Interview of Kanwal Rekhi

Interviewed by:
James L. Pelkey

Recorded August 17, 1994
Palo Alto, CA

CHM Reference number: X5671.2010

James Pelkey: I'm with Kanwal Rekhi at Novell. It's August 17 and Kanwal, as I indicated, I'll have this transcribed and I'll get this to you, and anything that I use I'll pass by you as well. I thank you for your time and I'm interested, as I mentioned, in how innovation becomes economic growth.

Kanwal Rekhi: All right. My involvement with this whole process starts in 1981. Before 1981, you know, I worked for Singer-Link for about ten years, a defense contractor, and I was in my mid 30s, and I was worried about being a lifer in the defense industry. I had see people ahead of me, and I'd say: "I don't want to be like that." The world was changing, and I was in an industry which was really static. So I left the job at Singer-Link --

Pelkey: And what company was that again?

Rekhi: Singer-Link. S I N G E R, Singer Sewing Machine. Link was the flight simulation people, they built all the flight simulators for the air force, navy and the airlines, and I was a system engineer in that. This is 1980 -- December when I said, you know: "I need to be out of here."

Pelkey: December of 1980.

Rekhi: December of 1980. "I need to be out of here. I'm 35. If I don't leave now, I'm a lifer here," and I didn't want to be a lifer there. Microprocessors were there, the Apple PC was a big deal, this is in the pre-IBM PC days now, remember.

Pelkey: Although in December of 1980, the PC wasn't out yet.

Rekhi: Oh, yeah, this was the pre-IBM PC days, but there was a big excitement about networking the workstations, distributed computing, the Z80, the Multibus, and Ethernet was being introduced but was not ready to introduce yet, and I got hired by Zilog because there was a wave of people who had left to form Bridge Communications; you know, Judy Estrin, Bill Carrico, and the whole people, so they were doing --

Pelkey: And was this the Benning era?

Rekhi: Yeah, Al Benning was still there. He was still there, and I joined when he was still there. He was the software manager and I was hired as the hardware manager in the same group, which was doing a personal computer along the Z80, and networking along ZNet, and a file server -- it was a machine called MCZ Micro Computer from Zilog, as a personal computer, and there was MSZ, Micro Server from Zilog. They were both Z80 based machines. The stuff didn't work at all. The computers didn't work, the network didn't work, the server didn't work.

Pelkey: How did you sell them?

Rekhi: Hard to sell them. When I first got there, I discovered that we had very poorly documented, very poorly -- you know, nobody knew what the machine was supposed to look like; nobody knew what the network was supposed to look like; nobody knew -- and I spent about three months just trying to understand, and then I said: "Oh, every engineer has a different computer on his desk." So I went through a learning process under which we really stabilized the machine and the network, and we were done, oh, six, eight months later, around the middle of 1981. So we had a very well functioning machine with a very stable file server. The IBM PC got introduced, and Ethernet got introduced all at the same time. And Zilog was falling apart as a company. Our CEO, Manny Fernandez, left to start his own company, and there was a person under him, Bernard -- not Bernard, Rolando Eschevera, he left to start his own company, and the person under him, what's his name, Doug Schwartz, he left to start his own company. Doug Schwartz started Integrated Solutions, which was a Unix 68K BME Bus Machine, and here I was with a couple of people doing all these machines with all the problems of management above us missing. So there was nobody there above us.

Pelkey: They were all off starting companies.

Rekhi: And we were doing a wonderful job. Finally we had everything they wanted to work working in the lab, and there was nobody --

Pelkey: To manage it.

Rekhi: Nobody to sell it. So in December of '81, we said: "There isn't going to be a revolution out there." We didn't take the PC very seriously at that time. We were much more networking people, so one day the 68 processor, the Multibus, star bus, the Ethernet already announced, and somebody is going to have to offer the distributed machines using star components, so the vision --

Pelkey: Similar to Sun.

Rekhi: Sun Microsystems. Sun Microsystems, by the way, was there and a start-up too, the same time frame.

Pelkey: Yes, now they were getting their boards, at this point, from 3Com?

Rekhi: No, they were -- Forward Technology, there was somebody called Forward Technology. Forward Technology had a 68K-based processor board which was very popular. Sun was using that, then, of course, Sun did their box, and Sun did their own boards also, eventually, but Forward Technology boards were the ones that were most popular. 3Com did the Ethernet boards for Multibus, for BME bus -- not BME bus, Unibus, and Q-bus, but they were dumb boards. They had no processor, no intelligence; they were simple. This was the days before the chips. I was a hardware wizard, and so I told people: "I can do an intelligent board using the MSI chips, and have a processor and memory on the board so where you could really board drive, process on the board itself, and not have to worry about doing it on that stuff over and over. So we could do a board, an abstract board, which a processor, memory, and Ethernet, and we can abstract out the bus. This can be put onto any and all buses. So we came up with a design for that purpose, and we did the board for all the buses -- Unibus, Q-bus, VME Bus, Multibus, even went to the IBM PC bus and Macintosh. Our idea was a bunch of protocols on board, which could go into any of those environments without having to redo the protocols, and so we were the first ones to do the intelligent Ethernet boards. Once we started doing the boards, they said: "You need protocols. What should we put on the board?" And if you looked at the --

Pelkey: Let me interrupt you for a second. You mentioned Sun. Did Sun become an OEM?

