Oral History of Alexander McKenzie

Interviewed by:
Marc Weber

Recorded June 6, 2011
Ojai, CA

CHM Reference number: X5863.2011

Weber: I'm Marc Weber of the Computer History Museum, and I'm here on June 6, 2011 with Alex McKenzie, pioneer of ARPANET and internetworking, BBN. And thank you for doing this.

McKenzie: You're welcome.

Weber: So let's start with your background. Where were you born and where did you grow up?

McKenzie: Okay. I was born in central Massachusetts in 1940, and my father was involved with the Radiation Laboratory at MIT during the war, and so we lived in Newton, a suburb of Boston. When the war was over, he got a job in New York City and the family moved to Hackensack, New Jersey, and that's where I grew up. Not only did I acquire a New Jersey accent in Hackensack, but then I went to college at Stevens Institute of Technology in Hoboken, New Jersey, and I polished my New Jersey accent there. I went to graduate school at Stanford University in Palo Alto, California. I went there with the idea of working on a PhD in Math, but I hated math at the graduate level and so I started spending all my time in the computer science department and I ended up, after a year and a quarter, with a Master's Degree in Computer Science, from Stanford. I didn't...

Weber: [INAUDIBLE]

McKenzie: Go ahead.

Weber: But going back, what were the subjects that interested you in earlier school, elementary school and high school?

McKenzie: In elementary school, I was interested in and good at pretty much everything, and when I took, in junior high school, some kind of an exam to figure out what I wanted to do with my life, and I think that it showed that I should be either a forest ranger or an agricultural scientist. But then, Sputnik went up and all the adults in my life said, "No. You've got to be an engineer. The country needs engineers. It's your patriotic duty to be an engineer. You're good in math and people good in math should be engineers." So I took their advice and became an engineer. I found, when I applied to engineering schools, that most of their course catalogs said well, in the second half of your freshman year, you can begin to decide what kind of an engineer you want to be. And so I asked all these adults who had told me I wanted to be an engineer, what kind of an engineer the country needed me to be as my patriotic duty, and they said, "Oh, any kind of engineer you want." I said, "That's not much help." So I went to Stevens because it allowed me one elective in the second half of my senior year and that was it. The rest of the curriculum was fixed. So I didn't have to make up my mind for three and a half more years. Then, they kind of shifted the ground under me by introducing a science program in my freshman year, and I decided to switch to that because I was now free of the influence of all those adults who were telling me to be an engineer. And Stevens got its first computer, an IBM 1620 card operated system in my sophomore year, and I took some courses that involved programming that and I found that I loved that and I was pretty good programmer. And so I decided to become a mathematician because at Stevens, the math department owned the computer. At Stanford, I think it was also the math department that owned the computer. So when I discovered that I really didn't like the graduate level math courses, I was easily able to take computer courses.

Weber: Your first exposure to computers was?

McKenzie: It was at Stevens. Probably it was my sophomore year, maybe my junior year. It was an IBM 1620.

Weber: And did you have technical sort of hobbies as a kid or a student? Did you build things? Or were you interested in—

McKenzie: Well, I was – I did. I did build things. But mostly, I would say my hobbies were outdoor stuff, which was undoubtedly why this thing that measured what I wanted to do with my life said I should be a forest ranger or an agricultural scientist.

Weber: And, perhaps, why we're doing this interview in an RV?

McKenzie: Perhaps that's some of it, yes. Right. So I did a lot of hiking and was heavily involved in Boy Scouts, and my parents had a small piece of property with a wooden shack on it, where my mother and my two brothers and I spent the summers. My father came up for a couple of weeks at the beginning of the summer and came back to fetch us at the end of the summer. But he didn't have a long vacation from his job in New York so most of the time he was back in Hackensack. And I spent a lot of time outside. I learned surveying through the Boy Scouts, so that I could find out how big the property that we had actually was. Because nobody knew. And I built a log cabin one summer and a treehouse another summer, things like that.

Weber: So yeah, which part of New Jersey was that?

McKenzie: Hackensack, New Jersey...

Weber: No. But the property?

McKenzie: The property was in New Hampshire and it was near Conway, New Hampshire. I guess, I should mention, also, that when I was a kid, I was very interested in diving, and I built a diving helmet out of an old hot water heater tank when I was about 12. And I acquired a rowboat to put on the lake in New Hampshire so somebody else could pump air to me while I experimented with my diving helmet. Well, people didn't want to pump air to me too much, so that got me interested in scuba, and when I was about 14, I built my first scuba outfit out of war surplus parts from Canal Street [New York City] and a couple of other things like that. So there's a little engineering there, too.

W1: And go-karts. Didn't you build go-karts?

McKenzie: No. That was my brother, Ken.

Weber: And your mother was in the technical field?

McKenzie: My mother graduated from Wellesley but she never worked. My father graduated from Dartmouth as an English major. But he was a radio amateur and he got a job during the depression

running a radio system on the top of Mount Washington, where they had just opened a world weather observatory for the International Geophysical Year. And he happened to have been there when the highest wind recorded on the surface of the earth (until about three years ago) was recorded -- 231 miles an hour. And he was mostly there to keep their radios working so that they could be in communication with the rest of the world and also doing radio propagation experiments with Blue Hill Observatory in Massachusetts, which was just kind of at the outer limits of where you could get good radio reception at the frequency range that they were using. And I'm not a radio expert, so I don't remember what that was.

Weber: So then, when you were exposed to computers, that's really what excited you. So go on from there.

McKenzie: Absolutely. Absolutely. So when I was an undergraduate, I wrote a project in assembly language to do symbolic integration and differentiation. Wasn't very good, certainly not at the level of MATHLAB. But it was 1961 when I did it so things were more backwards then. When I went to Stanford, my master's thesis was making modifications to the ALGOL compiler for the Burroughs 5500, which Stanford had just acquired in my second year there. But then, I wanted to be back in New England. I love New England. I'd done a lot of hiking in New England. We had this family property in New England, and I didn't really like being in California, where you could tell it was winter because the grass turned green. So I wanted to be some place where there was skiing and snow and things like that. So I looked for a job in the Boston area, and my father had some contacts with people who were hardware engineers at Honeywell Data Systems, and I interviewed there and ended up going to work in Wellesley, Mass for Honeywell. And my job there was building Fortran compilers, and I worked at that for about three years, and by then I decided that Honeywell was not going to be doing interesting things. They were not going to be at the forefront of anything. Their business model was to build whatever IBM had already built, only build it five years later, to take advantage of hardware advances so that it would go faster. So I looked around for a place that was doing innovative things, and interviewed a number of places but ended up going to Bolt Beranek and Newman, a consulting company in Cambridge. Because they had built a timesharing system, which was in use in the Massachusetts General Hospital, managing basically the hotel functions of the hospital. And my first job was to document that system, and that way I could learn about it. And so I did that, and then I began doing little consulting jobs on data automation things, and I did that for a couple of years.

And then I took a leave of absence in the Spring of 1970, and my wife and I went to Europe and went camping for six months. And when I came back in November of 1970, this thing that BBN had done in 1969 and 1970, which I was not involved with at all, called the ARPANET, was beginning to take off, and although it had been deployed in the field for almost a year, it still wasn't finished, of course. There were things that needed to be improved and the hardware and software people, who had built it and were trying to improve it, were being hounded by all the universities in other locations that ARPA was making connect to the network. And they wanted to know how it worked and what happened and how about this and what about this question and how about this hardware and what did they have to do to connect. My boss, Frank Heart, an engineer and a wonderful manager, said, "Well, you're just back from this long vacation. You don't have anything to do and you're pretty smart. So why don't you go ask all these questions of the hardware and software people, one more time, and then you can be the interface to the

outside world and they can be left alone." And so that's what I did. I learned not only a little bit how the hardware worked, but a little bit how the software worked and I did some of the documentation and I became the public interface, and that turned out to have been very lucky. Because a few years later, when interest in packet switching exploded around the world, I was the generalist. So I could go off to these wonderful places, see wonderful foreign sights and explain how the ARPANET worked to people who were willing to pay BBN to have me do it. And that was-- I turned that into a career.

Weber: Further back, you had already been working for Frank on the other projects when he switched you over to the ARPANET?

McKenzie: Yes, Frank...

Weber: Okay.

McKenzie: Although when I came to work at BBN, Frank was already there, I didn't know it. I worked for a man named Paul Castleman who was the head of the Hospital Computer Project. It turned out that Frank was his boss and I didn't see Frank for probably the first year. I heard him yelling down the hall every now and then. Because when he's enthusiastic he speaks very loudly. But I had nothing to do with that until I finished my documentation job for the Hospital Project and then Frank began to assign me to do other things.

Weber: So at the point that you switched to the ARPANET, you'd been working directly for him for a while?

McKenzie: Yes. Yes.

Weber: So what was he like to work for?

