Interview of Jonathan “Jon” Postel

Interviewed by:
James Pelkey

Recorded: February 18, 1988
Santa Monica, CA

CHM Reference number: X5671.2010

James Pelkey: When and how did you get involved with ARPA originally?

Jon Postel: Well, the ARPA -- I was a graduate student at UCLA, and -- when was that, in '67 '68, and I was working on a project --got to work on the project under Leonard Kleinrock, and Gerry Estrin. We had a project from ARPA to do some measurement of computer performance. So we were beginning to build some system support tools and programs to aide in doing -- basically, we were going to have one computer sort of probe into another computer and do some measurements of computer performance.

Pelkey: And what time frame was this?

Postel: This was like '67, and about that time ARPA was getting ideas for the Arpanet. Now, I didn't know anything about it. And they'd gone out to -- with an RFP to, you know, buy these IMPs. They turned out to be IMPs. And about the time that all of this got started, ARPA turned around and said to UCLA: "Say, that project you're doing about measuring computer performance, how about if we change it a little bit and have you measure network performance, cause we're going to have this network?" And so, that sort of like, you know, turned a big wrench into that project, but they said: "Ok. Let's do that." So we, you know, threw away all this stuff about probing into -- inside another computer and these (unintelligible) think about how to measure networks. And I said: "Well, we better have a computer that can talk to the network in the first place, before we think about how to measure it." So there was a couple of, sort of sub-teams to this new project, that -- one that just sort of said: "Well, how do we have a computer interact with a network in any way?" And another sub-team that said: "Well, once we've got that going, what do we do about measuring it?"

Pelkey: Ok.

Postel: And so I got involved with the group that said: "Let's get this computer programmed so we can talk to the network." So, then I start learning about, well, BBN is building this network with these IMPs and, how does a computer talk to an IMP and how does a computer talk to another computer on the other side of the country once this network exists? And, UCLA was picked to be the first site on the, and SRI was second and UC Santa Barbara was third and University of Utah was fourth. The idea was that they were going to build this four-node network and do a lot of experiments to prove that it works, and then after that they would extend to a 15 node network, and that was the grand scheme. We were going to have 15 nodes around the country and one computer at each site and that was going to be the network.

Pelkey: Ok. Do you recall how UCLA got picked as the first node?

Postel: No, I have no idea.

Pelkey: And who else was on your team that was thinking about --

Postel: Well, people like Steve Crocker, Vint Cerf and Leonard Kleinrock were really involved in thinking about the measurement problem and how to do some analysis of network behavior --

Pelkey: But Cerf was on the team that was looking at measurement?

Postel: Ah, yeah.

Pelkey: He wasn't connecting to the net?

Postel: Well, I mean, the two teams interacted a lot --

Pelkey: Ok.

Postel: So yeah, but Vint was very involved in the measurement side of it, and --

Pelkey: Was Steve on your team?

Postel: Well, Steve was sort of overall in charge of it.

Pelkey: Oh, Ok. So two teams kind of reported to Steve --

Postel: Yeah.

Pelkey: -- and then Steve kind of reported to --

Postel: Yeah, well --

Pelkey: Was that -- that sounds formal.

Postel: Maybe it wasn't that structured.

Pelkey: Ok. Because Steve was the senior person --

Postel: Well he was the -- well, they were all graduate students. I mean, that's the other amazing part of this whole thing, in all of the sides, the people who were doing the work, except BBN, were essentially graduate students. Alright, and stuff like, you know, these graduate students from UCLA would go meet the graduate students from MIT and talk about how this network was going to work. You know, and Steve tells a story, we sort of all expected that one day, you know, we would be doing this stuff and at some meeting the real professionals would show up and tell us that, you know, those were all dumb ideas and this is the way you should really do it. You know, this has all been some sort of an exam or something, but it never happened, and the graduate students -- the way the graduate students planned to do it is what happened, at least from the protocol that went on in the host -- in the computers at the end, the host computers. We had -- somehow we ended up with the term "host computer" as the, you know, the thing outside the network that connects, and we're sort of stuck with that still, so when we talk about hosts we're just talking about the computer now -- and in another state --

Pelkey: Do you think that term arose at that point?

Postel: Yeah, yeah. I think that a host on the network was the computer you attached to an end.

Pelkey: I wonder whom --?

Postel: Well, I talked with Steve about it one day and we weren't real clear; but I think it came out of conversations that we had.

Pelkey: Between Steve and you?

Postel: Between Steve and -- well, we just can't be very precise about any of that --

Pelkey: But you think it came out of the UCLA group?

Postel: Its fairly likely. Its just very hard for me to remember exactly where it came from.

Pelkey: I understand.

Postel: But I think it probably came from -- within the context of the UCLA group. And Steve, it was really Steve who started these RFCs and many other series of notes. He really had this sort of work habit. Any time you got involved in any project, starting a series of notes about that project, and in the UCLA group that was doing some other things about computer research, we had several different series of notes. And so, as we began to talk with people about how the host should use the network and the

details of how we should use the host to IMP interface and things like that, starting a series of notes that documented those conversations was very natural thing to come out of that group, and Steve in particular. So, once that series of notes got started; it sort of fell to me to keep it going, and to sort of make sure that contributions got listed and indexed and things like that.

Pelkey: When you say "fell to," was that?

Postel: Well, you know, it was -- it was a job that UCLA volunteered -- Steve volunteered UCLA to do, and I ended up doing it.

Pelkey: Why you don't remember, but --

Postel: Well, I don't remember, but it was a good thing. And of course in those days, there wasn't a lot of word processing things on computers and so a lot of the memos are written on typewriters of line and sent around in the US mail and so, I mean, we're up to a thousand and forty RFCs now, but the first few hundred are only available on hard copy. They were never on line.

Pelkey: That's interesting.

Postel: So, and every once in a while somebody comes back and -- well, we put some of the old classic on line, you know, just retyped them, so that they're there, but, you know, a lot of the old ones --

Pelkey: Still are in hard copy.

Postel: Still only in hard copy in somebody's file cabinet. And a lot of -- you know, we're over a thousand, but there are probably only about a hundred that are meaningful anymore. And they range across everything from announcements of a meeting to the definition of a protocol that's now a very standard thing --

Pelkey: Right.

Postel: So, all kinds of different topics and subjects are in there.

Pelkey: Were they called Requests For Comments in the very beginning?

Postel: They were called Requests For Comments in the very beginning and in the very beginning the idea was really that this is a document of a conversation or an idea and its probably a pretty crazy idea, so it really deserves somebody coming back with some analysis and critique on it, to say: "No. That idea's all wet, and here's the reasons why, and -- do this other thing." So they really were intended to be a back and forth flow of ideas. And they were really intended to be relatively informal and things that you weren't embarrassed if, you know, somebody later pulled up the RFC and says: "You wrote this note and its completely wrong, what a stupid guy you are." You could say: "Hey, you know, it was a just -- we were batting around ideas." It was like standing up at a blackboard and sketching out a program and something and somebody saying: "Ah, look at this. This IF statement's wrong," or something like that, and you say: "Oh yeah," and fix it, and nobody really thinks anything bad of you for that. It was really supposed to be pretty informal.

Pelkey: Do you remember when the first one was?

Postel: The first one was -- I got it. The first one was -- (leafing through files) -- RFC-1 was by Steve Crocker in April 1969.

Pelkey: Seventh of April 1969. And the subject was?

Postel: The subject was "Host Software" was the title, but it was an idea for how you could implement remote terminal access with sort of elaborate kind of computer terminals so that the translation of what the remote system thought the terminal looked like to what this terminal really looked like could be done in a programmatic way. And basically the idea is that, the user is sitting here with his terminal and he wants to use this computer far away, what he would do is send a kind of program to the remote computer that is a sort of programmatic description of his terminals properties.

Pelkey: Well how about that?

Postel: And then -- then on the remote side there'd sort of be an interpreter for this program, and so that when the remote program -- application you're running said now display this on the terminal, that that would be translated through this programmatic translation to something that would actually make a reasonable display on the actual physical terminal you had.

Pelkey: I think that would be called a terminal emulation builder?