Rekhi: No. We tried very hard to sell Sun, but Sun became an OEM of 3Com for a while.

Pelkey: Yes.

Rekhi: 3Com was ahead of us there. 3Com was there six months or a year ahead on the boards from us. So they became a supplier, but 3Com basically abandoned the marketplace as soon as the PC came out.

Pelkey: When they came out with the EtherSeries.

Rekhi: Yeah, PC came in the end of -- the end of '82. PC came in the end of '81, but 3Com went after the PC market at the end of '82, and that's when we were just beginning to sell the boards. So we had just arrived on the scene and 3Com was leaving the scene.

Pelkey: So Intertel is there.

Rekhi: Excelan.

Pelkey: Excuse me, Excelan, sorry.

Rekhi: Excelan was there. Excelan had a higher, faster, smaller board than 3Com, but they still did not have a processor and memory on the board. They had done -- 3Com had emphasized the cheapness, a real barebones board, and Excelan had done a slightly higher performance board, but still no processor and memory, and we were the first ones to come out with a processor, a full front-end processor.

Pelkey: So protocols.

Rekhi: Yeah, protocols. Now, what protocols do you use? If you look at the time frame, the XNS protocols were very popular.

Pelkey: Everybody else was using them.

Rekhi: They were supposed to be a matched pair with the Ethernet, and OSI was the dream everybody had. When I started to look at the problem, I said: "You know, OSI is not here yet. It's still being defined, and it might be two, three, four years before make them available." I look at the XNS and XNS didn't have any applications defined. They had the transport and the lower defined. Everybody was doing his own applications; Bridge, Ungermann-Bass, 3Com, even Novell used the XNS, Sytek, they were all doing their own applications and they didn't inter-operate. And my sense was that, if we were to pick up the TCP/IP, we could provide the instant solution in the environment, and the queuing was -- Xerox spent billions of dollars to perfect Ethernet, and it worked very, very well, and the government spent billions of dollars on the ARPANET to perfect TCP/IP. If you put them together, you have a very reliable, instant solution. There was a fear that TCP/IP was very inefficient protocol. It was designed for slow, long-haul networks, and wouldn't work very well on LANs because it has all that excess baggage on the error recovery, the unreliable underlying media, but my theory was that, if the media is reliable, those parts of the TPC/IP won't be exercised, and won't matter. And it will be just as high performance, so we adopted TCP/IP very quickly, and to our surprise, it was being also part -- the Berkeley 4.1C came and had TCP/IP built in.

Pelkey: When did you release your product?

Rekhi: We released our product sometime in '83.

Pelkey: In '83, so Berkeley Unix 4.1C, when did that come out?

Rekhi: That came in the middle of '83. Maybe the end of '82.

Pelkey: Ok. What did you say about '82?

Rekhi: See, '82 / '83 time frame Berkeley 4.1C came, Unix with TCP/IP built in, but around the same time, we were introducing our TCP/IP as a stand alone.

Pelkey: You introduced your product at what time?

Rekhi: Middle of '83.

Pelkey: So the Berkeley Unix was out at this point with TCP/IP.

Rekhi: Yeah, I guess maybe we became aware of that by the time we came out.

Pelkey: So it might have been the fall of '82.

Rekhi: That didn't drive us. We were driven by --

Pelkey: No, no, you didn't even know about that happening.

Rekhi: Yeah, but when we -- let's say Berkeley 4.1C was beginning to become available in larger

numbers to the users in '83, and they were playing with it and they were using it. There were many version of Unix out there at the time. If you remember, System III, Version VII, even the --

Pelkey: Xenix.

Rekhi: Xenix, Venix, 4.1, Berkeley 4.1 before the C days, and it was -- there were at least 40 companies doing 68K -- you know, one of the buses, standard buses, and varied versions of Unix to commercialize them. And as soon as we came up with our boards and software, we saw the opportunity of being suppliers of networking to all these people who didn't have the TCP/IP. So we focused on individual TCP/IP with an Ethernet board, a smart board, and the applications on any of these buses, we had the matrix. If any of these versions of Unix you had, if you are using any of these buses, we are the place to turn to get TCP/IP, and we signed up just about everybody, everybody except Apollo and Sun. We signed up NCR, Unisys --

Pelkey: Computervision?

Rekhi: Oh, yeah, Computervision. We had about 35, 40 OEMs.

Pelkey: Did you just get the Berkeley version of TCP/IP and fine tune it?

Rekhi: We did a lot of it ourselves, but it was the availability of the Berkeley applications. We did the Berkeley applications, so we are not out from the Berkeley code, because we weren't aware of it. The board, on-board --

Pelkey: Who wrote your TCP/IP then?