McKenzie: Oh, Frank was an engineer. Frank was a wonderful engineer. I feel that I was extremely lucky to have worked for him. He was a man with very strong opinions. But he was also completely open-minded, and he would allow the people who worked for him to debate until either he convinced them or they convinced him. And he was perfectly willing to change his mind if he was convinced that another argument was better. Frank, like most of the management at BBN, was strong in engineering and not so great in business, and probably because it was a wonderful place for people who were engineers, it's why we never really became a humongous commercial success with any of the new technologies we developed. I mentioned Frank was enthusiastic and yelling down the hall a lot of the time. Several years later, when I knew him better and we were on a plane trip together somewhere, I said to him, "You know, Frank, the first year or so that I was around you, it made me very nervous to hear you yelling all the time." And Frank said, on the plane, [YELLING] "Yell? Me yell? I don't yell." [LAUGHS] But he didn't yell when he was angry. When he was angry, he got very, very quiet. Frank had an amazing engineering talent. You could work on a proposal for a couple of months, which was typically the amount of time that one had to work on a proposal, and put together 100 pages of engineering detail, and give it to Frank to review before it got sent off to the-- as part of a bid. And he'd have it overnight and he'd come back and he'd put his finger on the one paragraph in the middle of those hundred pages where you had, kind of, tried to

blow smoke a little because you weren't really sure what you were going to do about that. And he could do that, time after time after time.

Weber: And so the team there was well established by the time...

McKenzie: The team was certainly well established, yes. The contracts-- the proposal was written in 1968, in the second half of 1968, and Frank put together the team to do the proposal that ended up implementing the system. So they had-- when I came back from this leave of absence, it was November of 1970. The team had been together for more than two years, designing and implementing the system. And I was not added to them because they needed-- any kind of technical help, I was added to the team to keep the outside world away from them and let them get about their business.

Weber: But did they make you feel part of the team?

McKenzie: Oh, yes. I knew most of the people who were on the team even from-- socially, from 1968 and absolutely, they were thrilled that there was going to be somebody who would save them from having to answer the same questions over and over again. So they were eager to have me learn everything that they could teach me and then, stay out of their hair, and that's what I tried to do.

Weber: So I have interviewed a number of the people on that team, individually. But tell me what was sort of the feeling of that group like? What were they like to be around?

McKenzie: Well, they were all very playful. They all had very broad interests and they were all very, very smart. Bernie Cosell was a debugger. He could take anybody's code and within a short amount of time, he could do debugging on it in real time operation. He'd gotten a lot of experience in real time operation in the Hospital Computer Project. Will Crowther was extremely smart, very good at inventing new algorithms and seeing how things could work and coding very, very tight code. Will, unfortunately, never really finished anything. He kind of lost interest when it was 95 percent done or so and wanted to go on to something else exciting. So Bernie was the perfect person to work with him to do the final debugging. But he didn't have to mess with Crowther's algorithms. They were great. Dave Walden was another guy who did software. He was very practical. He produced a lot of code very quickly. He had worked with Will Crowther at Lincoln Lab and he knew and respected Will. So he was able to follow along and implement stuff that Will designed and make it work perfectly with Will. The hardware team, I didn't know quite so well. They were headed by Severo Ornstein, a very bright hardware guy. He had other people working with him who I knew even less well. Kind of, the second in command managing the project was a guy named Hawley Rising. He had worked with Frank for a long time. Frank and he trusted each other and knew each other's habits very well. All of these people were extremely friendly, nice people to work with, and a lot of them had interesting side hobbies that they brought into BBN. Dave got very interested in juggling and he introduced juggling to BBN, I believe. He eventually became the Secretary and then the President of the International Juggling Association, in his spare time. Will Crowther loved to hack with computers. He was a caver. He and his wife and others had explored Mammoth Cave, and he used the IMP hardware and a flatbed plotter, off hours at BBN, to draw maps of the cave in three dimensions. Eventually, he wrote an amusing program for his daughter, which grew into the Adventure Game, based on his caving experiences, and a lot of puzzles that were peculiar to him. Bernie Cosell was an amateur

guitarist. You could be going to BBN at any hour of the day or night and hear Bernie sitting in his office playing the guitar. He eventually became a weaver and finally decided to get out of programming, and bought a sheep farm in Virginia, and is now a sheep farmer part time and a programmer part time.

Weber: And so you mentioned you would socialize with these people. Was it a cohesive group?

McKenzie: Yes. There were, during the late '60s and early '70s, a lot of marriages and a lot of births among the people that worked at BBN, and a lot of parties, not sponsored by the company. And I think, of course, everybody had favorite associates and people they were less interested in. But I would say yes, it was a pretty cohesive group. And certainly, the group of people who were called the IMP guys was a very cohesive group.

Weber: And you were already married by now?

McKenzie: I was married in 1968.

Weber: And you had met in...

McKenzie: Met at Honeywell.

Weber: Okay.

McKenzie: I met my wife at Honeywell.

Weber: So were you working crazy hours at this point or not?

McKenzie: Well, there were lots of long days. There were some people who chose to work in the middle of the night, either because they could get machine time then or because they were just night people, and whatever. I think that I lived a fairly normal 9:00 to 5:00 or 8:00 to 5:00 or 8:00 to 6:00 kind of working existence. My first child was born in 1972 and, maybe, I spent more time at home after that. It's not clear that I really did. As I said, you could come into work at almost any hour and find Cosell here playing his guitar, and there were other people, probably not so much among the IMP guys, who worked crazy hours. Will was married and had a couple of kids, and I think that he kept pretty normal hours. Dave Walden certainly tried to keep normal hours and Frank and Hawley did. The hardware guys, again, I'm not so sure, because I was not quite so close to them. But I think it was not-- there were not lots of people working all night or anything like that. There was a very limited amount of time to build the IMP. The contract was awarded to BBN on January 2nd, I guess – or maybe 1st – of 1969, and the first working IMP, a thing that had never been built before in history, was supposed to be delivered on September 1 to UCLA. So even though a good bit of design had been done during the proposal process, there was lots and lots that had to be done before any valid working machine could be deployed to the field. And then, another one was due on the first of October, another on the 1st of November and the 4th and then the 1st of December. That was the four machines that BBN was initially contracted to provide was those four, and they went in on schedule and they worked in the field. So, obviously, the people who worked on developing the machine had to work long hours a lot of the time. But I don't think ever-- there were never 24-, 48-, or, you know, 72-hour work shifts, anything like that.

Weber: And you came in November of '70, right? So about almost a year after the...

McKenzie: …after the first.

Weber: Right.

McKenzie: Well, a little more than a year after the first IMP was deployed, yes.

Weber: So they were still cranking out IMPs?

McKenzie: ARPA extended the contract as soon as they saw that a working system had been delivered to, and this was part of the-- an option on the original contract, which ARPA exercised, extended the network to 19 nodes and they got produced at the rate of about 1 a month. So there were IMPs being cranked out of BBN and delivered to the field in 1970, when I joined the project. And in fact, I think, probably the project went on delivering IMPs until about 1975, maybe at a slightly slower rate sometimes. But a lot of machines were passed through BBN. So there was a working network in the field when I joined. Of course, at the beginning, when ARPA gave this contract to BBN, it was a research and development contract and nothing like this had ever been done before. Nobody expected it to be delivered on time. It was a government contract, after all. Frank expected it to be delivered on time. That's the kind of manager he was, and so it was delivered on time. The places that ARPA had told, "We're going to build this network and then connect you up to it," figured it would be at least a few years before ARPA actually made good on this threat.

A lot of the sites that ARPA said were going to be connected to the network actively didn't want to be connected to the network. They felt that this was going to bring other people in to interfere with their work on their computer and take time away from their projects, and they didn't want to do it. Not every site was like that, of course. The first site at UCLA, the Network Measurement Center, [Leonard] Kleinrock, [Jon] Postel, [Steve] Crocker, and a lot of those people were very interested in getting access to the network. The second site, SRI, Doug Englebart's group in machine augmented human intelligence was very interested in having a network to try and spread the gospel of what they were doing. But other sites were not so enthusiastic. I don't think MIT Multics project was particularly interested. I don't think-- even the BBN computer center probably was not all that interested in being connected up to this network. It looked like a lot of hassle, and what was going to be the good of it? So nobody was getting ready for the network when it was first being developed. BBN had, under ARPA contract, specified how you connected a host computer, your current mainframe, up to an IMP and how they could exchange information. But nobody had worked particularly on how two hosts with different operating systems, different word sizes, different implementation philosophies – what were they going to say to each other and how were they going to say it? So when I got involved in 1970, there was a thing called the Network Working Group that was trying to develop protocols for host-to-host connection and for terminal access. But they hadn't made a lot of progress by November of 1970, and one of my jobs as the public interface of the BBN IMP guys, was to go to the Network Working Group meetings and explain how the network worked to anybody who had questions. And as a result, I got involved in the design of those protocols. Now, I'm not an inventor. I didn't invent anything in any of the protocols. But I'm a good writer and I'm good at seeing other people's ideas and how they fit together. So I ended up playing a reasonably important role in the Network

Working Group, in documenting the protocols and putting together bits and pieces of other people's ideas, in a way that made them work. And so, for example, I was the author of the document that was the specification of the first host-to-host protocol for the network, a protocol now called NCP, or Network Control Protocol. I was involved in the committees that did the Telnet protocol and in the file transfer protocol and a bunch of other things that didn't get so much attention and some of them which never even got finished, like a graphics protocol that had a lot of energy going to it for a while, but never turned into anything. Because graphics devices were just so dissimilar we couldn't figure out any common ground to build on.