Pelkey: Right, but remote -- you know, by send -- by downloading the emulation program, each session to be appropriate to the particular terminal you were using that session.

Pelkey: Oh, ok.

Postel: This is an idea that still has not been implemented. It comes up about once every two years, and I send a little note off saying: "Go read RFC-1," and everyone says: "Oh, gosh. I thought I had a new idea."

Pelkey: That's great. And so anybody could submit them.

Postel: Yeah.

Pelkey: And they were submitted, at first, to the UCLA group.

Postel: Yeah. All these are old paper things in the file. Yeah, so anybody could submit one. They were -- basically you would call up somebody at UCLA media or a secretary at UCLA and say: "I got an RFC and I need a number for it," and you'd get a number assigned, and then you would prepare it according to some rules that we had for format and style and stuff like that, although they were pretty broad. And then send it out to the mailing list.

Pelkey: Ok.

Postel: And this is when we were doing it in hard copy. Then it was the author's responsibility to pay for the postage to distribute them --

Pelkey: Well how about that?

Postel: But then, later on we got more organized and got things on- line and, there was a period when you could send in a message, sort of send an RFC as a message kind of on-line, or as a file on-line, and the network information center would distribute it either on-line or hard copy, depending on what you liked. And then later on they said: "No. No more hard copy distribution. You got to be able to pick it up on-line or else." And that's the way we do it now.

Pelkey: And do you recall when you went to only on-line?

Postel: Probably around RFC-500, some place in there --

Pelkey: Which would have been roughly what time frame? Early '70s or --

Postel: Yeah, early '70s, '73, '74. Sometime in then.

Pelkey: So this is kind of a logging of --

Postel: Yeah. I made a little graph of the date and RFC number. There was a lot, you know, between '69 and '75.

Pelkey: So by '75, you were up to roughly 700.

Postel: 700. And then they were fairly flat for a while. And now there's a much -- sort of the community of users has jumped tremendously in the last couple of years, with the spread of the TCP/IP and many more hosts and many more networks, so the rate is sort of picking up again.

Pelkey: That is interesting. So that -- it took you until -- from when you started in 1969 it was mid '70 when you got your first hundred?

Postel: Yeah.

Pelkey: There was a hundred a year for a while?

Postel: Yeah.

Pelkey: Interesting.

Postel: Anyway, so, the basic ideas has really been to document everything, to try and document the ideas, and even document bad ideas, so that -- and that's actually something we've fallen down on, that mostly people will write up a description of something that's been tossed around and its harder to get sort of the raw ideas in RFCs now. Mostly people bat around their ideas one way or another, and only one things settled down do they get to be RFCs. And one of the things we sort of feel plagued with is the bad ideas come up -- or ideas that we've just chosen not to do. Not necessarily bad but just different way of doing something and we already have one way of doing it, and they say: "No, we're just not going to do that because we just don't need three ways of doing this thing."

Pelkey: Yes.

Postel: Come up sort of every year, or every other year, and you can't point to some document that says: "Here's the complete discussion of that idea, including why we're not doing it." And so that -- that just came up again last week, something about some small point inside of computer mail is a way to handle errors, and so that just came up again. "Well, you know, that's an idea that we talked about two years ago and went through this whole argument about --" if somebody would have documented it all, we could just say: "Oh you guys that are having this idea, go read this memo," and we wouldn't have to go through the argument again.

Pelkey: When do you think the characterization of it being a form of just ideas, good or bad, versus it started to get filtered by people necessarily only wanting to stick pre-digested pre-argued pre-packaged ideas for comment?

Postel: I think it has to do with the size of the reading community, in that, when you felt that the number of people that were going to read this was five or ten at each of 20 places, you said: "Ok." So all of a sudden it’s an audience of 100 people or two hundred people --

Pelkey: The damage isn't that great if it’s a stupid idea.

Postel: Right and, you know, right! And they all part of this group that have worked together --

Pelkey: A culture.

Postel: A culture -- that have worked together for a year or so, you know, that's fine. But when you say: "Well, this might be read by a thousand people right away, and, who knows, it may live on for years and years -- I mean, the other thing is, in the early days, you know, the idea that this note would live on for ten years was, you know, nobody thought of that. You know, it’s going to be -- nobody's even going to be able to find this in six months. So, the idea that maybe a thousand people are going to be reading this and it might be around for ten years makes people a little bit more reluctant to just put crazy ideas out there.

Pelkey: Before going on with your involvement, I want to digress around this point, because this strikes me as a very important part and point, at least from my perspective. But, let me ask you another comment, before we come back to this. When you guys started up in ARPA, the Arpanet, E-mail wasn't part of the design.

Postel: Right, it was not part of the design. It was not part of the requirements, as I understood it.

Pelkey: Do you recall how E-mail started up?

Postel: Yeah. Well, the computing systems that people had were early timesharing systems. Nearly everybody that was involved had some kind of timesharing system. Now there was a few people that had batch oriented systems, but most of the people that were involved, especially the sort of the grad student computer science department people that were involved had early version timesharing systems, so that they were PDP6s running the system that MIT developed or the system we had at UCLA was a Sigma-7, and we had spent the year before that developing a timesharing system for that; or the BBN-10X system for a DEC PDP10 that was coming out. Multix was in progress, and of course the Multix people viewed the Arpanet as a bump on the side of Multix. But it was, you know, a timesharing system and that culture of people were involved. The people at SRI were just -- had an SDS-940, so that the people who were involved in this early development were people who were using timesharing systems. And just sort of at that time, one of the things that were happening on local timesharing systems, just within one system, was a very crude form of electronic mail. I mean, you wouldn't recognize it as the same thing as what we call electronic mail today, but it was possible for a user on one of those systems to essentially leave a note for another user, so that when he logged on, he would get this note. And it was very simple minded. Like maybe they had a "TO" field and maybe they had a "Subject" field and that was all, and so on, but the fact that that capability was available locally to users was part of the environment that people -- these people who were developing these systems were working in. And I know that at UCLA on our Sigma-7 system we were, you know, we had decided that we really needed this thing locally, and we developed a local electronic mail system, so that we could leave each other messages and people could pick them up when they logged in, so --

Pelkey: Did you call it mail then?

Postel: Yeah, I think so. I think we called it mail. So then -- and this whole thing about -- about how I got into networks -- there was a committee that was trying to come up with a description of how to do -- have transfer files between computers and what, you know, what command language you would use and how you would -- how one computer would tell another computer: "I want file XYZ and send it to me in ASCII, or send it to me in binary, or --"how would you control all of that? And so there was a committee that had several meetings, a number of people involved from all different sites, to come up with this general description, and then of course somebody was designated as editor, to go off and write up what the committee had said and setting that up as a spec. And, as I remember it, and you may get different people remembering different things, is that between, you know, edition -- the third edition and the fourth edition, or the second edition and the third edition of this spec, the editor slipped in a couple extra commands that were, one of which's name was "mail," and said basically, it worked just like the command that would store a file. If you wanted to send a file to another computer and have it store it under a particular name, you would say: "Store this file under this name." So here was another command that had that same sort of style, only the command was "mail," and the argument was, instead of a file name,

was the name of the user that should receive the mail. And then there was some data there that went in and that was the end of the message. And so that was the beginning of intercomputer electronic mail.

Pelkey: And do you know, do you recall who the author was then?

Postel: The editor at that version of that memo was Abhay Bhushan at MIT. And he's --

Pelkey: What is his name?

Postel: Abhay Bhushan.

Pelkey: Is he still with MIT?

Postel: No.

Pelkey: Do you know where he is today?

Postel: I'm not sure where he is today.

Pelkey: Ok. And who -- this committee had a larger charter, so I gather.

Postel: Oh, yeah. Well, the committee had the charter of doing this -- developing a file transfer protocol.

Pelkey: And who was chairman of that overall committee?

Postel: Maybe Abhay was the chairman at that time.

Pelkey: Ok. Now -- so mail -- and this --

Postel: So mail -- the electronic mail between computers, intercomputer electronic mail, sort of came out of -- came as a sort of option on the side of file transfer.

Pelkey: And do you recall what time frame that was?