Rekhi: We had two people inside and two people outside. There's a guy by the name of Bill Northlich and Bruce Borden who were hardware. Bill Northlich is still with Novell. They were contractors mainly, reported to us. Bruce Borden was one of the founders of 3Com and Bill Northlich was one of the contractors who worked with Bruce, so we hired them in '82 and '83 to do this work for us. And by the way, we were still doing some development work on the XNS and some development work on the OSI.

Pelkey: But for productizing, it was TCP/IP?

Rekhi: Well, we were doing it all together. TCP/IP was where we were getting some revenue. By '84 we were getting in big trouble because we were doing too much stuff with too few resources, and without the 35 OEMs, revenue wasn't that much. In the -- by the end of '84, there was real trouble because we were spending much faster than we were bringing the revenue in, and as it happens in any start-up, there is a day of reckoning, giving too much for the resources you have and, as a result, you're not satisfying anybody. So the company went through a reorg, and I was one of the founders. In the old business I was the head of development, I was the VP of Engineering, and one of my partners was the CEO, and he was fired by --

Pelkey: This is Jain?

Rekhi: This in Inder Singh.

Pelkey: Oh, yes, Inder Singh.

Rekhi: Inder Singh, and Dr. Inder Singh was fired by the board because we weren't making our business plans. So, once he was fired, I was made interim president, and this is where it became very interesting, because I was very worried that we weren't doing the right things and enough of opportune things, so right after I became president, interim president, we're doing nothing but TCP/IP and Internet, and we are going to go after the end users rather than the OEMs, and by -- this is now early '85 -- by that time there was a well developed CAD market, workstation, CAE, CAD, the Daisy, a whole bunch of CAD companies, and the PC had become fairly popular. The PC AT had come already and 286 machines were becoming

popular, so I defined the business very quickly by saying -- and, the VMS machines were very popular; DEC, by that time, had that, DEC has it now campaign, and VMS machines were very, very popular -- so we defined the business in sharp terms. "We will do the TCP/IP connecting from the PCs, the Unix machines of CAD types, and the VMS machines, and we'll offer you the basic two applications, the file transfer and virtual terminal, and that's our business and nothing else," and started with packaged boards and sold to the end users only. And we used to have, basically, a campaign that "DEC has it now," and Novell said, we'd say: "Well, we have it too, and we'll let you connect your PCs to your Unix machines, to your VMS, and do instant networking," and that turned out to be a winning strategy for us; very, very popular very quickly, and in '85, we took off like a rocket. The other guys, none of the others really adopted TCP/IP as fully as we had. Excelan adopted it in '86, fully. 3Com and Bridge adopted it '86 / '87 time frame, and --

Pelkey: So you really carved that portion of the market out for yourself.

Rekhi: Absolutely. By the end of '85, we could hook up any of the Unix machines on any bus to any of the VAXes, MicroVaxes or any size VAX machines, to PCs of all types, and we also had done it for al SX machines (??), the DEC's old minicomputers, and also we were working with Intel to do all of Intel's RMS machines (??). Intel had this big system business, so we did it for RMS machines also. And we also had, besides that one, what we called the standard OEM package, where we give you this OEM module of hardware, and you could build a bus around very quickly, and the source code, and we sold that to the people for the buses that we didn't offer, and we had about half a dozen of those OEMs also. So coming into '86, we had a very broad TCP/IP solution available, and both OEM and end-user.

Pelkey: So you ported the TCP/IP that you had in house --

Rekhi: To a PC board.

Pelkey: To a PC board. Did Dave Clark's efforts at MIT --

Rekhi: See, they had done what I call a lower TCP/IP in the host. You know, the PCs were very underpowered then. They didn't have enough memory; they didn't have enough processing power. We had the front-end board where TCP/IP ran on the board itself.

Pelkey: Gotcha.

Rekhi: So the advantage we had was the same code ran on all those boards. The boards were a way to isolate the protocols from the operating system and the bus environment, so we offered a very, very high-powered solution.

Pelkey: So the TCP/IP that ran -- the MIT PC version ran in the processor?

Rekhi: Processor.

Pelkey: Of the PC processor.

Rekhi: And used the PC memory.

Pelkey: And used the PC memory, or all 64K of it.

Rekhi: And so it was not very fast -- until the FTP software was developed and the 386 became available toward the end of '86. We were the first people to offer you TCP/IP on the PC which could support the multiple applications simultaneously. We had partnered with the --this is pre-Windows days, remember. The Windows was a pretty poor product at the time. This company, Quarterdeck --

Pelkey: Oh, yes, Quarterdeck.

Rekhi: Yeah, Quarterdeck had this way to switch windows, and we partnered with that and we offered TCP/IP stack, and multiple windows running different applications could use the same stack simultaneously, and the other state information was maintained on other boards, so when you switched from one application to the other application and back, your connections were not lost.

Pelkey: Was that successful?