Weber: Any others that didn't make it that were...

McKenzie: There was a network remote job entry protocol for submitting batch jobs to some of the big number crunching computers that were on the network. In particular, the UCLA 360/91 was envisioned to be a big number cruncher and the ILLIAC IV when it got finished (built and connected to the network) was envisioned to be another of these. But it's fair to say that most of the people in the Network Working Group were more involved with timesharing systems than they were with batch systems, and probably approached building the network job entry protocol too elegantly. And people who had batch jobs to submit, ended up, I think, universally implementing an interim protocol that the UCLA 360/91 guys under Bob Braden had thrown together while the committee was working on the “real thing”. And so I don't think the real thing ever got implemented. Also, in 1971, roughly, ARPA extended the BBN contract to build a terminal multiplexing device. They wanted to begin to allow people to use computers on the network without having a host of their own, and so they envisioned adding to the facilities of an IMP so it could also handle terminals. BBN built a thing called the Terminal IMP, or TIP. It could handle up to 63 terminals. It could handle a wide variety, but basically asynchronous terminals. Teletypes, and TI silent 700s were just coming in, and there began to be glass teletypes as we called them in those days, a primitive display...

Weber: [INAUDIBLE]

McKenzie: ...with a character generator on it. And that required a protocol for terminal handling that had to be implemented in a TIP, which had a very limited memory. An IMP was 16k, a TIP was 32k. That meant 16K two-character words, an eye drop, by today's standards, and that had to include the terminal buffering and the network buffering and all of that. So part of my job in the Network Working Group was to make the Telnet protocol, which was how a terminal connected to a host, be something small enough that a TIP could implement it, and I think it's fair to say that a lot of people in the rest of the network community resented that. Not me personally, but just the idea that they were all big hosts, after all, and they could afford to do fancy things with terminals, and we were saying, "Oh, we can't have that in the protocol. It'll be too hard to implement." But, eventually, it got implemented.

Weber: But the constraint was you wanted to keep the cost of the TIP reasonable; right?

McKenzie: The constraint was that – well, yes. The ARPA contract was to add to the IMP, to also give it terminal handling capabilities. And because the machine was absolutely limited to 32k 16-bit words, there really was – oh, and because the BBN philosophy under Frank was absolutely no moving storage devices

on an IMP. Because then we couldn't guarantee the reliability, and reliability was very important for these machines that were just sent out to the field to be taken care of – to not be taken care of. To be taken care of remotely. So we had no rotating storage, no tapes, on the TIP. We had 16k words of 16-bit memory. The program and the buffers all had to fit in that so it had to be pretty simple. And, even with all of this, I would say the network probably was developing slowly. And Larry Roberts at ARPA decided to exert a little more pressure on us all by scheduling a demonstration of the network in October of 1972 at the International Computer Conference in Washington, DC. And he announced that that was going to happen and we were going to have a ballroom full of terminals connected to a TIP, all talking to network hosts, and every network host was required to make interesting programs available to the conference participants. And that certainly got everybody's attention. And-- go ahead...

Weber: And before you get that it's a key conference, you were – okay, so you're doing the host-to-host. Were you also – and I mean, the NOC [Network Operations Center] was going and you later headed that. But at this point, were you responsible, at all, for the operations?

McKenzie: Certainly, by '72 I was. In 1970, I was not. I don't know when that happened. Maybe, late '71, possibly. Well, go ahead and let me come back to that. And just say this demonstration -- publicly announced – put everybody on notice that we better have the protocol designs finished and implemented and experimented with enough to know that a demonstration was going to be possible before we all got down there to Washington to help with the demonstration. That was an important forcing function and it did make people decide to give up on academic arguments, which of course, there always are about what's the right way to do things, and to settle on something. And I and Jon Postel and Steve Crocker, in particular, were busy trying to get documents written, and demanding that people hurry up and make up their minds so we'd know what to write down, that sort of thing. So that was very important for the maturing of the network.

One of the things that I discovered at the Network Working Group, quite quickly, was that although from BBN's point of view this was a research project in communication, ARPA had told all of these host sites, "You're going to start using this for your work." And it had started telling the people who were going to get TIPs, “We're not going to buy you a computer for this research we're sponsoring. You're going to get your computer resources over the network." And so most of the people that I interacted with in the Network Working Group wanted to consider the network to be a utility, like the telephone system or the electric system. They didn't want it to be as flaky as a typical research computer system was in those days. The original IMP team had built a lot of diagnostic routines and software into the IMP and built some diagnostic hardware into the IMP's interfaces so that there was a fair amount that could be done to an IMP remotely. BBN, from the beginning, had – well, from the fifth node, which was located at BBN. From that time, we were connected in – BBN was connected into the rest of the network and so had network access to the internals of the other IMP's that were in the field, and that BBN IMP was probably installed in January or February of 1970...

Weber: '70?

McKenzie: Of 1970, yes, I think so. Shortly after the first four nodes. So by the time I got back from my vacation, there was a Network Operation Center at BBN, and it basically consisted of the other IMP's in the network sending reports to the teletype at BBN and the teletype printed the reports out in some terse fashion. And in the morning, the engineers or the software people would go in and look and say, "Oh, yeah, something broke here. I guess, it's still down," and so forth. And after I had been in the Network Working Group for a while – six or eight months, so it's probably the middle of '71 – I began saying to Frank and the rest of the team, "You know, this network is just not going to be accepted by the people out there if it breaks as much as it does and if we don't get it going more quickly when it does break. And so I think part of my job, as representing BBN to the outside world, should be pushing on BBN to do something about this." After not too long hearing me complaining, Frank said, "Well, okay. You can be in charge of the Network Operation Center and you can do it how you want, and stop bothering the rest of us."

So I did become the leader of the Network Operation Center. We used people who were already on shift as computer operators, around the clock, for monitoring and controlling the network, and we got the implementers to fix things up so that these operators would have a set of procedures when they saw that an IMP at UCLA wasn't able to communicate over the communication circuit to the IMP at SRI. They would call up the phone company-- of course, the phone company didn't believe that anyone in Cambridge could be reporting on troubles with a circuit between San Francisco and Los Angeles. But after a while, we convinced them, and they would go work on the problem. Or if an IMP went down, the operators had a cookbook of procedures so they could call a janitor or a graduate student or whoever happened to be around in the middle of the night at a site, and talk them through the procedure of getting the IMP restarted. IMPs were always able to reload copies of the program from a neighbor IMP, so all they had-- all the people on site had to do was push the buttons to get a bootstrap program loaded from paper tape and the IMP would take off. So there were a lot of mechanisms in the IMP to allow remote management. I don't think that the programmers envisioned that this would be turned over to a bunch of basically operators. They thought they, the programmers, would be continuing to run this. But we did get the NOC to be treating the network as a utility and we got reliability up to probably about 98 percent or so for the IMPs, which was probably roughly 5 percent better than the reliability of any of the hosts, even the best of the hosts, connected to the network that were trying to provide service. Places like UCLA and UCSB were the early service providers. So that was probably good enough. It's not near as good as what you expect from the Internet today, but it was certainly way above what most computer facilities were offering in 1971, 1972. So we gradually expanded the functions of the NOC to a point where we did new software releases, as often as once a week. Although-- we had scheduled once a week, I think it was Tuesdays from 6:00 a.m. to 9:00 a.m. Eastern Time was reserved for network maintenance. And although we did not always use that time, we did manage new software releases from the Network Operation Center and we did debugging of systems and it ran pretty well.