Postel: Well, we could nail it down again with our trusty -- you know, I'd have to go find which one of these -- which edition it was -- but, you see, we were working on -- it was the editor of the file transfer and there's one here that's RFC-172 in June of '71

Pelkey: 23rd of June 1971.

Postel: So that's --

Pelkey: A very good chance that that could be it.

Postel: A good chance that could be it. That might have been a later version.

Pelkey: Now, mail took off.

Postel: Mail took off, and it was -- mail was not a requirement of the system, it wasn't a pre-conceived reason for building Arpanet, but in terms of the applications, you know that have taken off and made -- had the most use, electronic mail is by far the most interesting and successful application.

Pelkey: Why?

Postel: Well, it’s a mode of interaction, and that's different than previously existed. And it’s the difference between making phone calls -- you end up playing telephone tag. You don't have to play tag with mail, because it’s sort of left. You've left your message as a message, instead of leaving a phone message that says: "Call Joe," you leave your whole message right there. So that he picks it up whenever he gets into the office and reads his mail, and it also -- it’s much faster than letters. It’s essentially a very cheap way of sending letters that get there very quickly, and if you happen to both be in the office at the same time, you can send electronic mail and the guy at the other end gets it in like a minute, and can respond to it, and you can have an exchange of letters -- four or five letters in an hour. On the other hand if somebody puts something into one of these messages that you have to go think about, you can just think about it and answer him when you get done thinking, rather than, like I'm on the phone. Someone calls you up and you say: "Well, that's an interesting question. I'll have to go down to the library and do some research and call you back." Well, you don't actually have to say that's what you're doing. For some reason you don't answer the message right away.

Pelkey: How much -- do you think that the free form nature of it, it can be two words or it could be pages, and there was no protocol of mail itself, did that matter?

Postel: I think the fact that it was so flexible is a big help, and it’s not a straight jacket at all. The fact that its completely random ASCII text inside the text inside the message is a big help, because you can use your favorite text editor to compose the message and if you don't -- you're not constrained to have line numbers or meet any particular syntax of the editor, so that people can build mail programs, and look to compose and send messages, they can build their own mail reading program, because the rules of what the beginning of the message are so simple, so that anybody can have their own favorite mail program and its fairly quick. You get a new computer that doesn't have a mail program you can build one fairly quickly.

Pelkey: And that happened --

Postel: And that happened in dozens of different, sort of user interfaces to compose and read mail, an they all work together because the rules are very specific, but also very simple.

Pelkey: Now -- in terms of --

Postel: Now, of course, over the years, the mail -- the rules for what a message looks like have gotten more formal and more complicated and --

Pelkey: Do you think that's changed mail?

Postel: No. I mean, I think the additions have clarified things. Basically what happened, when people begin to include other information, a header and a message, and two or three different ways, and then sometimes people come along and say: "Well, you know, we're doing this three different ways. Let's standardize that extra piece." So there's more people with say, OK now anybody else that wants to do -- include this feature in your message, do it exactly this way, so that there's some hope that somebody can build a mail reading program that can parse that automatically.

Pelkey: Ok. Now, if I might -- another question -- do you -- mail, obviously became, from what I understand, became the dominant user of Arpanet, in terms of traffic and so on, which as you say and as I've been told, was accidental if you will.

Postel: Yeah. It’s accidental. It wasn't in the plan, but, like I say, there was some sort of underlying culture for it in that the people who were using these early timesharing systems had done it locally.

Pelkey: And that community of people still communicates very heavily over mail -- its part of your communication mode.

Postel: You know, I'm fairly hopeless. If the computer goes down and I can't do electronic mail, might as well go home. I mean --

Pelkey: Now let me ask you, if I can, to contrast the RFC phenomenon of it being very free form and loose and allowed, when there was a small group, ten people at ten sites where people felt free to just go off and say something, and E-mail, in terms of its implications in terms of when -- one is that, in E-mail, obviously you can control the directory size as to how many people see the message that you're sending out --

Postel: Well, actually, not necessarily. We've organized electronic mail in groups in several different ways. We have some systems - - some facilities that work in what I guess we'd call, informally, a mail exploder, in that somebody maintains a list of destination individuals, and they essentially put that under a group name. Then, people who want to be on the list contact this list maintainer and get added to the list, but you can send to this group name, and have your mail essentially forwarded to everybody on the list, and you don't even know who's on the list. You have no idea how big the list is. And we have several lists like that that talk about technical issues, sort of subject area lists, and so it’s very easy to send a message that goes to hundreds of people. And you may think that you’re sending a message to maybe 20 people but it’s actually going to 100 people, or you may think you’re sending a message to a hundred people and its going to 1000 people.

Pelkey: Yes.

Postel: So that's a very interesting phenomenon, and we do that -- I mean we get away with that because we don't pay for it.

Pelkey: Right.

Postel: The individual that sends isn't paying for it, so if you had to -- if you sent a message to one of these things and all of a sudden you got a bill for --

Pelkey: $10,000

Postel: Right. 1000 times 22 cents, or whatever, you would only do that once, so we get a lot of stupid messages on these mailing lists because people don't think about how many people it’s going to.

Pelkey: And some of the odd contrasts, in terms of against in a conceptual level, but drawing upon your experience because you have an experience base that I don't have, but in commercial studies now, companies where they have E-mail, they find that they get a point where there's -- some people think there's too much E-mail going on, that people spend too much time doing E- mail and they're reading stuff they shouldn't be reading, and it’s a waste of time.

Postel: Well, if they're reading stuff they shouldn't be reading it’s a waste of time, but it’s just like any other source of information, I mean, those same people probably spent time reading The Wall Street Journal and Datamation at their desks, right. And, you know, well maybe they should read that article or maybe they shouldn't. How can you tell what they're reading, the relevant articles out of Datamation or Digital News or whatever it is, right. So, you know, people who are going to waste their time reading the wrong articles are going to waste their time reading the wrong articles, whether its electronic mail or whatever.

Pelkey: What kinds of phenomenon --

Postel: But, yes. I get a lot of messages, I got a lot of messages because I'm on somebody's mailing list, and you really have to decide, well does this mailing list have enough content that's valuable to me to make up for all the ones I have to look at and say: "No, skip that. No, skip that. It doesn't start the right way." So that's a judgment you have to make.

Pelkey: Have you found that within the community of E-mail users that there are in fact these shifts of groups and that happens frequently and people chose, I don't want to work on that problem or that group doesn't give me enough value --

Postel: Right, right. You get on a mailing list and say: "Yeah, this is an interesting subject to me, but most of the discussion on this list is at the wrong level. Either down on the hardware level, and I'm not interested in that, or its philosophical, rather than practical, or whatever. Or maybe you're involved in policy and you get on a list and you say: "These guys are all talking about the practical aspect and I want to talk about the philosophy." So you migrate around and find a set of lists that are relevant to what you are doing.

Pelkey: And you find that lists, if you will, come into being when there's a problem --

Postel: Oh yeah.

Pelkey: -- and do they die or they kind of have -- they just kind of keep going on?

Postel: Well, they don't -- they die in the sense that there's no messages sent to them after a while, but they -- It's easy to leave the mechanism in place, so what you'll see is, as a list that's been around for a long time, may fall into disuse and people don't send messages to it anymore, because people on the list have already fought through all the issues that were relevant, and then some new guy comes along, and tries to start talking about those issues someplace. And somebody will say to him: "That issue was talked "

Tape side ends

Postel: Some new guy will send a message to an old list, and everybody else on the list will say: "We already settled that problem. Why don't you go read the archives?" So that's one of the other things that can happen with these mailing lists is that it can be sort of a file associated with a list that gets a copy of all these messages, so that when some new guy comes along and starts talking about something you've already talked about, you can say: "Go read all the -- go read the file of all the old messages before you talk about that."

Pelkey: Right. Do you find that there are certain people who are, I guess, a moderator of a list, or a keeper of a list that know how to maintain a quality or flow of the mail?

Postel: Well, Ok --

Pelkey: -- where it gets more productive?