Rekhi: Very successful product, very successful product, because that made the PC work very well with that larger environment. And people could use, in one Window, maintain their connections to the Unix machines, and another window to do their local spreadsheets, and the card became very, very popular. And, as a matter of fact, in the '86 time frame, we worked with -- we were trying to sell them with Windows. Windows would not maintain the state. When Windows would switch from one window to the other window, all the connections would be lost. In '86, we spent a lot of time with Microsoft, helping them evolve Windows so that applications, when they were --

Pelkey: So the multiple windows would stay open and usable, as opposed to closing down.

Rekhi: Useable, yeah. We expended an enormous amount of time with Microsoft, so there was a close partnership with us and Microsoft in the '86 / '87 time frame, and Windows -- it had become useful by, on that count, Windows 2.1 release. I don't know the exact name, date, which made that happen was done with our help. Coming at the end of '86, you know, we had become a very successful company. We were doing $20 million plus revenue, making $3 million of profits. I think we had done $22 million revenue and 3.6 million of profit, and then we went public in February of '87, and I made that famous speech at the -- the First Interop in Monterey, where I had bought a Porsche and my speech developed that TCP's paying for Porsches and BMWs. And we were the first people to make money on the TCP/IP, the first commercial company to make money on TCP/IP. There was the --

Pelkey: Now, the Monterey Conference -- the first Monterey Conference was in '86, wasn't it?

Rekhi: '86. The first one was just the seminar conference. He had a seminar setting. Probably March of '87 he had the one where he had allowed the commercial participation. Dan Lynch -- by the way Dan Lynch did work as a consultant for Excelan in '86 before he started Interop.

Pelkey: Oh, OK.

Rekhi: I had -- he came and worked for me for three months. He was very restless as a person, and he wanted to do something, and he told me: "I think this TCP/IP thing is going to be a big deal, and it's going to require an industry structure to support it," and he talked about doing a newsletter, doing all sorts of seminars --

Pelkey: Connections was --

Rekhi: Connections was -- and I offered to help fund him. He said if I need -- not big time, like a few thousand dollars or something. He said he will take me up on that if he needs it, but he did the first seminar, and it was a big hit. It not only paid for itself, he made money, so he never needed any funding. So Dan Lynch's Interop was launched in the fall of '86. First at his seminar he was expecting 200 people and I think 800 people showed up, and then, of course, he really capitalized on that one and did the next one, and he charged more money --

Pelkey: And had a lot more people come.

Rekhi: A lot more people come, and the rest is history. So, looking back, you know, we found the method to develop TCP/IP given they had modest activity, and my theory was that the government has spent a lot of money to perfect TCP/IP and produce the environment on the Arpanet, and we could use it. In the early days, when we were pushing it, we were the only people pushing it, but by '86 / '87, the other guys had become players. Microsoft and Excelan had to produce, and they would supply them with

TCP/IP, but nobody else, really.

Pelkey: You supplied Excelan?

Rekhi: No, Microsoft, but those deals got nowhere, because every time -- after deals were signed, they were giving us less and less and making us do more and more, and so partnership with them was really very one sided. I went to talk to Leonardo in '87 to see what he was going to do with TCP/IP and said, well, he doesn't need to worry about TCP/IP. He's got his hands full with the PC business he's in. And I kept telling him, you know, there's market there of wider connectivity, and he said: "Pal, that's your opportunity. My opportunity is, you know, the PC networking, PC file server." But, by the end of '88, they had become aware of what we were doing. We did extremely well in '87, we did $39 million. In '86 we had done 22, in '88 we did $66 million, so we are profitable every step of the way, ten to 15% profitability. In December of '88, they did call. And he said: "Well, you know, you were here last year talking to me about joining forces. How about now?" By that time I really -- Mike sounded very sour, very sour, because they kept saying: "We'll bundle your stuff free. You'll make your money in the after-market by updating and upgrading everybody." And I said: "Why would I give you my stuff I'm making money on right now free," and so the relationship was very, very sour with them. So when they and I met in December of '88 and he proposed that we join forces, and his basic proposal was if we joined forces we had the PC networking through Netware and TCP/IP would give us everything else we need. We had a large number of building blocks we'd need to provide the whole network, and that turned out to be a very, very sour partnership.

Pelkey: Now, that happened after NET.

Rekhi: NET was done in spring of '88, and it was being done by the board sort of over my objections.

Pelkey: Oh, OK.

Rekhi: There were some troubles developing. We had brought a CEO from outside.

Pelkey: Richard Moore.

Rekhi: Dick Moore. And he got involved with an employee -- a sex scandal. He was a married man, and that turned out to be very bad, and the board felt that it was easier to merge the company than fire Dick Moore and go through a process. But that thing got out -- this is not public information --

Pelkey: I understand.

Rekhi: So what happened was the merger was being done to solve that problem, but while we were doing the merger, Dick Moore -- Moore got pissed off and they fired him anyway after the merger was announced -- so he was fired, or he left, anyhow. Probably he left. And, so I became the CEO again, the second time. The first thing I told the board, if they do that merger, I'm done the day after the merger, because I don't think this merger made sense. I don't like the people. I don't like Bruce. I don't like the style, and I didn't tell them why, but I said every time I looked at them, their way of doing business, it just doesn't smell right to me. And one of the board members, Victor Mason, came up to me, he says: "If you don't want the merger we won't do it." So we didn't do the merger. And I told him, you know, this merger was being done for the wrong reasons. It doesn't feel right, and I'll bet you any money that either Microsoft or Novell will applaud us, if we keep on doing what we're doing, because they won't need what we are doing, and that turned out to be true. And, I don't know if you know about NET. They were not playing straight, right, and board members were very happy in the end.