There's one incident that comes to mind, where the NOC got a lot of criticism for the management of the node in Hawaii. There was an IMP installed in Hawaii probably early '73, I'm guessing, although I'm not positive, and it was at the University of Hawaii. But ARPA had promised to the Commander-In-Chief of the Pacific operations that there was this wonderful new communication tool that he should do experiments with. And so CINCPAC [Commander in Chief Pacific] fleet was planning some demo day of

ARPANET demos on some particular day, and the principal investigator at the University of Hawaii said, "You know, you guys are terrible at maintaining this TIP. And it goes down every day. Every day it goes down, and if that happens on the day when this big demo was going to happen, you know, these people are going to laugh at us and laugh at you and laugh at ARPA, and you'd better get this fixed." And in fact, the PI at University of Hawaii brought this up at a principal investigators' meeting, sponsored by ARPA, and basically was yelling at my boss, Frank, for not doing a good job. And Frank, of course, came back and said, "Well, what's the story, Alex? You're in charge. Why is this so bad?" Well, it was true that every day at about the same time, give or take an hour, the Hawaii IMP--TIP, actually. It was a TIP--would go down, and the demo date for CINCPAC was getting closer and closer, and finally, I said, "I don't know what else to do. I took one of the network operators." I said, "I want a volunteer to go down to Zayre's and buy an air mattress and a sleeping bag and go home and get some underwear and jeans and go to Hawaii and sleep next to the TIP so we can find out why it's going down. And so, at least, if it goes down, we'll have somebody right there to bring it back up again." A guy named Dick McDonald volunteered and he went down to-- off to Zayre's and got the sleeping bag and the air mattress and went off to Hawaii, and about two hours after his plane took off from Logan Airport in Boston, the TIP went down in Hawaii, right on schedule you might say. And the operators called the site and they got a graduate student who was often working in that room, and he said, "Oh, yeah, something bad happened. I dropped my screwdriver in, and a lot of sparks came out." They [the NOC operators] said, "What? What were you doing with your screwdriver?" "Oh," he said, "I'm working on the host interface out here and we need a six-volt regulated power supply and the TIP has one. So I've been tapping off of that. But today, when I went to connect my wires up, I dropped my screwdriver and there were a lot of sparks that came out of the TIP." So they said, "How long have you been doing this?" "Oh, well, I've been doing this ever since the TIP got here. About the same time every day, I come in and start up my project." So Frank got to call the PI at Hawaii back, while Dick was still in the air, and say, "We've solved your problem, if only you can keep this graduate student out of the computer room." And he [Dick McDonald] got there, and I think I made him sleep by the TIP at night but he could have his days at the beach until the demo was over. The TIP, of course, never went down again.

Weber: So Norm Abramson was the...

McKenzie: This was not Norm.

Weber: Oh, no?

McKenzie: It was not Norm Abramson. It was Frank Kuo.

Weber: Oh, okay. Yeah. And the grad student, do you want to name?

McKenzie: I don't remember his name. I don't remember his name.

Weber: And when you had arrived, I mean, the NOC was not-- there was no one else that was a formal head of the NOC, before you? It was just whoever was...

McKenzie: I think that's pretty much true.

Weber: I mean, it existed officially in the ARPA contract. But there was no chain of command, per se?

McKenzie: Well, I think Hawley Rising was kind of in charge, but Hawley Rising was-- I mean, it was only one of many things that he had to do and he continued to view the whole thing as a research project. You know, that we're doing R and D; we're not being a telephone company. And I would've had the same viewpoint, except I was out there exposed to all these people bitching at me at the Network Working Group meetings. So I had to adopt a different point of view for self-preservation. And Hawley had certainly no interest in trying to change the way the NOC was run into a utility type operation. It was not-- he was an engineer, a hardware engineer, and that's not what he thought was interesting. So when Frank put me in charge of it, I wasn't really displacing anybody else who thought it was their job. In that sense, there was no chain of command. I mean, there was some chain of command and probably it was Frank, Hawley-- I don't know who after that.

Weber: So effectively, you were the first network...

McKenzie: Essentially, I was the first, yes.

Weber: The first network administrator anywhere. Because...

McKenzie: For packet switching networks...

Weber: True.

McKenzie: Plenty of other networks...

Weber: Yeah, a hundred and something years before then.

McKenzie: Yeah, that's right.

Weber: Okay. And one thing I forgot to ask you earlier. When you were switched over to network, effectively by Frank, was this something that interested you before or was it neutral? I mean, did networking have any particular meaning to you?

McKenzie: No. It had no particular meaning to me. It wasn't clear how long this job was going to last. I'd had a number of other short-term jobs that Frank had given me, and it seemed like it could very likely be another short-term job. I hadn't paid a lot of attention to what was going on with the network when the IMP guys were building it in '69 and deploying it in early '70, before I went off on my leave of absence. So no, it didn't mean much to me, one way or another. I trusted Frank to give me the most interesting jobs he had that he thought I could do. So I'd try this for a while and if it didn't work out-- so I was very lucky. No planning, at all. Just luck.

Weber: And the Network Working Group, tell me a bit about some of the main personalities and the bigger conflicts.

McKenzie: Well, <sigh>...

Weber: What was it like working with them?

McKenzie: The Network Working Group consisted of a representative or two from each place that had a host. And it spanned the gamut from people like Bob Braden, who was trying to run a service center at UCLA with the 360/91 (he was not in charge of the center, but he was their rep to the Network Working Group) to people who were working on research computers that could be up or down or, you know, modified on a day-by-day basis. So there was a wide range of experience and interests and professional background of the people. But everybody, as far as I can remember, was smart. And while there was a great tendency to take an academic point of view and argue about what was ideally correct or what we might do best, how best to do something, when Roberts came around and put pressure on to get something done by scheduling the demo in October of '72, then everybody buckled down and was willing to make compromises. Some of the notable personalities involved in the group in the early days were, certainly, Steve Crocker who started it, and Jon Postel who worked with Steve at UCLA and was a very smart guy and very inventive, unlike myself. Mike Padlipsky who was a very humorous guy, although sometimes frustratingly pedantic, who represented the Multics project.

And maybe the biggest battle that I can remember from all of the things that we did in the Network Working Group was the battle in designing the Telnet protocol (which, to repeat, is to attach terminals to a host computer) between the hosts and terminals which thought line-at-a-time, and the hosts and terminals that thought character-at-a-time. Maybe the most vocal of the line-at-a-time people was Padlipsky, representing Multics. And there were quite a number of character-at-a-time people, but maybe among the most vocal were people representing Tenex that was developed at BBN. The character-at-a-time people had the view that they were supporting a thing like a teletype. Every time you pressed a key on a teletype, it transmitted the character to the host. Every time the host transmitted a character to the teletype, it printed on the printer. There was no direct connection, at all, between the keyboard and the printer at the terminal. So everything went through the host. The people – hosts that did character-at-a-time things did character-at-a-time editing. Tenex had command completion, where as soon as you typed enough characters to identify a unique command at the level you were operating, the system would fill in the rest of the characters of the command for you, and things like that, and expected a very dynamic set of interactions between the terminal user and the computer. Line-at-a-time systems certainly, I think, almost all of the IBM systems, except maybe the timesharing system at UCSB, were line-at-a-time systems, and the Multics system was a line-at-a-time system. They didn't envision using a teletype. They envisioned using a Selectric typewriter, where the keyboard was directly linked to the print mechanism and as a user typed characters, they went to the printer and they went into a buffer. And when the user typed a new line character, then the buffer was emptied to the computer. So the computer had no opportunity to see what the user was doing until the user completed a line. Well, to make either the Tenex oriented teletype talk to Multics, or the Multics oriented IBM Selectric terminal talk to a Tenex, was a tremendously difficult job. And it was probably the most anger-provoking, bitterly contested, set of design issues that I can remember from the Network Working Group. Everything else was a piece of cake, compared to how you were going to make these things work. Yet, eventually, ways were found to make them work, and that is due to the essential good will and engineering competence and desire to get done on the part of the people in the Network Working Group.

Weber: And it was finished by approximately when?

McKenzie: Well, it was certainly – Telnet was in use at the October '72 demo. So I imagine, probably, roughly, the middle of '70-- oh, the end of '71, probably, is-- we can look up the dates. But from memory of something 40 years ago, that's about as close as I can come.

Weber: And how often did you meet in person, the Network Working Group?

McKenzie: Well, Network Working Group only met in person because there was no networking at that time. And we, at least, met in conjunction with the spring and fall joint computer conferences, and my guess is that the full Network Working Group met twice a year in those early years, and that subgroups responsible for doing the design of a particular protocol met more often. Probably, no meetings happened more frequently than one every two months. My guess was, maybe, four or five times a year for a meeting, if you were on a subgroup.

Okay. At the 1972 demonstration of the ARPANET in Washington, another quite important thing happened, which was that Larry Roberts instigated the gathering together of people who were involved with all of the different [packet] networks, the ARPANET being the biggest extant network. But there were other network experiments going on and there were additional people involved with these networks to get together and talk about the next step, which was getting these networks interconnected. And the people there were receptive to this idea and formed a group, which became known as INWG, or the Internet Network Working Group, or sometimes the International Network Working Group. And principle people involved were representatives from ARPANET and from the CYCLADES network in France, under Louis Pouzin, and the NPL network in the UK, under Derek Barber, and the European Informatics Network, which was just being planned at that time-- I'm sorry. I said NPL was under the direction of Derek Barber, but it was Donald Davies, of course, and Derek Barber was the director of EIN. And I think, probably, the French PTT, the British Post Office, the Canadian Bell System, and a few other people who were interested in this new packet networking, were there. Larry volunteered for Vint Cerf to become the Chairman, and said that he, Larry, would support Vint's involvement through ARPA contracts. Vint, I believe, was a professor at Stanford at that time and just gotten there, really. And Larry also volunteered the Network Information Center at SRI to disseminate documentation developed by INWG.