Postel: Well, that's what I was beginning to talk about. There are several different modes of these lists. The ones that I've been talking about have been what I'm calling sort of "exploders," no moderators, no editor, there's -- if anybody sends a message to this mailbox name, this group name, it is immediately forwarded to everybody on the list. There's no chance of any interference or influence or control or editing or quality control at all. Ok, another mode is a moderated group or moderated list, where when you send a message in, it goes into a mail box or a file that the editor of the list or the moderator of the list reads and he can throw away the stupid messages, and clean up or delete paragraphs from messages that are not relevant to -- and then put out what's sometimes called a digest.

Pelkey: Ok.

Postel: And he may do that once a day, or once a week, or whenever the traffic on the list merits, of the good parts of the messages that have shown up in the moderator's mail box.

Pelkey: Right.

Postel: And then everybody sort of gets this, so it’s more like a magazine that's -- or a letters to the editor column, for the editor that's selective about what he passes on.

Pelkey: Right.

Postel: So there you get maybe a better discussion. You get higher quality news, but you don't get the immediate feedback. Say like, with one of these exploder things, you can send him a message, and you'll probably get an answer back instantly, because somebody will --

Pelkey: Somebody somewhere --

Postel: Somebody somewhere is probably on line at the same time. So the digest thing has been very useful and a very important way of gathering and distributing information, but it takes a lot more work. Somebody's basically got to dedicate him/herself to being the editor of that thing, and it takes a lot of time

Pelkey: You mentioned that there were three types, the exploder; the digest --

Postel: What is another kind? Well there's another thing that people do. They can maintain a list of people who are interested in a subject, sort of privately, and send messages to it, so that it goes to all these people, but that list is sort of hidden, and nobody else can send messages to that exact list of people. So that's sort of like a private file that can be used as a mailing list. And I guess there's an example of that, in that one of the things that we're doing currently is, there's a list of a couple hundred people who are fairly heavily involved in the research on networks and protocols and mostly associated with sponsored projects, one way or another. And people send in a monthly report of a couple of paragraphs from -- maybe twenty different projects do this -- and we assemble that into a monthly report of, basically a concatenation of all these submitted reports, and then send it to this list of 200 people. And that list is essentially our private list. It's not -- it doesn't work as an exploder or a digest.

Pelkey: Right. I have a view that kind of in this book that the organizations of the future -- part of what makes an information based company is in fact this electronic communications, in that the way groups will start to communicate ad hoc around what I call perturbations of the organization, who start to give meaning to the behavior that's outside the organization, the perturbation, and that, these are going to flatten organizations immensely, in that organizations are going to much more concerned with the process of intra-organizational communications and so on, than they are with the kind of paradigm of, "Go collect all that information outside the organization, model it, make better decisions from that -- " It's going to be a very different kind of way of looking at information, and information is -- how do you capture the information and judgment and experience and so on within the organization to respond to events that are outside the organization? More of a biological metaphor --

Postel: Yeah, I don't know. I don't know how to speak to that, because I've always been in this research world, where, from the researcher's point of view, its much more interesting to keep in touch with people doing similar work, across -- ignoring any organizational boundaries --

Pelkey: Right.

Postel: -- so I'm much more --

Pelkey: But in fact there is an organizational boundary. The organizational boundary you've chosen to define. Our researchers are interested in what your interested in.

Postel: Say, for instance, I'm at ISI that is part of USC. I don't have very much concern about ISI or USC! I have a lot more concern about the network research community. That's what I do, so in terms of my electronic mail, probably a much larger percentage of my electronic mail involves people outside of ISI

Pelkey: That's right. Probably 90% of your mails relate to the interests you're interested in and 10% is USC or --

Postel: Right.

Pelkey: Whatever.

Postel: Well its probably a little -- well, not very much has got to do with USC, but maybe something's got to do with ISI.

Pelkey: Right, ISI, right. One other questions along these lines, different -- I'm told that there was an amazing amount of traffic in the early days of intra-IMP traffic that never even got onto the net --

Postel: Right, right.

Pelkey: -- which was an unexpected behavior.

Postel: Right. One of the things -- basically, the earliest model, the earliest thinking about this network was that it was going to be one big computer in each city and that was all, so you needed one IMP to connect to one computer per city, and before anything got fielded, and fairly early in the design stage actually, I guess, people said: "Well, you know, they might have more than one computer in a university, so maybe we ought to provide for up to four computers per IMP, and repartition the address space and design changes to make that happen, and very early in the game the second computer at a site showed up and a lot of the traffic was between those two computers at that site. And so, in some sense, having an IMP at your site was equivalent to having a local network --

Pelkey: Right.

Postel: -- so the intra-IMP traffic within a site was the first local network that most people had.

Pelkey: Fascinating. Coming back to -- I want to come to the issues of protocols in a second, but you went from UCLA and, finished your graduate work and -- when did you leave UCLA and what did you do then?

Postel: Well, I finished my PhD at UCLA, and I went to work for MITRE in Washington.

Pelkey: What year was this roughly?

Postel: I have to go look all these things up. I can't remember any of this.

Pelkey: The immediate man.

Postel: Right. No I still remember some of --

Pelkey: Early '70s.

Postel: Yeah, probably '74.

Pelkey: Ok, '74.

Postel: So I went to work for MITRE and, of course MITRE doesn't do -- somebody said MITRE does studies. And so I worked for MITRE for ten months and then there was this little company -- well, there was a company that was running a data management system, a data management computer for the ARPA office in Washington, and they had -- it had been programmed to be accessible from the network computers in the ARPA building, but had -- they had taken some shortcuts and had some special tricks

built in so it was not accessible from anywhere else on the network. And, this was the way ARPA wanted it. They didn't want any of this data in this computer being stolen by any of the undergraduate students at any of the universities. So, but then later on they discovered that once in while they were going to travel, and one of the ARPA --

Pelkey: They wanted to access it.

Postel: - and they said: "Hey, I've got the network. I can access my data management system." So they wanted that all fixed, and it was probably time to redo it anyway, so I took a, essentially a summer job with this little company, reprogramming their --

Pelkey: Their IMP.

Postel: -- No, the network in this data management computer, which was the host connected to their IMP. And, that was sort of mine and (unintelligible). And then at the end of that time, at the end of the summer, I went to work for SRI in Doug Englebart's group at SRI --

Pelkey: That was around '75, '76?

Postel: Yeah, '75 '76. And they were just beginning to be involved in this big project called the National Software Works, which was supposed to somehow put an overlay on all these network protocols to build a multi-computer system that looked to the user like one computer, and be based on computers of different kinds. So you were supposed to somehow be able to log in to this National Software Works as if it was a single timesharing system and say I want to compile a -- I want this file to be compiled as a, by a FORTRAN compiler, and the system was supposed to look around and find some computer that knew how to compile FORTRAN programs and send the file over there and get it compiled and I don't know what else. Ok. But, it was a great idea but I don't think the resources were quite up to the job.

Pelkey: Right. But that was a creative group at that point.

Postel: Yeah. There were several contractors involved to do this National Software Works, but the -- but this was -- in Doug Englebart's group with his on-line system of document preparation and Hypertext. I guess is what it's called these days, was, you know, very powerful document preparation system, was supposed to be available through this National Software Works, and that was their involvement. And they were also looking -- and they had a very powerful front end user tailor-able interaction scheme. And the idea was that they would be able to take that and make that into sort of the user front end for this National Software Works.

Pelkey: Oh, ok. Is that where the black and white and the mouse and those things came from?

Postel: Yeah. The mouse came from there.

Pelkey: How about that. Who else was in that group, do you recall?

Postel: Oh, Lots of people: Lots of people.

Pelkey: Anybody else who -- from the communications side? Who either went into or came out of kind of the communications where the ARPA sort of experiences or?

Postel: Well, I don't know communications. A lot of the -- we had a joke that it was the Xerox training center, in that a lot of the people from there went to Xerox PARC. A lot of the early people at Xerox PARC were --

Pelkey: Now you can have a second order derivative joke of they staff DEC

Postel: Right.

Pelkey: DEC PARC.

Postel: DEC and Apple and so on.

Pelkey: Did you -- did a group of you interact with Vint Cerf at that point in time?

