have in our pockets today where they weigh ounces, far more power preserving, and far more portable. But the technology then did have a military taste to it.

STEVE CROCKER: Yeah.

LEONARD KLEINROCK: The other thing is you talked about these networks being robust. Now there's two ways to get robust. One is to design redundancy in the network to protect against failure of one part. The other way to get robustness by accident is to say if these networks are going to scale, you can't put a lot of control and function in any single node. You've got to distribute the functionality and control and thereby, you do automatically get robust behavior.

What led us to that was very different than saying we needed the robustness. It came about because of the need to scale and therefore, the nature of our design.

STEVE CROCKER: I have a sense that Jim wants to…

JIM TRENGROVE: Well, I want to give – if you have a question, here's your chance to ask a question. Come up to the microphone here.

Dr. Kleinrock, I'd like to ask your version of the story of the sending of the first message and the abruptness of the end of that mission after the first two letters. What was your understanding of it? And did you understand it as a success or a failure or let's just keep trying?

LEONARD KLEINROCK: Well, there were two events. The first event was when the IMP arrived over the Labor Day weekend in 1969. The Tuesday following the Monday of Labor day, we connected the IMP to the host over a 15-foot cable in the room.

That was an event where people were present. ARPA was there. UCLA was there. The engineering school was there. We had BBN there. We had GTE. We had ATT long lines. We had Honeywell. We had the folks from SDS that manufactured our machine. Everybody was there and everybody's ready to point the finger to the other guy if it didn't work. And sure enough, the bits flowed and we were all very happy.

The event you're talking about took place after the second IMP arrived up at SRI in October. We had done ours in September. And with first connecting high-speed line of the internet at a blazing 50 kilobits per second.

STEVE CROCKER: Not 56.

LEONARD KLEINROCK: Not 56, it was 50. Exactly, and people make that mistake often.

Sri connected – 400 miles up north – connected their IMP to their host, so now we had two hosts. Now we had a network. One node a network does not make. Two nodes, it does.

The decision was to simply test it by doing what the network was designed to do, namely allow remote access across a network to a time-sharing system that somebody else had. So you sit at a terminal I machine, use the network to allow you to log in to that remote machine – in this case, one hop away – and see if you were a local terminal.

The host did not necessarily have to know that this user coming in was coming over a thing called a network. It looked like another port coming in from a local terminal. So what you have to do is log in.

We had Charley Kline at our end, we had Bill Duvall up at SRI. And just to make sure this thing worked, they had a telephone connection. Now the irony here is just dripping. We were using a telephone to prove out packets which is about to displace the telephone network, so they could communicate. The point is Charley typed the L, and he asked, "Will you get the L?" Came back, got the L.

Charley typed the o. "Get the O?"

"Got the O."

Charley typed the G. What happened? Crash.

As most of you know the story, the very first message on the net was "LO" as in lo' and behold.

Now, we weren't as smart as Samuel Moss and Alexander Graham Bell and Neil Armstrong. Those guys had good messages. They understood PR and public media. We didn't even have a camera. There's a small written record of what happened there. But the message we did end up with has to be the shortest, most succinct, most prophetic powerful message you can imagine. By accident. That was October 29th, 1969.

STEVE CROCKER: That's cute.

JIM TRENGROVE: Steve, you mentioned, and in Walter Isaacson's book, you said you were concerned that somebody from the east was going to come out at some point and either shut all this down or raise hell.

STEVE CROCKER: As Len described, there was initial resistance in the computer science community that DARPA was supporting and many of the places didn't want the network to intrude on their fine facilities. So the network as it worked out was grown from the west coast to the east coast – UCLA, Santa Barbara, SRI, and the University of Utah.

As a consequence, those of us who were in those facilities started to talk with each other and say, "Okay, what are we going to do with this? How are we going to design protocols and so forth?"

At least speaking just for myself, I found it all a little peculiar that we had no formal instruction. We didn't have any guidance. There was nobody laying out what the plan was. And we were self-propelled and I think we were having some productive conversations, but I kept expecting – or at least fearing – that somebody would show up, somebody wearing the label of adult would show up from the East. I didn't know whether it'd be from Washington or from Boston, but I thought from one of those places somebody would show up and say, "I don't know what you kids are doing, but here's what the plan is. Why are you proceeding on your own?" Nobody did show up as it turns out, so we continued on our own.

After several months – the first meetings took place in August of 1968 in anticipation of this forthcoming network, and then over the course of the next several months we visited each other's laboratories and we started to sketch out broad ideas. At first, there wasn't enough specifics about what the details were going to be in the network – we knew the broad outlines – so we could concentrate on the higher-level thoughts.

In fact, we said, "These are 50 kilobit lines, which may seem fast, but they're slow compared to the computers. Maybe one of the first things we want to do when you start an interaction session is download a special program that would take care of some of the interaction and you'd run that locally, and that would speed things up."

So that was about a 25-year advance anticipation of Java and ActiveX and now Java Script and so forth. That was kind of cool. We took a few swipes at trying to design such language, but it never materialized. That was the thought, and out of that fearfulness about seeming to be presumptuous and not wanting to do that.

We had assigned ourselves after in the spring of 69, I think a meeting in March, we said, "Look. We've been dealing with a number of these ideas, it's time to write them down," and we assigned them to ourselves. You write about that. You write about that. You write about that.