Louis Pouzin was influential in getting INWG adopted, as it were, by the International Federation for Information Processing. At that time, they had a Technical Committee 6 on data communications, which was chaired by Alex Curran of Bell Canada or Bell Northern Research. I don't remember which. And Alex Curran and Louis got IFIP made into a working group under Technical Committee 6 of IFIP. Because IFIP is an international organization, it has standing to present information to the CCITT, which makes telephony standards, and to the International Standards Organization, which makes other international standards, as technical experts. And so our-- the adoption of INWG by IFIP gave INWG standing to go to CCITT meetings and make presentations and get their documentation and see what they were doing, and the same with ISO. At that time, in 1972, I don't believe there was any ISO or CCITT activity going on, with regard to packet networking. But there was certainly the concern that it might be going on soon, and in fact, before long, CCITT began working on the standard that turned into X.25, which was for interfacing a host computer to a packet network.

So in any case, there was this group formed at the '72 conference. Several documents were exchanged, and then in 1973, in June or July, in conjunction with the National Computer Conference in New York City, INWG met again and a subcommittee of INWG got together and, basically, put together the first draft of a protocol for allowing host-to-host communication over an interconnected set of packet networks, where it was believed by INWG that each of the packet networks would work differently, have different message sizes, different timing constraints, different degrees of reliability, and all that sort of thing. And this protocol was based on the idea that a gateway would be a small host interconnecting any two of these networks and it would be able to do minimal header processing of a host-to-host message, and this is what eventually developed into TCP and IP. So that was a very Important meeting happening at the '72 conference, even though it had nothing-- it was not specifically related to the '72 conference, and credit to Larry Roberts for having pushed on everybody to have this meeting and get organized.

Weber: So it was like a parallel meeting...

McKenzie: Yes. Yes. Yeah, an evening or two evenings, we spent meeting when everything else was done.

Weber: And before we move on, tell me your experience at the conference itself.

McKenzie: Well, I spent almost no time at the conference because my job was to be in the demo room, ready at a moment's notice to get the Network Operation Center back at BBN to do things, if things needed to be done. And, indeed, we did make some modifications. Experimentally, we made a modification to the IMP to IMP routing protocol to suppress or to take longer to declare circuits going into -- the IMP to IMP circuits going into the ballroom to be unusable, and it was a big mistake, and we got-- we experimented a little bit live with the first flood of traffic from the conference, and we saw that it wasn't working and we took that patch back out again and went back to the old tried and true routing, which worked better. Not perfectly, but better. I was involved a lot with getting terminals connected up to the TIP. Bob Kahn had been given the job, by Larry Roberts, of organizing this demo. Bob and Al Vezza who was at that time at MIT, went off to a whole bunch of terminal manufacturers and got them each to donate a terminal for the demonstration. Many of these terminals, none of us had ever seen before. The TIP, as I said earlier, had a very limited capacity for dealing with the vagaries of unusual terminals. So we got all, or at least almost all, of the terminals to work with the TIP and it all went pretty well. The demo itself put a huge load on the ARPANET. There was more traffic in those days-- in fact, there was more traffic that month, even though the demo only went on for a couple of days, than there had ever been before in the ARPANET and it all ran pretty well. Except, there was one network crash in the middle of the demo, which happened to be when Bob Metcalfe was showing the bigwigs from AT&T around, and AT&T went away convinced that packet switching was a failure and would never work. Bob was very fond of telling that story and remembers vividly how embarrassed he was and how angry he was at the AT&T people for making this judgment.

Weber: And the turtle, there was the...

McKenzie: Oh, well, yes. There was-- the Terminal IMP allowed one port of the-- or one-- a terminal connected to one port of the Terminal IMP to direct output to another port, and this was primarily

envisioned for use by having a line printer connected to some port of a TIP and having an interactive terminal direct its output to the line printer. But this was also-- this facility was also used by a project called the Logo Project, which developed a language for use in teaching very young children about the ideas of computation and recursion and subroutines. And it did this by having children build elementary programs for controlling a device that crawled around on the floor, and it was called a turtle. And, actually, the turtle had a pen in the middle of it. So in addition to telling it to go forwards or turn right or left, you could tell it to put the pen down or pen up, and the idea was to get children to program the turtle to draw pictures. And at some point-- so the Logo Project was there and there was a demo of the terminal and some interactive terminal redirected its output to the turtle, which had a radio receiver on it and, all of a sudden, the turtle was going berserk, doing strange things, forward and backward, and not behaving, at all. And a little investigation disclosed that a line printer output had been diverted to the port that the turtle was connected to, instead of diverted to the port that the line printer was connected to. So the turtle was getting line printer output and trying to interpret it as turtle commands and didn't make any sense, at all. I don't know if the turtle output went to the line printer because that was less visible. But...

Weber: What was the atmosphere like at this conference? Were people very excited by what was going on?

McKenzie: Well, certainly, all of us who were there-- I mean, there was somebody there from every host site on ARPANET, supervising or willing to help with demonstrations of that site, for anybody who wanted to use their demo. And we were all highly enthusiastic. We were running on adrenaline and we were very enthusiastic, and except for the one crash when the AT&T bigwigs were present, people generally, I think, came in to that ballroom skeptical and went away believing that they'd seen something important, that they'd been present for something important. None of us, of course, could've ever guessed in 1972 what the Internet would be like in 2011. But there was a big feeling, I believe, on the part of most of the people who came to the conference, that they came in skeptical and they went away believing they'd seen something that really had something behind it. That it might actually work, in spite of the naysayers and the fact that it really was a small research project. So that, of course, recharged the batteries of all of us who were giving the demos and made us feel even better. Seeing people coming in, kind of looking around skeptically and going out smiling and enthusiastic made us feel good, too, and that redoubled the energy in the room. There was also a slideshow playing in the hall that was on continuously and showed a carousel full of slides with a tape recording accompaniment describing what you were-- as a soundtrack, basically, describing what the demo was about and what packet switching was about. I believe that there was a movie being made by – under Bob Kahn's supervision, which I believe was supposed to have been finished for the conference, and which was supposed to be playing in an auditorium somewhere that had dramatic music and moving images instead of slides, and a lot of other things. But it's my recollection that if that movie was – well, I know the movie was not finished in time for the conference. So my recollection is it was intended to be and it was just late. Perhaps, I'm wrong about that. I know the movie was finished at some point. It's available on the web now.

Weber: Heralds of the information...

McKenzie: Yes. It's something like *Heralds of Resource Sharing*...

Weber: Heralds of-- yes, yes, yes, *Heralds of Resource Sharing*.

McKenzie: Yes.

Weber: So that was made for that conference?

McKenzie: It's my belief that it was made for that conference, but it wasn't finished in time to actually be there, and that the slide show was put together as a stop gap.

Weber: But as far as you know, there was no footage shot of the conference itself?

McKenzie: I'm not aware of any, no.

Weber: I asked Bob Kahn...

McKenzie: No. Personal video cameras and cell phones that can take movies were not thought of in those days, and I don't believe anybody brought in a...

Weber: <inaudible>

McKenzie: ...bigger camera capable of-- with enough illumination to light things up and all of that. I don't believe so.

Weber: <inaudible>

McKenzie: I've never heard of it, anyhow.

Weber: No 8-millimeter home movies, you know of...

McKenzie: If any got done, I never heard of them.

Weber: So go back, then, to the International Networking...

McKenzie: Okay. Well, as I said, the INWG, or as it later became known, IFIP Working Group 6.1, met in June or July of '73 and drafted a protocol for Internet working, was documented in a document distributed by Vint Cerf, who was one of the people at the meeting, in late July. Then, he and some of his graduate students at Stanford and Bob Kahn worked on a-- or, I guess, and a couple of guys from the University of Illinois worked on refining that and produced a paper, which actually many people would consider to be the first real draft of TCP/IP specifically, which Vint and Bob presented at a conference in Sussex, England in September of '73. I think it was '73. September of '73. Heather was there for that one, too. But not at the conference and she didn't develop any teeth, so that was okay. There was an INWG meeting at that-- it was a NATO Institute for Technical Subjects focusing on data communications, or something like that. I don't remember exactly the title.

Weber: And Peter Kirstein was one of the people...