Postel: Yeah, right --

Pelkey: Was it at Stanford?

Postel: Yeah, no, when I was at SRI Vint was at Stanford, and he was just beginning to do the first experimental work on TCP.

Pelkey: Did you -- was there -- did the two groups interact where you --

Postel: A little. Not a lot.

Pelkey: Did you interact with Bob Metcalf at that point in time?

Postel: No, well, I had interacted with Bob Metcalfe earlier when he was still at MIT and when he was at Xerox developing the first Ethernet at Xerox, there was some discussion, and he probably talked -- there was probably more discussion between Xerox and Stanford than otherwise (unintelligible).

Pelkey: Ok.

Postel: But we knew about the Ethernet fairly early on.

Pelkey: Ok.

Postel: When it was still very experimental.

Pelkey: Right. And then, how long were you at SRI?

Postel: I think almost three years.

Pelkey: So now its around '79 or so?

Postel: Yeah, and then I came down here.

Pelkey: Came down here. And during all this period of time you were the RFC chief?

Postel: There were, yeah, there were probably a couple of periods of six months when each -- that somebody else was trying to do it, but, you know, they didn't do a very good job or something like that, I took it over again.

Pelkey: Did, by popular request did you get back involved --

Postel: No, I --

Pelkey: -- or did you get frustrated and you just --

Postel: I just got frustrated and took it back. It’s important. Basically, I thought it was important that it keep going and, I think that's one of the things that has made these ARPA protocols so successful is that some new guy who wants to learn about it CAN get his hands on the documents.

Pelkey: Right.

Postel: And the fact that it’s a shared thing, it doesn't belong to any one company in any sense --

Pelkey: Right.

Postel: -- has been a key thing and attractive to a lot of people.

Pelkey: Yes. Now, along that line, do you have any views or recollections around -- there was a period around -- it might have been '74, my dates are fuzzy, when ARPA wanted someone else to take the management of this over. They went to AT&T, so I gather, and there was some discussion about splitting it off. Maybe it was earlier, it was, I guess the time when TeleNet, BBN created TeleNet, and there was some discussion about, rather than everybody going off and doing their one, why don't we create -- why don't we let everybody take this and maybe all these networks can finally come together and we'll have one big network in which all the protocols are the same, versus having a number of people go off in different directions, and they become incompatible with -- and so on. Do you remember any of that?

Postel: No. Not really. I don't think I was really involved in that, well planning that. I mean the motivation for TCP was that ARPA was beginning to experiment with other kinds of networks --

Pelkey: Radio networks --

Postel: Radio networks, in fact in satellite networks. And that the realization that you need to communicate between a computer on one of these kind of networks to a computer on a different kind of network, so that you need a protocol that spans across these networks of fundamentally different physical types or media types, and -- so that was really the motivation for starting TCP.

Pelkey: (Affirmative).

Postel: And I suppose -- well, I wasn't very much aware of the issues of ARPA's control of, or versus spinning off to some other sponsorship or other management of the networks then. I know later on they managed to get the day to day worries of the Arpanet handed off to DCA.

Pelkey: In '75, yeah.

Postel: In '75. Of course that was a total disaster, but -- cause DCA is just not very responsive.

Pelkey: Well, DCA still manages Arpanet right now.

Postel: Right. If you want -- you want another IMP in a different place or you want more wires in the Arpanet or something like that, you've got to go to DCA. DCA will eventually get around to it.

Pelkey: Now, Paul Baron refers to Arpanet as an experiment that went bad, in that when it first started off, it was meant to be low speed, just a some -- really done in software, to see if, in fact, you could prove that you could have a packet switching network, and it still exists today. I mean it found a life of its own and --

Postel: Well, another phrase for it is a success disaster.

Pelkey: Yes. That's another expression, same concept.

Postel: It worked --

Pelkey: Too well.

Postel: It worked too well. It worked much better than a -- it worked to well that they use in part exceeds its design.

Pelkey: And part of the reason it worked so well was E-mail.

Postel: Part of its popularity is -- yeah.

Pelkey: People didn't want to get rid of it because of E-mail.

Postel: The underlying network did packet switching very well, and it supported this --

Pelkey: Application.

Postel: -- this E-mail application that people fell somewhat in love with E-mail that the demand for use of the network far exceeded its design.

Pelkey: Yeah. And then I guess the government went off and did DDN and whatever they did --

Postel: Well, that's a very amazing story. That is an actually incredible story, in that Arpanet was being built and designed and beginning to appear successfully, and DCA says: "Ok, we need to -- we have this AutoDin network, and we still have the AutoDin network, I guess --

Pelkey: Now this is what time frame are you talking about?

Postel: -- and -- it was probably early '70s -- and they said: "Ok, we'll go out -- We'll put out and RFP and we'll try to buy a network kind of like the Arpanet. It'd be a modern packet switching network to replace AutoDin, so this would be AutoDin II." So they went out with an RFP for AutoDin II and they got bids from three or four people and, there was one by Western Union who did AutoDin I, and -- with some help from some other companies -- and they proposed to build a packet switching network and sort of all the same principals of the Arpanet, but of course they were going to -- they weren't BBN so they were going to build it differently. And they were -- You know, they had something like two years and a hundred thousand dollars or two hundred and -- I don't know, some huge amount of money, a million dollars or a hundred million dollars or something to do it.

Pelkey: Right.

Postel: Two years came and two years went and it still wasn't running. So some place along there they got -- and I wasn't involved in the -- there got to be a committee in DOD to review this and decided what to do about it. And extremely remarkable thing in the history of this county is a major defense program was stopped. AutoDin II was cancelled, and the decision was made to not do AutoDin II but to install a copy -- an exact copy of the Arpanet. That became MilNet which is the fundamental part of DDN.

Pelkey: Isn't that something. I hadn't heard that.

Postel: Ok. Go find the story of AutoDin II.

Pelkey: That's a great story. I will.

Postel: Steve Walker. Steve Walker is the guy to ask.

Pelkey: Where can I find Steve Walker?

Postel: He's now at a company -- he was in the office -- he was in office of Secretary of Defense for Command Control and Communication Intelligence at the time, and now at a company called Trusted Information Systems.

Pelkey: Oh, he's with Steve Crocker.

Postel: Well, Steve Crocker's here on the west coast. Steve Walker's on the east coast, but yeah.

Pelkey: Good. Thank you. That's helpful.

Postel: You should go -- tell me the story of AutoDin II.

Pelkey: Good. I appreciate that. Let's come back to the issue of protocols now.

Postel: Ok.

Pelkey: BBN was doing the IMP -- software for the IMP and the IMP to IMP protocols.

Postel: Right.

Pelkey: And each IMP site would do the host protocol -- of connecting the host to the IMP.

Postel: Well, ok, BBN defined the interface. Here's the IMP. So it’s a box and its got a wire sticking out of one side which goes to a host. So there's an interface here.

Pelkey: Ok.

Postel: This interface is defined by BBN.

Pelkey: Is that the spec 1822?

Postel: Right. BBN 1822, that's report # 1822 from BBN is what is called 1822, and -- so that defines the interface that the host has to use. And then the IMPs talk to other IMPs using the IMP to IMP protocol, which is completely controlled by BBN and they could change it any time (unintelligible). Ok, so then the host can talk -- use 1822 to sort of hand information to the IMP, which -- and the instructions in that, across that interface will cause that IMP to deliver it across a similar interface to some other host some place else.

Pelkey: Ok.

Postel: But now the host has to decide -- so basically the host says: "Here's some bits. Here's a little bit of 1822 information on the outside of it. I give that to the IMP. That will cause the IMP to deliver that bundle of bits to some other host." Now, should these two hosts have any agreement about the format of those bits, what the interpretations is, that's a host to host issue. IMPs don't care. They say: "Hey, these are bits. I don't care about what the format of those is." So now the guy who has this host and the guy who has this host have an agreement about what the format and the interpretation of those bits are. And, like, probably it’s from some program that's running on this host and its probably intended for some other program on that host. Well how does -- what do you -- you need to attach something else to those bits that say: "This is from my program X to yours program Y." So how do you do that then? So that's what host to host protocol is all about. And that's -- you really want not this guy and this guy to make an agreement, and these two guys to make an agreement and this guy -- that's sort of N squared. So you want to say: "Let's get everybody together and have one agreement about what host to host protocol is, and then once I've got it implemented to talk to host A, I can talk to host B and C and, when host Q comes along that we haven't even heard of yet, we just force him to do the same thing, and then we can talk to him without changing anything."