Pelkey: Oh, I can imagine.

Rekhi: And they were very happy, because I was able to smell it, but I wasn't able to --

Pelkey: Articulate what it exactly was that didn't feel right.

Rekhi: They played so close to the edge financially. My style was that you make money by building products and selling them well, and financial engineering is not something that should be part of it, and they were Harvard MBAs, all of them, and they kept saying: "This is where the real money is made, you're playing, not -- you're playing neater. Be sharp on the financial side," and I kept saying: "If you think so. You never know -- "

Pelkey: Right.

Rekhi: "-- When they're going to be all over you." And to me, you build the technology, parts, you hit the sweet spot, you have the customers come, and that's how you make money. And these guys were, you know, capitalizing all the software development. Their numbers were better numbers than ours, because we did zero capitalization of software development, and they did 50% development. So I would expense off my R&D and they would capitalize it, and their numbers would look better, and I told our board: "But those are funny numbers. Real numbers, the way I have --" and, at the end, when they did come and talk to us, then we had this nice growth of Novell. I don't know if you know it. If you look at Novell's growth like that, and they acquired us in summer of '89, Novell's growth has tripled, and within the next two years, the stock was up eight-fold, not a bad deal for our company.

Pelkey: Not if you were holding stock.

Rekhi: Oh, yeah, I had stock through the whole process. I believe in the process, and it was very good merger.

Pelkey: Now, you sold your board business off.

Rekhi: Right after the merger, we said that Novell -- we decided we were being less focused on software, and that's where the future is. And the basic theory Noorda and I used to have, which I strongly believe, is that the hardware business is not something that can build value in, long-term value. The whole business is driven by the underlying silicon cycle. The underlying silicon cycle moves, all the other guys have to run to stay in place. They have to do all sorts of reengineering to stay in place, but the software guys, their value doubles. Now the software runs on twice as powerful processors and supports twice as many users, so the investment you make in software, they get more returns and better and better as the silicon cycle moves. The hardware guys are always running to stay in place. So his special theory was: "Sun is a hardware business, and we focuses on software," which we did. And by the way --

Pelkey: You sold, what, the standard microsystems?

Rekhi: No, we didn't sell directly, we licensed out the technology to have it in place, and so we never sold the business to anybody. We licensed it out. Where we licensed the idea and the hardware designs, (unintelligible) Forward Technology, FTC, yeah, Forward Technology.

Pelkey: So you have maybe ten to twelve licenses.

Rekhi: Yeah.

Pelkey: Now this is in '89?

Rekhi: '89, '89, yeah. So they will do the hardware and mail customers, they will sell this, so they will pay us 15% royalty on the boards, and we'll collect half of the royalty on the software, so they will do all the work, and so, as a matter or fact, Novell collected over $45 million worth of royalties on that stuff in '90 and maybe that same number in '91.

Pelkey: Wow. What a smart deal.

Rekhi: And there was zero effort on our part. So all the hard to sell business started to turn into royalty

business very quickly, and at this late date, we are still collecting some royalties, because we were designing for some of the Air Force Tomcats and Navy Tomcats and they have a long life. But, by and large, we are -- we focused our effort to do the TCP/IP on the Windows, on the desktop, and on the server, and all sorts of other things you need to do to make not that environment that environment. So, now in that environment TCP/IP is a big deal now --

Pelkey: Well, yeah, but it wasn't. You were right at that point in time -- in '83, Xerox also stops giving out -- they say they're not going to give out the higher level protocols.

Rekhi: Oh, see, but that was a big plus for us, because if they had given that (unintelligible) no way to use protocols, so everybody else's applications wouldn't -- If you bought from Bridge, you were stuck with Bridge. If you bought from UB, you were stuck with UB. And if you bought from us, that was your other alliance -- if you bought from us, Sun interoperates with us, Berkeley Unix interoperates with us. We offer you a solution right now. We had VMS, and by that time, the big guys had come with the TCP/IP solutions, and so they -- what had happened was we were so progressive, by the time other people came, we had close to 100,000 boards shipped by the end of '85, that you had to interoperate with us, and one of the things we had done was we had adopted 4.1C as the worked started. That was the one we had to interoperate with, and 4.1C moved to 4.2. People still had to work with us because we were the dominant --

Pelkey: Do you recall when 4.2 came out?

Rekhi: 4.2 came out in '85.

Pelkey: It was really the 4.1C version that came out in late '82, probably.

Rekhi: Late '82.

Pelkey: But in '83, people were starting to get the licenses on Unix at Berkeley and playing with it and trying it and so on.