McKenzie: Peter Kirstein was certainly an organizer of it, yes. Right. It was held in Brighton, England. My wife and my daughter went, hoping for some beach weather, but September in England is not great beach weather, and the British like their beaches covered with pebbles rather than sand. So it's not even comfortable to lie there and look out to sea. So the two of them went off to Spain, leaving me in the conference, and we did have another INWG meeting there. Then, in early 1974, people working with Louis Pouzin, particularly Hubert Zimmerman and Michel Elie, produced a draft also based on the July '73 INWG design. But with a different emphasis focused more on bigger transmission blocks. TCP and IP used a transmission unit of 8 bits. And the Zimmerman and Elie proposal used a concept of letters, which were multiple blocks of about 255 characters each. These two proposals were debated at the next couple of INWG meetings, or maybe just the next one INWG meeting, and it looked like at a meeting there was never going to be a conclusion. Because too many new people came to each meeting with different ideas or needed to be walked through the whole design process again and rethink old thoughts that really had been pretty much settled. So it was decided by Vint and Louis, I guess, that there would be a vote, and it would be by letter, not in a meeting, to pick one of these two proposals and make it the official IFIP Working Group 6.1 presentation to CCITT, in particular, which was working on X25 and maybe the ISO, if they got interested. (The open systems interconnection project had not yet been started at that time.)

It seemed to me like it was a waste of time to try and choose between these two proposals and that it would be better if the two factions could be brought into agreement. So I drafted a proposal that was a combination of the two, and sent that around for comment, and then I worked on lobbying Cerf and Kahn on their side and Zimmerman and Elie on their side, and the people from the European Informatics Network, which was pledged to use whatever protocol was decided on, to try and get together and come up with one standard that we could all agree on. We could use my compromise proposal as a basis or we could start anywhere. And there was a meeting in London in September, I believe, probably of '74, consisting of Hubert Zimmerman, Vint Cerf, Roger Scantlebury from NPL/EIN, and myself, and I was really only there to take notes and write down the conclusions. As I said earlier, I didn't really invent any protocols, at all. I just could write them down well. And we decided we would meet either until we had a compromise everyone could agree to or until we decided that wasn't going to be possible. And we very quickly decided it would be possible, and in probably less than a week of meetings, arrived at a compromise, which got written up.

And then, when that happened, it was decided that the only thing that the INWG membership would vote on by letter ballot was whether or not to submit this compromise to the International Standards bodies, rather than to try and choose among several contenders. And so there was a vote and INWG decided that this was its position, this compromise proposal, and in January of '75, Derek Barber (who was the new Chair of INWG starting January 1 of '75) announced that the result of the voting was that this proposal was approved for submission to International bodies, and sent it off to CCITT, who paid no attention to it, of course, which was not a big surprise to anybody. Because this was a datagram network kind of proposal and CCITT was really working on a virtual circuit basis. They felt that their customers wouldn't accept the idea that they had to do all this duplicate detection and reordering and retransmission inside their end equipment. They were used to the PTT delivering everything and so that's the way they would go, and they were probably right, furthermore, in 1975. That probably was the only way customers would've gone.

Weber: A lot of the conflict, as I understand it, was end-to-end versus...

McKenzie: Yes. That's right.

Weber: <inaudible>

McKenzie: Yeah, certainly so. Yes. The INWG approach was really that the network would do-- would make a great effort to deliver packets, but it wouldn't make any effort, at all, to try and keep them in order-- keep them ordered, avoid duplicates, fill in gaps. That should be up to the end systems. Because it seemed to us in INWG, even though the ARPANET was built differently, that when you connected together a whole bunch of networks built by different people operated according to different principles, you really couldn't count on much. And if it was really important to you to know that everything was there only once, but exactly once, you better do it yourself. The downside to that is that for every computer connected to the network, or every operating system connected to the network, there has to be a software package that can do that. And in 1975, there were many, many, many different operating systems, all in contention. It may seem strange in these days when the only choices are, basically, Unix or Windows, but there were dozens, or perhaps hundreds, of operating systems in 1975, and it would be a lot of work to write complex software that would run reliably for every one of them. The PTT's customers couldn't do it. The PTT's themselves couldn't do it. And the manufacturers didn't look likely to do it. So the PTT approach of saying, "It'll all be taken care of in the network," was probably the correct choice if you wanted to get customers in 1975. In any case, the CYCLADES network people immediately agreed to go with the IFIP proposal and EIN immediately agreed to go with the IFIP proposal. But everybody was shocked when ARPA said, "Well, we think we're too far along with TCP/IP to make any changes, so we're not going to do it." And that really was kind of the end of INWG. I mean, it lasted for another ten years, or maybe fifteen years, doing various things but the enthusiasm was pretty much out of it because if the ARPA sponsored networking and the European sponsored networking were going to work differently, what could we do? And the PTT's were going to work yet differently again. So there didn't seem to be too much room for INWG to do anything more and the enthusiasm, as I say, pretty much died away.

Weber: So then Cerf and Kahn made that call presumably?

McKenzie: Well, yes it was Cerf and Kahn who made the call. I have no idea what pressures, economic pressures, budgetary pressures, other congressional pressures, nationalistic pressures, may have been on them. It was a shock to us though because Vint had been one of the authors of the compromise, had seemed to agree with it enthusiastically. And according to my recollection, he said, at the time, that it was Bob's decision, that Bob just said that they couldn't lose the time or the money to back up and do anything. Since then, Vint has said “well probably I made the decision 'cause I was in charge of the programs”, but that doesn't explain why he ever went through the effort of making the compromise if he did -- if he knew that this was going to be the result. So it's something I don't understand really; not the causes, but the results were clear. The Europeans and DARPA went different ways and DARPA had a lot more money. So their way, in the end, won. Not that it was better or worse, but it was better funded. I used to think that this set back the thing that we now know as the web by a long time, five years maybe, by having things being done differently in Europe and the U.S. And still differently in ISO because then we had the open systems interconnection wars and so forth. As I look back on it from this point, it looks to me

like the web probably was inevitable. I think software was there essentially for free when a graphics user interface existed, when the hypertext was implemented a lot of places. And someone-- and so when the network in the U.S. became privately owned rather than government owned, and a whole lot of other things. So maybe what protocols people were working on in the 1980s didn't really matter that much, at least didn't matter that much in terms of slowing down or speeding up the worldwide adoption of packet technology. So it was a disappointment at the time that we weren't all marching forward together, the Europeans and the Americans, to the same beat, but didn't happen that way and eventually the better funded project took precedence.

Weber: But had ARPA agreed would there ever have been the pressure to go off and do all that in OSI?

McKenzie: Well, I don't think – that's another fair question that I've asked myself many times. John [Day] argues in his book [Patterns in Network Architecture] that OSI – the battles of OSI had to be fought because they were really not about the protocols per se, but battles for influence between the computer manufacturers and the PTT's and also between every other computer manufacturer and IBM. And that's certainly a lot of what the OSI debate was about and maybe it had to happen regardless. It would have been nice if there had been an internationally agreed and implemented and in-use protocol. Maybe it would have saved some of the money that got wasted on OSI, but it's not real clear.

Weber: You were aware at the time that this was the kind of turning point [INAUDIBLE]?

McKenzie: Yes, I think everybody in INWG was aware that – as I say, the enthusiasm went out of INWG. The organization went on for quite a lot longer working on other things, but from my point of view, and I think from the point of view of most people, the enthusiasm was gone. We no longer believed that the international community of researchers was going to do something to lead the world forward. It was now two camps.

Weber: You describe in these meetings – it must have been interesting to go to all these international meetings. You met a lot of people. Describe some of the INWG meetings.

McKenzie: Well, the INWG meetings – certainly, under Vint's chairmanship, the meetings were held in interesting places. Vint liked to travel and we went to several – probably the most unique INWG meeting took place on a – basically a ferry, an overnight ferry, between Stockholm and Turku, Finland. And we had the use of the thing that was usually the disco bar reserved for night and morning for our meetings and that's where we met, in dimly lit, plush chair, small table, disco glitter balls handing overhead (but not operating), room. The meetings were attended by two sets of people: A set of people who were actively researching, whether they were academics or builders; and a set of people who were just sitting and listening. And the debates were intense among the people who were actively participating and I don't know what the people who sat and listened got out of it. There were also ranges of opinion about what we should be talking about from the bit pushers to the theoreticians. Probably the most theoretical papers were presented by a guy named Gesualdo Le Moli from University of Milan, who was inclined to write papers about if the process alpha and the process beta wished to communicate over a channel gamma, then the channel gamma needs to have the characteristics omega and so on, and then would describe these and it was certainly not an implementation document. Then on the other end there were

experimentalists; Bob Metcalfe came to a lot of the meetings and brought descriptions of the protocols that had been developed at Xerox PARC, the protocols, and mostly for use over Ethernet. Cerf and Kahn – or Cerf was there for most meetings. Kahn, I think, only came to the one in Brighton in September of '73. But Vint was there with graduate students and other people that worked with him, and people like myself were there talking about implementation details. A lot of the focus of the meetings was along the lines of, “we can't let the PTTs push us around. The PTTs are going full bore for virtual circuit switching and it's not going to work. And how are we going to get the world to see that it's not going to work?” In addition to trying to work out what will work, which was a big job, how are we going to convince people that in the long run this other approach just is not going to be practical? Louis Pouzin was a big cheerleader for the idea of keep the network really simple. It's probably his idea to begin with. He's certainly the one who coined the term datagram. And he was there from the very beginning, in the October '72 meeting, and from there on.