Pelkey: Now, is that the network working group?

Postel: The network working group was essentially the committee that started defining host to host protocol.

Pelkey: Right, ok. And RFCs --

Postel: RFCs were sort of the notes of the network working group.

Pelkey: Was Steve Crocker the first chairman of the network working group?

Postel: Yeah.

Pelkey: And that -- within that networking working group is when you defined this concept of these layers?

Postel: Right. Well, like I said, I think -- my view is that was a real trivial idea, because, like I said, all the people who were working on this were coming from the context of these timesharing systems or at least computer operating systems, and computer operating systems have the idea of layers of abstraction or layers of functionality. They've had that for years before networks came along. So think about the operating systems that has -- you know it has some real low level code that runs devices and stuff like that and its got some idea of disks divided up into pages and its got some idea of a file made out of pages and its got, you know, system calls to do the file system, so its got these layers already built into it. So, all the people -- all the culture -- all the people who came to this networking program to define these host to host protocols were operating systems people who knew about layers of stuff in the system. So the idea that -- the idea that the protocol was built up out of layers was perfectly obvious.

Pelkey: That's right, intuitive.

Postel: Intuitive, it was -- well, sure, you have the stuff that the host absolutely needs to talk to the IMP and then you have the stuff that the operating system in this host needs to have to talk to the operating system in that host, and you have the stuff that this application program needs to talk to that application program, so you got layers. And you just -- that was the way you do it. There's no other choice. So -- and this whole -- like the ARPA community sort of suffered for a while from, you know, when the ISO guys got started they said: "Well, we have this protocol architecture. And we have all these layers." Well my goodness all we had was three layers. There's five layers or, depending on your caste, we never thought the number of layers was particularly important. But we have -- the ISO guy said: "We have this protocol architecture." And so we were always getting asked the question: "Do your ARPA protocols follow the ISO architecture?"

Pelkey: Right, right. Do your layers match up?

Postel: Well, maybe, it sort of maps the question is sort of: "Do your layers match up?"

Pelkey: Right, right.

Postel: But the question was, well wait a second, I mean, what is a protocol architecture? What does this mean? Do you have layers? Oh! So, I went -- so a couple of years ago somebody went to a conference and presented a paper on the ARPA protocol architecture. We didn't think that the protocol architecture was significant enough of a question to need a definition. I mean it was just obvious that you had layers.

Pelkey: And cross all the guys in the network working group.

Postel: So, yeah. So now we have an ARPA protocol architecture that we can talk about, but, it just basically going back and giving that name to what we thought was a real simple idea in the first place.

Pelkey: What were some of the issues that the group of you struggled with in the network working group in the first year or two as you got -- did you ever get -- I guess this is meant to the ICC, the scenarios for using Arpanet in October of '82, '72 --

Postel: '72.

Pelkey: - was a big event.

Postel: Yeah.

Pelkey: I mean, that -- you kind of came out of the closet at that point in sort of a sense there --

Postel: Well, I don't know -- It was a big event in a couple of ways. One is, yeah, I mean, basically -- well we didn't come out, I mean, it wasn't like anybody was keeping anything secret because, I guess in those years NCC -- (unintelligible) NCCs were big events. And NCC-70 and NCC-72 there was a whole session of five papers on Arpanet in both of those. And at other conferences in between there was Arpanet papers, so the community was publishing lots of papers about Arpanet things all along.

Pelkey: Ok.

Postel: But at thing, we basically put on a demonstration. I mean, when people go to conferences you give a paper about Arpanet and they sort of go away and say: "Well, how real is that? I mean, are they just talking about papers, or are they -- does this thing actually run? Does it run as well as they say?" So, ICC was, like, here it is live demo. So, that was sort of a major milestone in giving a live demo to whoever randomly showed up in a real concentrated way. That was one aspect of why it was a spectacular thing, but the other part of it was, that it was a (unintelligible) forcing function to all of the people involved and get their act together so they could participate in the demo. So lots of work went into it in the month ahead of time of making sure that your host at UCLA and your host in BBN and your host in MIT and your host in wherever else had some applications running that could be accessed over the network in a reasonable way.

Pelkey: Affirmative

Postel: So that was really also another important factor.

Pelkey: Ok.

Postel: You know, a lot of testing went on and a lot of work went on to make that a success.

Pelkey: So that milestone -- I guess there was a lot of pride amongst the people that had been working so hard in '69 and '70 and '71 to kind of pull together and show off their stuff to each other as well as to the --

Postel: Well, that's another thing that sort of came out of this whole thing, I mean, is that, even before the demo, even before that, that one of the things you would see happening is that, anyway, if your a systems programmer and responsible for developing new applications on your computer and you go visit somebody else's computer center and you see them running a little program -- Ah, ha -- a program like that, but I can probably put something together that does that. So you go home and you make a program that does something similar and maybe a little better --

Pelkey: Right.

Postel: Well, -- and people don't actually move around that much, so that would happen occasionally.

Pelkey: Right.

Postel: But with this network, as soon as you were able to do remote access and you could log into somebody else's computer and see what applications they have, you say: "Hey, you know their application to do mail has this really neat feature," or "Their application to list out who's logged onto the system lists this other factor about what room they're in," or something like that --

Pelkey: Right.

Postel: -- Well, I'm going to go fix our to do that. So there was a really spectacular growth in terms of software improvement in terms of, you know, "Hey, I can make my system do -- "

Pelkey: Functionality improvements.

Postel: Functionality improvements to make what goes on in Tenex to be as good as what goes on in Multix to make it as good as what goes on in this other system, and so on. So a tremendous feedback of good ideas about how the software should work and functionality of applications happened because of this community working together and the fact that there was this new kind of access that you could check out other people's systems.

Pelkey: Right, and you guys were doing that -- checking out other systems before the conference --

Postel: Oh, yeah.

Pelkey: -- but. Ok.

Postel: And so that -- so there were a lot of people that worked together and knew each other through the network.

Pelkey: Now, at this point in time, the community of people who were really actively involved was a couple of hundred maybe?

Postel: At most.

Pelkey: At most. Maybe a hundred?

Postel: Yeah.

Pelkey: Yeah, yeah. Now, during that period of time, from '69 -- when you established the network working group in '69, to the scenarios for the Arpanet, dealing with the host to host, were there, I mean -- you kind of, say -- the layering was kind of obvious, but were there technical issues or were there design issues --

Postel: Yeah there were.

Pelkey: -- that WERE a real struggle, that --

Postel: Yeah, there were real issues. I guess, one I remember was the fact -- was there was an idea of trying to be able to support the . . .

Tape side ends

Postel: -- the idea that you can have a program running in one computer and it could hand off the computing functions to a program running in another computer -- supposing you originally somehow started using this application program on computer #1, and you got to a point in that application where, you know, it needed to access some other data or something like that, so that the best thing to do was to, essentially refer you to another program that runs on computer 2. Well, it could do that. It could just type out on the terminal: "Hang up and go call computer 2 and use this other program." But the idea that you could somehow do that automatically behind the scenes -- that the network protocol in this system could sort of call up the network protocol in another system and say: "I'm going to hand you this connection," and connect it up to this other program. So its sort of like, sort of unplugging the wire from one socket and plugging it into another socket --

Pelkey: Right.

Postel: -- sort of behind the user's -- without the user knowing about it, was called a reconnection, a dynamic reconnection, and that was an issue that Steve Crocker and several of us talked about argued to have it designed into a very early draft of sort of this host to host protocol. And other people just felt that that was too complicated and not needed. So there was a big argument at that, and eventually it got dropped and wasn't included in the protocol.

Pelkey: Ha!?

Postel: So there's a sort of an example of a technical issue --

Pelkey: Right.