I took one on host software, and then I also volunteered for what seemed to me a trivial clerical task of I said I'd organize them. That caused me quite a bit of pain over the next few weeks because every time I started to write down what I thought needed to be written down about organizing these notes, I found myself balking. I had trepidation and also felt a lot of pressure to do it.

In the middle of the night one night, I forced myself to get past that and I said, "Oh, I know what to do. We'll write down the rules as simple possible." Write anything you want. It doesn't have to be complete. You can write questions without answers, a design without an implementation etc. All you have to do is put your name on it, a title, a date, your institution and so forth.

I thought, "We want to number these so that we keep track of them." I said, "I'll give you a number as fast as you write it, but you have to write it before you get the number," because I didn't want a lot of holes in the series.

And then in one of these quirky moments, I said, "You know, we'll just slap a label on it that insists that these are unauthoritative. We'll call them request for comments – every one of them, no matter what was in it – to signal that we invite response to that." That was the birth of request for comments or RFC series.

I thought it was a temporary hack that would last a few months, that eventually, we'd have some formal documentation. When I was asked to write the introduction to RFC 1,000, I had this feeling that we had started a Sorcerer's Apprentice kind of thing and we couldn't turn it off. Some few years later, I was contacted by the Oxford English Dictionary for the authoritative source for RFC, which is now in the OED, one of the most peculiar and unexpected twists and turns in life.

JIM TRENGROVE: Dr. Kleinrock, we have a question here. Go ahead.

UNIDENTIFIED MALE: When you designed the internet, the internet had to be designed in such a way that there was a lot of dependence on telephone lines. Today with all the advances in technology, do you visualize a possibility of an internet which is completely independent of telephone lines in every form?

STEVE CROCKER: Well, you have to use something. You need to get those bits there somehow. Your choices, I think, are basically either wires or fiber or something along those lines or a radio, and you've got a multitude of those options. There have been designs to work with carrier pigeons and other media, but at very low bandwidth.

LEONARD KLEINROCK: The network right now does live on these various forms of media. Certainly, you know, light very important, wireless very important, and cable. So I'm not sure what you're driving at.

Those early lines were noisy. Telephone lines had terrible burst noise characteristics. Therefore, we put a very simple solution in – some error detection in the form of a CRC, cyclical redundancy check code. At least you know when an error is made and you retransmit it. And that was the simple solution that we adopted at the time. It was very effective and it's still used in many cases. It's much harder to correct errors than it is to detect them.

JIM TRENGROVE: Dr. Kleinrock, it's been a few years—

LEONARD KLEINROCK: Len is the name.

JIM TENGROVE: All Right, Len. It's been a few years since the two of you worked together at the lab in UCLA, you've had plenty of time to reflect on the importance of the work that you guys did. How do you view it? Does it evolve differently as you give it thoughts [over] different times?

LEONARD KLEINROCK: Well, one thing I could say is I'll quote Steve. "It all rolled out exactly as we planned with tongue and cheek every time." Of course, there've been a lot of surprises. Each of us had our own vision as to what this thing might become.

My particular vision was, in fact, printed in a news release in July of '69, two months before the IMP arrived. It talked about having always on, always available access, anybody with any device could get on it at any time, and it would be invisible relating the analogy to electricity. Electricity is beautifully invisible. It's a plug in the wall. You want power? Plug it in, you're going to get it. You don't care how it's generated.

I wish and pray that the internet were like that today. It's far too complicated to use. These standards are all over the place, the human computer interface is still severely lacking.

So did it do what we expected? In some ways yes. In many ways no.

But the thing that surprised me and I think many of us was the dark side of the internet. We could spend many, many hours talking about all those aspects, but for those early days, what we haven't talked about is the culture that was growing up around this community of people putting this network together.

It was one of openness, sharing, trust, creativity. We knew everybody on the early ARPANET. We trusted them and people behaved well for 20 years until 1988 when the first worm appeared and we said, "Ouch." We said, "Oh, well. That's an aberration." Big mistake.

Six years later, 1994, the first broad-based spam message by Canter and [Fitzgerald] spread all over the network. And that time we said, "Ouch. And this is really a problem." If you recall what we did, we started sending e-mail back to those green card lawyers who were advertising their services on the network. We sent mail and said, "You shouldn't do this. Shame on you. Stop! Shut down."

We sent so many e-mails back to their server, we took their server down. So the unintended consequence of the first spam message was the first denial of service attack. But the point is this dark side has come up.

Look, what was the Internet at that time? What is it now? You sit in your basement, short, poor, dirty, a hovel with banana skins and Coke bottles all over the place, and you can reach out to millions of people immediately with no effort in money or time anonymously. Well, that's a perfect formula for the dark side of the internet. It's the source of many of the problems we have today – serious, serious problems.

I think the better question is what should we have, could we have done back then to ameliorate this? And that's a difficult one. In my sense, we should have put in strong user and strong file authentication and then turned it off immediately to let people join the network without any impediments. And as some of these problems began to arise, we could then have slowly cranked it up. But that's not the entire answer. This is a serious issue, the dark side. And that's part of the disappointment, but certainly is just one aspect of the Internet.

JIM TRENGROVE: Steve, Len, thank you very much. I appreciate it. It's been a great hour. Thank you.

[applause]

STEVE CROCKER: Thank you. Thank you all.

LEONARD KLEINROCK: Thank you very much. You humble us.

[END OF TRANSCRIPTION]