Weber: People listened to him?

McKenzie: It was mostly preaching to the converted, to the true believers. He gave a lot of papers at a lot of conferences in Europe pushing the idea that the PTT's were going about it the wrong way, and I think he probably convinced some people. As with so many things, it really turns into religious arguments. If you come from this side of the street you believe this, and if you come from that side of the street you believe the opposite, and logic is not going to convert so many people once they've really decided that they're believers in one thing or another.

Certainly we ate well at the meetings. There were a lot of attendees at the meetings who knew restaurants that it was necessary to visit during our time in whatever city we happened to be in. And so those of us, like myself, who were not quite so with it just got to tag along and eat really good meals and that was great. And we got to see interesting places, although of course the meetings would go on eight or ten hours a day and then you'd go back to your hotel and write up the notes from the meeting for people to read over the next day to see if they agreed that you'd written it down properly, and usually on half a bottle of wine or more, after dinner, so not so much sightseeing, but it was fun. I enjoyed it all, even though occasionally-- more than occasionally – the meetings would just have to go over the same things that they'd gone over previous meeting and the meeting before that because the attendance had changed, there were new people or somebody raised an objection that hadn't been thought about the last time or whatever. And of course that gets pretty dull, but I enjoyed it and it was a good life for me and I enjoyed it in the end.

Weber: And how many people at a typical meeting?

McKenzie: Typical meeting probably had at least ten and maybe up to twenty people actively involved in the discussion, and another ten or so bystanders who were just listening.

Weber: And how many days you were gone?

McKenzie: Typical meeting was probably two and a half days, something like that. The one where we had the meeting on the boat between Stockholm and Turku, we started off with a meeting at the

university in Sweden, in Stockholm, and then we had that meeting and then I think we had a follow up meeting the next day at the university, so it wasn't all on the boat.

McKenzie: Later on I went to ISO meetings and those tended to go on for about a week or maybe nine days. There were a lot more people, a lot more documents, a lot more position papers that you really ought to be familiar with before talking about them, but they only got distributed on the first day of the meeting. And not-- the people were not uniform-- the people who were arguing were not as uniformly bright as they were at either the Network Working Group or the INWG meetings. Not to say that there weren't bright people, but wasn't uniform.

Weber: And more of a culture clash I imagine.

McKenzie: And often more of a culture clash too. Although they also liked to eat well and they also met in interesting cities, so wasn't all lost.

Weber: And do you remember just offhand where-- how many meetings there were and what cities for INWG?

McKenzie: How many meetings there were..?

Weber: For INWG and in which cities, or were there too many to count?

McKenzie: We probably met a couple times a year. The first couple of meeting I can remember, they were, as I said, in October '72 in Washington D.C. at the ICCC. Next one was in New York in conjunction with the National Computer Conference in either November or – sorry, June or July of '73. And then there was the Brighton meeting in late September of '73. There was a meeting in Hawaii at the University of Hawaii, maybe in 1974. Heather loved that one because I got there a little early – she didn't go, but I got there a little early and I went bodysurfing with Metcalfe and Cerf and a few other people and I tore the ligaments in my knee and I came home on crutches. And when Heather saw me at the airport – there were no security barriers in those days – she said to my wife, "Look at Pa, he's swinging on sticks." And from then on whenever I went to a meeting she wanted to know if I was going to come home swinging on sticks. She was only two years old at the time, so it was a big thing for her, yes. And fortunately I didn't come home swinging on sticks anymore.

Heather: So this is Heather, the daughter, and my favorite thing about the fact that my father was gone so much when I was growing up is that I had a collection of dolls from every place that he visited, and I also had a soap collection from every hotel that he visited, so I had a grocery bag full of little soaps from all over the world, you know, South Africa and Hawaii and everything. And I knew where every single one was from. And then I had a cabinet, a glass cabinet, full of little dolls with native dress from each country.

McKenzie: I did other traveling besides the going to INWG meeting and ISO meeting later on. Travel – in the mid to late '70s, everybody in the world wanted to know how packet switching worked. And a lot of the people who wanted to know how it worked were willing to pay to find out, and they were willing to pay – to hire BBN, because we were a consulting company, to come and talk to them about – consult or inform them about -- packet switching. And because of my good luck in being the network generalist from the

very beginning, and then staying in touch with the user community and knowing what other people were doing, I was the logical person to go on a lot of these expeditions. So I got to go give lectures in Santiago, Chile and in Bombay, India, and then Rennes, France, and Mexico City and a lot of other places. South Africa -- BBN was making a proposal to, actually through, a South African company; they were making a networking proposal to a bank and I got to go on that and see my first wild elephants and giraffes, and visit Victoria Falls. So I had a most enjoyable time doing this. Then eventually, in 1979, BBN had a contract with Olivetti to help Olivetti build a network that their sales people had already sold, but which nobody at Olivetti headquarters knew how to build, and I got to go live in Italy for a year consulting to the Olivetti people, and my whole family went with me on that trip, and we all developed a great love for the people that we met there and the places that we got to stay. So that's been another rewarding experience that came out of all of this.

You were starting to ask me a question about Metcalfe.

Weber: The story that he and Dave Boggs have told us, they knew how to do all this internetworking stuff and they couldn't say a lot because of confidentiality.

McKenzie: That might be true. I think – if Bob and Dave say that they had confidentiality restrictions, who am I to argue with them? Certainly they were doing a certain amount of internetworking in Palo Alto, so if they tell you that they'd had an experience with it, no doubt they had.

Weber: But from your end, I mean, how did you feel like they were getting really concrete input? It was more Bob, I think, that came to the meetings.

McKenzie: I don't think Dave Boggs ever came to any of the meetings and Bob-- I don't remember anything real specific, but I vaguely remember that he brought up particular points that we had to think about in doing the protocol design and perhaps that was based on direct experience or perhaps just on the fact that he's a smart guy.

Weber: And as far as you knew at the time, it was really Ethernet to Ethernet type stuff they were doing?

McKenzie: I have no idea what they were doing.

Weber: So let's go forward to..

McKenzie: Olivetti stuff?

Weber: Yeah.

McKenzie: Okay. Well, I mentioned that I got to go to Olivetti, to Italy for a year. Olivetti was in the business of making mini computers and terminals more than anything else. And they made a line of bank terminals and there was a Danish savings bank consortium that wanted to fix things up in Denmark so any savings bank customer could walk into any savings bank anywhere in Denmark and transact business as though they were at their local branch where people knew them. And the Olivetti salespeople in Denmark saw this as a wonderful opportunity to sell a lot of Olivetti banking terminals to all the savings

banks in Denmark and their salespeople said sure we can connect them all together with a network. And eventually the Danish savings bank consortium bought the Olivetti solution. Then the salespeople, I presume, called up headquarters in Ivrea, Italy, and said, "Hey we've just sold a network. Somebody there better figure out how to build it." BBN had been conducting negotiations with Olivetti for some time about possibly starting a joint venture in Europe to do networking, and it wasn't going anywhere because probably neither side was offering realistic terms to the other. I certainly don't think BBN's terms were that reasonable from Olivetti's point of view. In any case, the Olivetti people who were involved with that negotiation said, “hey, let's put this aside for a minute. We have a concrete problem now. Send us a couple of consultants who can help us design this network that our salespeople have just sold in Denmark.” So again, because I was a generalist and knew quite a lot and because the Danish PTT had built an X.25 network and I knew a fair amount about X.25, which a lot of the other people at BBN who were just involved with internetworking didn't know, I was a good person to go, and I was eager to go, so I went. And I spent a year in Olivetti's headquarters in Ivrea, which is about halfway between Turin and the Mont Blanc Tunnel right on the edge of the Alps, helping design a set of protocols that would work over an X.25 network and could be implemented in the Olivetti terminals, which were mini computer-based terminals, to accomplish these jobs. And I was there with another BBNer named Ira Richer, and he was helping them more with the network management issues and then after-- I went home after a year, Ira stayed on. Another BBNer, Tony Michel, went and spent some time there, and I think by that time Ira and Tony were helping Olivetti-- were selling BBN products to Olivetti, to OEM, and were helping Olivetti to market these network products around Europe and the Near East. So I think we were of some help to Olivetti.

As a side note on this Danish banking network, the Olivetti part may or may not have been on schedule, but the Danish savings bank had a bunch of IBM System 360 computers and had contracted with a frontend-- an IBM frontend manufacturer to implement the corresponding protocols in the frontend. And that company really never got started on doing what they were doing, although they came to a lot of progress meetings that Olivetti and the savings bank had with them, and reported that everything was going fine. So Olivetti delivered, or allegedly delivered stuff on time to the bank and there was no way to test whether it worked or not because the frontend stuff never arrived. And eventually the savings bank consortium hired me, through BBN, to go to Dallas, Texas and figure out whether the savings bank stuff ever was going to be delivered, and I went with a guy who was more expert on IBM systems than I was, and our conclusion was probably it wasn't ever going to be delivered, and if it was it was going to be real, real late. I don't know what the savings bank did after that.