Postel: -- where there was a lot of debate, and memos were written back and forth about whether this was good or bad or whether its worthwhile and, you know it was one of the ideas that got left out along the way.

Pelkey: Right.

Postel: -- and again, its one of those things that people come up with from time to time, saying: "Hey, wouldn't it be neat if we could do this," and we say: "Yeah, it would really be neat. Why don't you go read RFC, you know, 33. It's in there."

Pelkey: Where there -- an issue of within, from the Spec 1822 from this wire thing, within the IMPs, the protocols in there really haven't surfaced in terms of having an impact, other than people who might have worked on them --

Postel: Right, right.

Pelkey: -- It's really been the host to host that's gone off to TCP/IP and --

Postel: Yeah. You see, everything inside of this BBN 1822 boundary --

Pelkey: Which they called a subnet at that point, right?

Postel: Yeah -- has been sort of BBN's thing. And they published some information about it in BBN reports, but nothing about that has been published in RFCs.

Pelkey: And there's probably lots of stuff (unintelligible) in BBN reports.

Postel: And they've given papers about it generally, I mean not the specific details, but they've given quite a lot of public papers about what it's like to run this network and sort of general statistical behavior and what the routing algorithms are that they use and all that kind of stuff. So, I mean, they haven't been

particularly secretive about it, but on the other hand the exact details I don't think have been published in a public way.

Pelkey: Now, E-mail kind of came about, but FTAM and VTAM were really specified, I mean, in terms of--

Postel: Well the, yeah, the file transfer -- basically we identified remote terminal access and file transfer as functions that were needed, and set subcommittees, you know one guy from each of the --

Pelkey: That was defined in the network working group early on, or was it kind of --

Postel: Yeah, well, and I think it was -- it came out of the very earliest discussions. I don't know exactly, but maybe people from ARPA were involved in saying: "Yeah, you have to do those things," or whatever, but they were fairly early on identified as things that we wanted to do. I mean, basically, the idea that these are all timesharing systems and one thing you want to do is use that system in another city as it you were there from your terminal over here. So its remote terminal access, which we ended up calling TelNet. And --

Pelkey: Where did that name come from?

Postel: Well, that name came from Steve Carr who was the guy at University of Utah, and I don't know that it actually stood for anything --

Pelkey: Ok.

Postel: There are some later papers where its, just, made into the acronym of something, but I think TelNet came first and the acronym -- using it as an acronym for something came later.

Pelkey: Ok.

Postel: But the idea was that somehow you would get on your computer and run this little program that would somehow use network protocols to get you connected to another computer where you could log in as if you were there, and then something had to be done about the fact that that computer expected you to have a certain kind of terminal, with certain kind of physical properties, and in fact you had a different kind of terminal, cause you had the kind of terminal your computer has, and they had to do some translation between the physical properties, the programming specs and the physical properties you have --

Pelkey: Right.

Postel: -- and so we came up, and very early in SID we came up with this idea of the virtual terminal, which was -- the idea, this is the standard terminal, and what you do at your end is you map from the terminal you expect to the standard terminal and at our end we'll map from the standard terminal to what we really have, so we each do sort of half of the translation, but this makes it so that you do it once and I do it once and it turns -- its not an N squared problem anymore.

Pelkey: Right. Now, did you parse up, like, FTAM and VTAM to different sites so that --

Postel: No. Each side had -- we formed committees, and it would be perfectly acceptable for each site to have a member on each committee by subject, so that the TelNet committee could have a person from every site and the file transfer committee could have a person from each site. And there may be -- there might have been some place that didn't care to be on the file transfer committee or didn't care to be on the TelNet committee, but --

Pelkey: Now, at this point in time you've gone from like four IMPs, and it went to 15 pretty quickly, right?

Postel: Yeah.

Pelkey: And when you (unintelligible) the scenarios, were there more than 15 at that point in time, or roughly, do you recall?

Postel: I think there were still about 15. I mean, there might have been 20.

Pelkey: So there were about 15 sites who were working these issues at different levels.

Postel: Yeah, right. And then there was other things like, at UCLA there was the group that had this Sigma-7 timesharing system and there was the group that had this IBM 360 mainframe batch computer, and they had clearly different points of view about things, so that, in fact, there might be a guy from the IBM system and a guy from the Sigma-7 system both on the file transfer committee.

Pelkey: Ok.

Postel: So that you might have two people from a, quote, site, from a univ -- the same university, even, or company even though, they represented different -- fundamentally different operating systems, or something like that.

Pelkey: Gotcha. I don't here much mention about the University of Utah. Was that an active site? I mean, no one has mentioned anybody from the University of Utah. Your -- the TelNet application is the first time I've heard about that site.

Postel: Yeah. They actually didn't do very much. One of the things that happened there was, I think, that they were one of the early sites that got a copy of the Tenex operating system, and the work on developing software for the Tenex operating system was done by the other group at BBN. BBN is a big company and has several divisions. One division was doing the IMPs. A separate division was doing this Tenex operating system. And, so the network code to go in Tenex was done by the BBN group doing Tenex, and Utah people just got a copy of all of that. So there wasn't really much point in them being terrifically active in developing protocols.

Pelkey: Now, what happened to UCSB?

Postel: They did some, actually, very good stuff fairly early.

Pelkey: Some guy named White was there?

Postel: Jim White. White has actually been very much involved in the X-400 stuff. The X-400 wouldn't be the way it is if it wasn't for Jim White.

Pelkey: Do you know where Jim is today?

Postel: I think --

Pelkey: Is he with Larry Roberts up at DHL?

Postel: That might -- he went someplace.

Pelkey: But UCLA and BBN and then MIT joined, which was one of the 15 sites, was an active site.

Postel: Well, people from Santa Barbara were involved in some early stuff. People from SRI were fairly -- somewhat involved.

Pelkey: Was Carnegie-Mellon an active site?

Postel: Carnegie-Mellon was not that involved.

Pelkey: And so the real sites were UCLA, MIT, BBN and then, lesser, SRI.

Postel: I mean, one way to do it is sort of go through this index of RFCs and see where the authors are from. And that would give you a clue as to who the most active players --

Pelkey: And if we did that exercise, you --

Postel: I think we'd come up with that list --

Pelkey: From the sense of talking to people that's the sense I get is that's kind of what it was. And then, the TIP sort of came about because, in fact, people started to say: "Wait a minute. I don't have a host or I need to travel or -- "

Postel: Right, I mean, basically, yeah. People said: "Well, I need to access to these computers but I don't have a big timesharing system of my own, and I need to get access to some computing resource, so, can we somehow have a mini-host that does the stuff to make it so that I can do remote terminal access and nothing else." And so they went to BBN and said: "Can you do that," and they said: "Yeah, we can do that and we can actually build it in the same box as the IMP." And that was, that was -- became a TIP. And that became -- that actually became a technical -- led to some technical argument about the way the file transfer protocol was defined in that, well supposing you want to be able to control a file transfer from a terminal on a TIP. So, can you sit at a TIP and sort of connect to the file transfer server on a timesharing system and type commands at it, so that the protocol -- the control protocol for file transfer is an ASCII command language, so that you could access there and type it if you had to. Of course, it would be very painful to do it because --

Pelkey: Two other general kinds of questions, and you've been kind with your time, the first one is: X-25 -- and here, packet switching, in terms of the network, really had gone with ARPA, and yet X-25 kind of developed separately and is very different than ARPA.

Postel: Right.

Pelkey: Do you have any views on that, or how that happened or why that happened or --

Postel: It happened because X-25 is a CCITT thing. CCITT is a organization of telephone companies. So all the people who go to CCITT meetings are with telephone companies, you know, public telephone systems of each country is -- most other countries it's a government agency, but, the telephone company people. And so they think about things in terms of circuit switching and, the fact that it's as much packet as it is is sort of remarkable. And that's why it is that way. But, somebody you should talk to about X-25 and how it came about is Larry Roberts.

Pelkey: Yes, I intend to.

Postel: He was very involved in that.