Rekhi: By the way, we also supplied a lot of software to Berkeley itself, because they were (unintelligible) and they had PCs and they had other machines, and we offered the TCP/IP connectivity to them. And I went to -- we had become the supplier of TCP/IP to universities on non-4.1C / 4.2 machines. If you had System III, you had Version VII, any of the other Unix, Santa Clara Unix, we had TCP/IP on that one. I went to Japan in February or March of '86, and we had a Japanese distributor partner we had signed up in 1985. He took me to Tokyo University, and they were all wired up with TCP/IP, and it was all us.

Pelkey: Amazing. There was some competition on the software side --Wollongong and NRC was in it --

Rekhi: Wollongong, NRC -- Wollongong and CMC.

Pelkey: Yeah, CMC.

Rekhi: Wollongong was not a good player. They were under-funded, and they were working on the hardware. I don't know if you remember, they had this product line called Pegasus, which they were -- they had -- they also had this product line called Unis, the Unix on the VMS, that's how they started, right?

Pelkey: Yes.

Rekhi: And, from there, before they did TCP/IP -- Wollongong, by the way, was a TCP/IP licensee and OEM in early '85.

Pelkey: Wollongong did.

Rekhi: Yeah, they became an OEM and licensee of ours, but they were doing this product line called Pegasus, where they were going to have this Unix bus, but they had built around the OEM, bought the OEM from Intel which will let you offload the software development from VMS machines, so you become the front end of the big Unix machines and Unis machines, and TCP/IP will connect that to the machines, and you get the service here, locally. So they had spent an enormous amount of time working on that one, and funding was not enough for them to see it through. And the economics was not very good, either, because he was going to save some money, but as these machines become cheaper, Microvax, they lost all this job, so they were not -- in fact, in '85. They -- we never felt them as a -- CMC was the main competition for us in '85 / '86 year. We had a board set which was based upon an Intel processor --

Pelkey: And they make the LANs.

Rekhi: They had the LANS, which were 68K, and we had the Intel chip set, on Ethernet, they had the AMD chip set. And I don't know what happened to CMC. We were very afraid of them. We were very scared of them, because Intel was very hard to develop on.

Pelkey: Yes, yes, and, it took more board space, as I recall, to use the Intel chip than the Lance chip.

Rekhi: Yeah, segmentation of the chip was a problem. Intel's had been – they developed the logic on 286, and 186 was a big problem, because, I mean, for us, it was a big issue. So CMC had a 68K board with the clean memory setup, but I think what happened to CMC, they over-hyped their solution. Their solution -- they were making claims that were not sustainable. I had saved the ad sayings they had done. They had run the ads saying that their chip gives them 2 million bits of performance throughput, and they ran the ads in May saying that theirs had 1.4, and they ran the ad in July saying theirs had 1.1. Same ad, but the numbers had changed, and I had them framed and I sent them to all of our salespeople. "Show it to customers." So they headed south.

Pelkey: Good for you.

Rekhi: And this is from the ads they had published themselves, because they had been challenged and they reduced them, so they had some credibility problems in '85, '86, and we, on the other hand, were very understated people. We would sell, and --

Pelkey: You were much more -- I mean, they did TCP/IP as well, but you had a much more sophisticated software strategy.

Rekhi: Oh, yeah. We were absolutely software driven, and they were much more hardware driven.

Pelkey: Yeah.

Rekhi: And our basic principle was that we had the exactly same design. Our hardware design was abstracted to the point where exactly the same software loaded everywhere. We didn't have to change a bit. We bet the firm on TCP/IP. They had TCP/IP, but they also were doing, if I remember right, XNS.

Pelkey: Yes.

Rekhi: And they were starting to -- we never bought XNS at all, even though there was the IEEE effort, OSI effort. I had trouble, absolute trouble doing OSI. In '85, I had made a couple of decisions. Ungermann-Bass and GE formed what would be called INI, to go after the factory automation and OSI and the Token Bus environment, and MAP

Pelkey: Broadband.

Rekhi: Broadband, right, and I had to say whether I would do that or focus on what I am doing. There was fear -- not shared by me, but industry-wide, that TCP/IP is a short term solution. There's going to be OSI everywhere, for the factory and the office, MAP TOP and, you know, the factory will have the 802.4

and then, afterwards IBM came with the Token Ring in '85, and so I had to decide to stay sharply focused on what I was doing, or broaden my focus, and I chose to stay very sharply focused on TCP/IP and Ethernet, and I said: "I can do them now, in this market. If I succeed here, I'll have enough to do." And there was a problem with my board. They kept saying "Token Ring is in your future," and I kept saying, you know: "Who knows what the future is going to be?" And we had a board member who was IBM's chief scientist for twenty years. What was his name? Gene Amdahl and this guy were the architects of the 360. Bob Evans. He was the -- he had retired from IBM and he had joined Hambrecht & Quist, and he was on our board. He was not a board member, but was a board visitor from Hambrecht & Quist. And he kept saying, IBM's going to steamroll us if I don't do the Token Ring, and if I don't do the OSI. And I kept saying: "I'm solving the problems right now. Customers are paying me the money. And I don't need the others since what I'm doing is what they want. If that means go over there, it will come at the expense of what I'm doing here, and I don't have the resources. And I kept saying -- this was an insight I had which I think sustained me. I kept saying this problem is going to be solved in the field and not in the committees. I kept telling the board that TCP/IP will get better and better the more real world problems it faces; it solves and satisfies customers. The other ones are going to stay in the lab and I cannot image they can get refined and -- and that turned out to be a very smart insight. I wrote a paper, which basically said that -- in '87, I think it was published in Novell's -- there was a guy who used to publish a newsletter from Minneapolis --

Pelkey: Yes, yes. Thurber?