Weber: And then?

McKenzie: And then? Well then when I came back from Olivetti, back to Cambridge, I worked for a little while managing a project to build a network, a banking network for Citibank in New York. We were building that with hardware that we built that didn't really work very well. Eventually they got tired of having their demos fail and stopped working on this project. It's ironic that a week before they decided to shut down their work with us, one of the guys who worked with me, a bright engineer named Lou Fernandez, who died recently, figured out what the hardware problem was and we could have fixed it if they'd given us another month. We would have had their network working perfectly. But we didn't get that

chance. Anyhow, after that I worked for BBN for a while managing a project that BBN had providing technical assistance to the National Bureau of Standards. By that time, and this was mid-80s, early '80s, mid-'80s, the ISO open system interconnection project was going full force and the National Bureau of Standards was charged by the government with representing the United States' interests in the standards-making process, and the United States' interests were predominantly defined by the Department of Defense which was really the only part of the U.S. government that had any experience with packet switching. And the Department of Defense had recently adopted TCP and IP as their standard protocols, so what, in fact, this contract meant was that BBN, in helping the NBS, needed to try to get OSI to contain TCP/IP (or close approximation) protocols as part of their protocol stack. I came into this project when it probably had been going on for about a year, and the method of operation was that the U.S..

[pause to change recording equipment[

McKenzie: Okay. We were talking about ISO, I believe. So the way that the U.S. Government's needs and views-- representative to ISO which is the American National Standards Institute, so anything that we did on behalf of NBS had to be done through ANSI, and that meant going to the ANSI meetings and convincing the people there that something needed to be done. And then volunteering to draft the position paper that they were now convinced was worthwhile and then drafting the position paper and getting ANSI to forward it to ISO; and then going to the ISO meeting and presenting the position paper and trying to get the international community to accept it. Convincing ANSI was often just as difficult as convincing ISO, so in this project, we had one person working on network layer, one person working on transport layer; probably there was somebody working on session layer, although I don't remember who or what happened there. And I worked on the presentation layer, and in fact for a year I was the rapporteur for ISO for the presentation layer, not because anybody particularly liked my ideas but because that was a U.S. spot and the U.S. had fought for getting that spot and when the last U.S. guy had to leave the spot because he changed jobs or his supporter no longer was sending him to the meetings, they had to look for another U.S. person to take it over and I was available. So I did that. In any case, the work went on for quite a while.

I think that we did some good work for NBS, not so much in this representing TCP/IP to the international community, but rather because the people, who worked on the project before I even got to it, developed a specification language based on C, and then built a compiler to compile specifications directly into working code. So you could test whether the specification was well written or not by compiling it and then trying to run two instances and see if they talked to each other. And I think that was the main contribution for that project. NBS did have-- with our considerable help-- have influence in getting an IP-like version of the network layer defined and a TCP-like version of the transport layer defined. In the end, that all came to nothing because the whole ISO effort came to nothing. And even if that hadn't been true, the way the ISO effort was structured was that there were lots of options that one could choose for any layer. So it was not at all clear that two implementations totally compliant with ISO standards would be able to exchange any information at all if they chose different options. But it was a job. After a while, we had done pretty much what we could for NBS and they had budget reductions and decided to take on the rest

of the work in the committees themselves and let our contract terminate and that was okay because we'd probably done about as much as we could.

And the next thing I did at BBN was work on academic networking. BBN had originally had the contract to build CSNET. We had done a lot of work on NSFNET although we did not build the network. We did ancillary work. Near the end of the '80s, New England was the only region of the country where there was not an NSF supported regional network because the ARPANET presence in the northeast had been so strong. But the ARPANET was about to terminate and New England needed a network. A BBN team, under my direction, approached MIT-- because a regional network in those days really did need to be run by a university-- about helping them and they said well we've already been talking to Harvard and BU about what to do. And we said well we've put together a proposal. And they said oh that's very interesting. Well why don't you give us the proposal and we'll see how it fits and after a little while they issued an RFP for a network which we could rapidly respond to, so perhaps we were just lucky or perhaps other factors were at work. And we won the contract to build the-- what was called the New England Academic and Research Network, NEARNET. Laura Breeden, who was one of the leaders on the project, had suggested at first that it should be the New England Research and Development Network, or NERDNET, but even she realized that was better as a joke than as a viable concept.

We did build NEARNET and ran it for-- on behalf of the consortium of MIT, Harvard and BU, for several years. And eventually they-- those universities were getting pressure from auditors to not be involved in so many extraneous affairs, even though this was a non-profit network. And BBN made an offer to buy the network from them and they accepted that offer. And that established BBN as an internet service provider directly and then we could market directly to commercial customers and other-- which NEARNET had been doing already, but had been more focused on academics. And so that went forward and shortly afterwards was renamed BBN Planet, and Planet then bought the southeastern regional network, SURANET and perhaps another network on the west coast, and became a national internet provider and was very successful at it until the dot com bubble burst and orders for lines exceeded the ability of BBN to pay for them, or BBN Planet to pay for them. By that time I was gone from the company, so that didn't have any effect on me. I retired from BBN in 1996. Actually I was fired from BBN in 1996 because the new company management was trying to get rid of all the expensive managers, whose absence wouldn't be noticed for a few years, in order to make the bottom line look better so the company could be sold. And the company was sold to GTE and then it broke into pieces and Planet eventually went bankrupt when the dot com bubble burst and the research and development part of BBN continued along, relabeled as GTE labs. Then it became privately held by some of the principles and eventually, within the last year, was sold to Raytheon, another New England Company, who interestingly, as far as I understand it, are managing their BBN division from a headquarters in Texas rather than from right up the road in Massachusetts, which perhaps ensures that they'll go along the same path that they've always gone along, of relative independence for somewhat longer. But in any case, I was gone from all that in 1996, and I don't miss it much. Certainly I miss some of the travel, but I've managed to do a fair amount of travel on my own since then, so that's not a hardship either.

Weber: With the OSI era, what do you-- if OSI had prevailed, how would something like the web look different today?

McKenzie: Well, of course you're asking me to speculate and engineers are not very good at speculation. Certainly none of us could have speculated back in the '70 what the current network would look like. A big problem, as I said, with OSI was that every layer had a number of different optional ways to implement it. And you could be OSI compliant without being able to talk to some other computer that was also OSI compliant. That would have got worked out eventually, surely, or it wouldn't have been of any use at all. But OSI-- the OSI effort was directed at big manufacturers and big providers. It's difficult for me to see how even if OSI were adopted, how small, innovative companies would have been able to do much of anything in the short term. Of course in the long term, if OSI was successful, it would provide an infrastructure for transferring data around the world, and it would not have tied the hands very much of any information provider. And so the web might look today, or five years from now, very much as it does now, or will five years from now.

I think that the reason that TCP beat OSI is because it was already there and OSI wasn't. And when asked the question well what if OSI had won, you have to imagine a time when there was an OSI implementation about as big as the TCP/IP implementation, so there would be something to win. And it's very difficult for me to envision, given the speed that OSI was going and the corporate bureaucracy approach to defining it. And implementing it was hardly even in anybody's consideration at the time it was being defined. It's difficult for me to envision a time when there could have been an OSI implementation that was as big as a TCP implementation. A lot of governments were pushing for manufacturers to do it, and most of those manufacturers are gone now. I think a thing that made networking in general possible was having only two operating systems around, two real operating systems: UNIX and Windows. And once the world was reduced to a very small number of operating systems, I think what the vendors or the manufacturers of those operating systems chose to do was what was going to be the real standard. And it's not at all clear to me that either the people working on UNIX or the people working on Windows would have gone for OSI. I know that Bill Gates was driving Microsoft toward the Microsoft network and he, earlier than OSI, realized that the internet TCP/IP technology was winning. And in what I consider a brilliant managerial move on his part, turned Microsoft around and said no we're not going to do our own; we're going to do TCP. Would he have done the same thing with OSI? I think he would have fought them. I think he would have said Microsoft does – is doing it better;we're not going to do this. So it's really hard to imagine-- for me to imagine a world in which the OSI protocols triumphed, even though I spent a couple of years working very hard to try and make them good. I don't think they were really very good, and I think Microsoft and people who supported UNIX would have decided what was going to happen if it hadn't been for TCP/IP.

Weber: Any other areas you want to cover?

McKenzie: No, but it's been a lot of fun talking about the olden days. I always enjoy it. Thanks for the opportunity.

Weber: Thank you. Appreciate it.

END OF THE INTERVIEW