Pelkey: The other issue is: do you find it strange that, my view is that after -- when I started this off I thought, ARPA, wide area networks, X-25 -- I thought that's what I was going to find. What I find is, ARPA, TCPIP, local area networking. And that the home -- that the intellectual home of a lot that was done with ARPA really I think finds its way into local area network and had a more profound --

Postel: Well, yeah, I think this -- you have to go back to who were the people? What was the culture? The people involved in ARPA networking are from computer science departments. They're from timesharing systems, operating systems, ok. They're not traditional data communications people.

They're not from telephone companies. Right? They view what they get from the telephone company as a wire.

Pelkey: Yeah.

Postel: Ok? And the fact that they have to get it from the telephone company is slightly distasteful, ok? Right?

Pelkey: Objectionable.

Postel: Right. But that's the only source of long wires.

Pelkey: Right.

Postel: And the fact that -- and the telephone companies come out and they say: "Well, ok, now we can do -- we have T-1 multiplexers, and drop and insert" and all this kind of stuff, and its all just sort of a headache for the networking people, data networking people. They say: "Why do I want all that stuff? I just want a wire from here to here and I'm going to multiplex it with packets."

Pelkey: Yeah, yeah. It is interesting. Do you have -- I've kind of exhausted the kinds of questions that --

Postel: Ok. Well, I would like to tell you about something --

Pelkey: Good.

Postel: That probably you haven't even heard of yet.

Pelkey: Thank you.

Postel: There are -- there is a project here at ISI, its been going for some time, called Moses.

Pelkey: Never heard of it.

Postel: Ok. Let me show you my toys. This is not -- several years ago, people said: "We want -- we'd like to do some research on VLSI, and we would actually like to get some chips back." And the problem with making -- getting a chip, is that is costs a lost of money. And what costs a lost of money is the mask. When you make a chip, you don't make a chip, you make a wafer, which is something like this, I'll show you a wafer in a minute. And the expensive part of making a wafer is making the masks for the layer -- you sort of have the mask and you do something so that you put some material down according to that mask and then you have another mask and put some material down. So, these masks, I don't know what this is but it's some sort of piece of glass and its etched with an electron beam microscope. So the set of masks is fairly expensive, and you use a set of masks to do a wafer map. This is a Moses wafer. Now, typical industry wafer would have, you know, a hundred copies of exactly the same. Now, this Moses wafer has lots of different things on it.

Pelkey: My goodness.

Postel: And, so the whole idea here was to figure out how to share the cost of making a wafer amongst people with lots of different -- lots of different --

Pelkey: Low volume, special run-off chips.

Postel: Sort of specialty chips, special chips. So, to make a wafer costs $60,000 or something, ok. And so like if you have a random idea for some new VLSI design and you'd like to see if it works, you got $60,000 to find out if it works? Well, no, but you might have $1,000.

Pelkey: Right.

Postel: So if you can share that with 60 people, it'd be really terrific.

Pelkey: Yep. It would be.

Postel: That's what Moses does.

Pelkey: That's fantastic.

Postel: Ok. Now, but that's -- and then there's a key to the wafers (unintelligible). More to the point of the previous discussion, is how does Moses go about accumulating all these designs and (unintelligible). What they did is they said: "Well, look, we don't want to, you know, have people send us tapes and we don't want people to send us boxes of cards or whatever it is that describe their design." I mean, people sit down, when they design a VLSI circuit, they sit down with some sort of design system, CAD/CAM system or something, and they sit there at a workstation. They type things in or they draw pictures with light pens or something, and at the end of that, a file is written that is a description of that design in some language. Well, it turns out that, over at Cal Tech, some people came up with a language called Cal Tech Intermediate Format, CIF, that is a description language for VLSI designs, ok. So what the Moses project said is: "If you send us a CIF description of your chip, we'll put that together with a whole bunch of other descriptions that other people send us and lay out a wafer, and we'll take care of whatever we have to do to get that manufactured and all the chips sawed off and legs put on them and sent back to you, but the way we want you to send us that CIF description is via electronic mail," ok. So if your on the Arpanet or you have some way of sending us electronic mail, and lots of different ways, we'll -- and we're going to write a programmatic thing that will process the mail. So you send us a message, and we're going to have a little program that runs that notices that we got your message and starts processing it, and its going to send you a little note back saying that we got your description and since you asked for four micron CMOS fabrication, the next time we're going to have a run of that kind of fabrication is this date, so you should expect to hear from us shortly after that.

Pelkey: That's fantastic.

Postel: Or, if you want to make changes, you can make changes up to this date, because we won't have committed before that. And so, if you discover -- you know, you're looking at your design tools, that you actually made a mistake and you want to send in a replacement description, you can do that up to this date. Use this code number when you do that. So there's all this sort of automatic processing of these electronic mail messages to describe these designs. So this is a very interested use of electronic mail to do a -- a very custom business. Basically, this is a, Moses project is a broker for acquiring very customized items.

Pelkey: Right.

Postel: And it’s done all, from the user's point of view, its done all through electronic mail. I mean, he has to, at some point he has to call up and talk to somebody in person and get authorized and, "Yes, we can really do this," and there's the right way of paying the bills exists and stuff like that. In general, for the Moses project, what happens is that the government agency gives ISI a lot of money and says: "These people are allowed to make chips." And so then we have to make sure we go that. So that's the Moses project, and it's a very interesting thing. It's been very successful. It's been running for several years. Over the years they've moved from one technology to another, you know, like four micron to three micron to two micron (unintelligible) CMOS and NMOS and they're trying to get organized to do gallium- arsenide now and all kinds of stuff like this. Ok, now just recently, we've started another project that we call FAST. And the idea of the FAST project is an exploration of just computerized business, computerized commerce. Let's take the same idea that Moses has, but instead of doing this for really weird custom things, let's do it for very standard things.

Pelkey: Right.

Postel: And we'll say: "Well, how do we start this, and, you know, what is the initial community of interest? Well, let's start with the Moses users and let's start with standard electronic parts --

Pelkey: Right.

Postel: So if you have something that you know, some standard electronic part that's in the TI catalog or the Intel catalog or, you know, catalog of a vendor, you send us an electronic mail message that says: "I want to know, I want a quote on a 1234N73 transistor," or whatever the numbers are, I don't know the numbers, and send an electronic mail message to us, our -- the goal is our program will process that automatically, get in contact with some major distributors, get their inventory quantity and price out, send you back an electronic mail message that says: "Here are three quotes. You can get from this guy in two weeks for this price, you can get it from this guy in one week for this other price and the other guys didn't have any," or something like that, ok. And then you can look at that and say: "Ok, well, fast delivery is more important to me than price, so I'll check number 2 and I'll order that. Deliver it to me." And then go back in to him, and say: "Ship that. Drop ship it. Don't ship it to ISI. We don't ever want to see it. Ship it from your warehouse to that address, send us the bill, then we'll send the guy that ordered it a bill."

Pelkey: Right.

Postel: And because ISI, or the FAST project is, processes lots of these orders, it gets a big customer to the vendor, even though it has all these funny addresses that everything gets shipped to. So we get a big discount, because we're big customers. So then we can pass that discount on to the user.

Pelkey: Right.

Postel: Ok. So, computerized commerce. Automated brokering system, via electronic mail. So that's what we're trying to do.

Pelkey: How are you commercializing it?

Postel: We're not -- this is an experiment, a crazy experiment, ok. Well, if it works, then we would expect, I mean if it REALLY works, then probably somebody, maybe some people here, maybe some completely other people we've never heard of will come along and say: "We're going to go run a copy of this as a business," and ISI will stop doing it and it'll be a wonderful business. It remains to be seen whether or not there are enough customers on electronic mail, yet, that, you know, to support a business. It remains to be seen whether or not there's enough margin in the discount to pay for the operations cost of it. So there are a lot of unknowns about whether or not it can be a business, but right now we're doing it as a crazy experiment. But the other thing that's different about this one than the Moses thing, is that the guy who gets the parts has to pay for them --

Pelkey: Right. Yes, yes.

Postel: -- directly, so there's a big accounting thing.

Pelkey: I'm going to take this off line, cause I want to follow up on this in some of the things I'm working on. Let me in conclusion say thank you very much. You've been incredibly helpful. I look forward to listening to this and learning more from the tape.

END OF INTERVIEW