Rekhi: Yeah, Thurber, Ken Thurber. He published my paper in '87, which basically said 'It's hard to imagine that TCP/IP can be stopped by anybody, given it's getting refined in the real user, real world, and evolving under real conditions as against the lab setup.' So he published that paper in one of his newsletters.

Pelkey: Well, that's a wonderful story. Great timing.

Rekhi: Timing -- luck and timing is everything.

Pelkey: Yeah, but also a great deal of common sense on your part.

Rekhi: Oh, yeah. Two or three. One was my original thesis that TCP/IP and Ethernet, mating them together, would be a sensation. We had applications. That turned out to be a very, very good insight. And the other one was, it was the end, and focused sharply on one or two things solid.

Pelkey: And not try to do everything. Plus, a lot of the other players got -- Novell pulled 3Com towards becoming a minicomputer type company -- a client/server company.

Rekhi: Oh, yeah. 3Com and Microsoft. 3Com really got (unintelligible) by Microsoft.

Pelkey: When I talked to Bill Krause the other day, he was saying that, when Novell came out with -- before 3Com came out with 3Plus -- Novell came out at the time, 3Com's software was always looking at volumes, and Novell's software came out and you accessed files. And he said: "We didn't appreciate it at that point in time, that different paradigm, how important it was.

Rekhi: Also, 3Com kept seeing themselves as a hardware company, a lot more and for a lot longer. Microsoft and Novell decided to focus on software only. The file server, file access was one of the insights Novell had, but the other insight Novell had was the value of the software-only solution, because they used try, with him, at that time, a handful of those digital watches. I don't know if you know the early digital watches. Remember all these chips companies became digital watch companies -- Fairchild, AMD, National, they had digital watches. All of them, right? And Hugh (unintelligible), he said: "There's no value here. The more you do, the less you get." And his best insight was there's no way you can sustain value in hardware designs because the silicon cycle is going to obsolete you.

Tape Side Ends

Rekhi: That's what he said.

Pelkey: Now who said this?

Rekhi: Leonardo? You see, at the end of the day, you're going to use Intel chips, and you're going have to run Microsoft and Novell Software. There's very little room to innovate anymore, because if you do add some features, and no software users, you have extra cost with no benefit. And Digital (unintelligible) was an example we used to use. So, the basic theory he had put forth was either you have to be in the software business, or certain business if you want to have franchises, the one in the middle is the tough business to be in, and that turn out to be absolutely true, and 3Com was stuck in that hardware, (unintelligible) hardware. They believed enough to dump his hardware business.

Pelkey: That's right.

Rekhi: It was 50% of his revenues and he dumped it.

Pelkey: StarLAN. Did you compete?

Rekhi: Yeah, we, I ignored StarLAN. Even the issue of -- I don't know if you heard a word called "YAFO."

Pelkey: No.

Rekhi: Yet another fucking opportunity.

Pelkey: [Great glorious laughter].

Rekhi: YAFO. Yet Another Fucking Opportunity. Because what happens is you steel your opportunities. My basic theory was that a top grade solution will win. StarLAN to me appeared underpowered, so --

Pelkey: Do you remember if that was '86?

Rekhi: StarLAN was a big thing in '85, '86. We did fool around with StarLAN in the CAD, to see if you could do StarLAN without Geneva Software, and we were able to do it. We did do a StarLAN board for the PC, because, what happened, our designs were accepted enough so that you could change -- there was an Ethernet compatible chip available from Intel which was StarLAN, but never sold any. We put it out, and maybe we sold 200 of them, but we could not do it on the Multibus or VME bus or Unibus. We had bought $100,000 worth of (unintelligible), so we left the hardware very quickly, and we had done that one, but refocused and stayed on that TCP/IP, Ethernet, because then we had, at the time of the merger with Novell, 120 software people.

Pelkey: Wow.

Rekhi: Doing TCP/IP, and maybe about 15 people doing hardware. So even thought we had the reputation of being a hardware company, we were really very sharply focused on software.

Pelkey: When you joined Zilog, Summit was already spun off.

Rekhi: By the time I joined, they were being shut down. As a matter of fact --

Pelkey: Did you ever have an opinion as to why Summit chose to work with Metcalfe and not with Ungermann-Bass to put out their Ethernet?

Rekhi: That all predated me. When I came on board, there was this memo from somebody that Exxon is shutting down Summit, and there's a lot of equipment available. If any of us need anything, we should go

get what we want. So, as a hardware manager, I went and got some strobes and things and I just brought them for Zilog.

Pelkey: This is shortly after you joined Zilog.

Rekhi: That was in May of '81. You see, I left December of '80 from my job at Link.

Pelkey: Do you remember the name of the engineer who was at Zilog and went off to Corvus to --

Rekhi: Yeah, yeah, I know him. I can't think of his name. [Phil Belanger] We tried to hire him, by the way, at Excelan. I'll get you the name. The OmniNet people, they got confused also, they went into that Unix workstation business with a thing called Concept, and that was the other -- One of the problems I had with Inder, my partner -- you know, Inder, myself and Navindra Jain are the founders -- Inder was very wildly focused. We wanted to do the workstations, he wanted to do the servers, he wanted to do the OSI.

Pelkey: He had Fortune coming in with the Unix workstation --

Rekhi: He was real sure that, if you don't do it now, this market is going to be taken, and I kept saying that we don't have the resources. If we don't do what we need to be doing here -- and even in the networking, he wanted to do the Ethernet, Token Ring, Token Bus, OSI, XNS, TCP/IP -- and if it ended up you were doing lots of stuff and lots of under-funded, half-manned projects, and you know what happens when you are spending money and you have nothing to show for it.

Pelkey: In 1980 -- at the end of Christmas of '82, I was introduced to Sorcim by Jack Melchor, and in the spring became president.

Rekhi: You were at Sorcim next door to us then.

Pelkey: So this is a walk down memory lane for me, coming down, because if I went to the original building, and originally I looked at the fax and, oh, you're in building 6 now. But I haven't been down here since we sold it to Computer Associates in '84. But the same thing was at Sorcim. I mean, Richard Frank. They had 60 software engineers and we had 300 projects. It was --

Rekhi: -- a mess.

Pelkey: Because of -- there was an Indian who was a cofounder --

Rekhi: Oh, yeah. I can't remember his name, but I met the chap because he knew Inder. He had ported Supercalc to the MCZ machine. Zilog was also an issue -- they were doing their own OS when CP/M had become standard, and one of the things I did at Zilog at the end was add a CP/M board to the MCZ machine. So we had CP/M front end with us in our server, but there was nobody to sell. There was nobody to sell.

Pelkey: One can think about, at Zilog, I mean Zilog had their own product, but you had Ungermann-Bass, and you had Bridge, and you had yourselves, and you had this guy who went over to OmniNet. Metapath came out a little bit later, which was a nothing thing.

Rekhi: When I look back at myself, when I came out of Singer, I had been doing flight simulators for eight, nine years, and I lived and breathed these big, large systems. Also, I was in the world of (unintelligible) buses, star communications, I had never done anything like that. DRAMS, processors, so for me, I went to Zilog, that was my first time when I started to see things from a totally different perspective, but I was very driven in hardware. Nobody could touch me in making hardware work. So –

Pelkey: What was the source of your confidence?

Rekhi: Oh, I did a large number of systems, huge systems. I started working in '67 on mainframes and minis, and I moved from there to the large math processors. Flight simulators were, by and large, huge pipelines. It's easier to take something that works and make a solution out of it now than invent some air. And I also became very aware in '85 -- end of '84, OEM business was not a way to win in the marketplace at all. You have to -- the user has to approve your value. One of the problems I had -- we had 35 OEMS, and --

Pelkey: And you had no idea what was happening.

Rekhi: And the OEM sales guys weren't able to sell. Users, they will guide us into helping them sell, and if I have to help you sell, what do I need you for? That was our sense. When I did the first VMS end user package, and we sold it for $15,000, and we said: "If you go a board, software, and a one year free warrantee on the software," and everybody thought I was crazy. 'How can you sell a $15,000 worth of network connection over the phone?' By the way, this was sold over the phone. So we pioneered selling the packaged solution over the phone, and the way we did it --

Pelkey: And this was when, '86?

Rekhi: '85. So we go and run ads that say, TCP/IP, Ethernet for VAX, $14,995. And I run the same ads, TCP/IP, Ethernet with PC, $1,495, but it's TCP/IP and these huge machines, and the ads will run all over. We were spending about $30,000 a month on ads, which was a lot of money back then. And we guaranteed that if the stuff didn't work, we'd send somebody out to make it work for you. And if it's your fault, you pay for the trouble; if it's our fault, we pick up the tab. And we never had to pick up the tab. Never picked up the tab.

Pelkey: Wow, that's great.

Rekhi: And that turned out to be a very smart move, because what happened now, as I was selling to the OEMs the same damned stuff for $995 and I'm pricing it $14,995 for the end user. And you could make a lot of money. And you have the money to support and service. Not having a sales force in the field was a very smart move.

Pelkey: Oh, absolutely.

Rekhi: But you had a lot of work to make sure the product did work and did stay in the field. We had an excellent reputation for quality and support in those days.

Pelkey: This has been extremely helpful. Thank you very, very much for your time.

END OF THE INTERVIEW

