Oral History of Laurence “Larry” Boucher

Moderated by:
Tom Burniece

Recorded: January 27, 2015
Mountain View, California

CHM Reference number: X7398.2015

Tom Burniece: I’m Tom Burniece, a volunteer at the Computer History Museum. I happen to be the chairman of the Storage Special Interest group and I’m here today to interview Larry Boucher. Larry is an ex-IBM’er who was a founder of Adaptec, Auspex and Alacritech, three very famous companies, all started with an ‘A’, and I’ll end up talking about that.

Larry – please introduce yourself by talking about your family background, where you were born, where you grew up, where you went to school and we’ll go from there.

Laurence “Larry” Boucher: Thanks, Tom. I was born right here in San Mateo at Mills Hospital and grew up in San Carlos, went all the way through high school in San Carlos and then joined the US Navy. The Navy sent me to college. Originally, I was planning on going to Berkeley but when I was going through the review process of front of a chaplain, a general in the Marine Corps and an undersecretary of the Navy, the Marine Corps guy asked me where I’d been accepted and I said, “Well, Oregon State so far but I’m waiting for Berkeley.” He said, “Are you interested in going to Berkeley or are you interested in being in the Navy?” and so I wound up at Oregon State. <laughs>

Burniece: You’re saying you chose the Navy over Cal Berkeley!

Boucher: Well, I was already in the Navy; I’d enlisted in the Navy <laughs> and they were bothering to send me to school so I figured if they want to send me to school and they figured that Oregon State is good enough then that’s where I’d go so—

Burniece: You were actually drafted into the Navy?

Boucher: Actually, I went in before they started the draft. I joined the Navy in 1960 and then in 1961 the Navy sent me to ET [(electronic technician)] school. I went through boot camp and ET school and then [after attending Naval Training School in Bainbridge Maryland] in 1961, I went to Oregon State and was at Oregon State for two years. Each year they give you a physical and at the end of my second year they discovered I had a deviated septum. You can’t be an officer if you have a deviated septum so <laughs> I got washed out of the program and went back into the regular Navy. I was on what was called a kiddie cruise. I went in before I turned 18 so when I turned 21 I got out of the Navy and wasn’t in for that much longer after my two years at Oregon State. I finished the Navy and went back to school.

Burniece: Did you go back to Oregon State or someplace else?

Boucher: No. I went to San Jose State …

Burniece: San Jose.

Boucher: --and I continued engineering. I was in EE (Electrical Engineering) and at the same time I went to work-- actually I first went to work for-- I’ll leave a couple of these out, I went to work for an electronic distributor and then later on I went to work for Eitel-McCullough followed by Varian when Varian bought Eitel-McCullough. I was building traveling wave tubes and klystrons and magnetrons. I hadn’t finished my engineering degree so I was sort of a glorified tech, sort of in between if you will, and after spending a

year or so there I discovered because all the business was government business it’s basically feast or famine and engineers are sort of treated like dirt and the guys that really run things are the marketing guys. So I decided at some point this doesn’t make any sense, I’m switching into business, so I switched from engineering to business, and started working on a bachelor’s in business. Then I got married and for a little while I was going to school full time; then during the summer I decided to get a job for the summer and I went to work for IBM.

IBM gave me the opportunity of either a business job or an engineering job and I said, “Well, I’m sort of going towards business now” and they said, “engineering pays better” <laughs> so I took the engineering summer job and I didn’t leave for 11 years. <laughs> So yes, what happened was they hired me to do the magnetic deflection channel on an electron beam microscope and when I finished it they decided they really didn’t want to let me go, but there are problems after you get to six months. I wasn’t there yet but by the time I got to six months they had to figure out a way to keep me, which somehow they did, and after-- I don’t know-- a couple of years they hired me on a full-time regular basis. That was nice because then they started paying for school. I finished my business degree, got an MBA and then immediately went back to finish my engineering degree. By then I had decided I really liked engineering.

After the magnetic deflection channel they put me to work on the 2305, which was a replacement for the 2301 drum. It was a fixed-head disk with either one or two heads per track. These were fixed heads one or two per track, the idea with two of course being that you could transfer data twice as fast because we needed something that would go really quickly. So the two versions, 2305 mod one and two were one and a half and three megabits a second, which at the time was pretty amazing. The whole thing, which stood in a much-larger-than-refrigerator-size box, was either two and a half or five megabytes and this was definitely just for the new products that we were coming out with. For the 370s we were going to virtual memory and the 2305 was going to be the paging device for main memory.

Burniece: You went with Varian first and you were working on it sounds like microwave stuff. At that point were you really hooked into that world and the RF world …

Boucher: Correct, yes.

Burniece: You were a tech but you were still working on interesting stuff.

Boucher: Yes. It was a past a regular tech but I was basically working on getting final approval for the products. In fact one of the devices we were working on was a klystron for Sperry for their down-looking radar, and Sperry insisted that it be tested very thoroughly; it had to meet a certain set of specs. And so my job was to make sure that they all fit and then when they came to inspect be able to prove to them that they would meet spec.

Burniece: That sounds so much like what I would think about the way you think, from what I know about you, so you would be really hooked into EE at this point. You’d be thinking about a career in RF or something and yet at some point you got disappointed enough with the way they were treating you to say, “I’m going to business school” and that’s an interesting decision. Tell me a little more about that.

Boucher: Well, basically I wasn’t planning on getting out of electronics. I had been in electronics since forever.

Burniece: Did you do that as a kid?

Boucher: I started out with the standard crystal set and then built a five-tube superhet, so yes, I was really interested in engineering, and how I got into programming or computers at all is sort of interesting, because the first two years at Oregon State their only programming class at the time was in the upper division, so I never took a software class; I never saw a computer. Then when I went to San Jose their only computer class, a Fortran class, was a lower division class, which at the time was the introduction to computers. By then I’d fallen back in love with engineering. I had never fallen out of it; I just didn’t like the way they were treated. In the RF world, engineers working for the government were sort of dog food. At IBM it was a hundred percent different. You felt like you were one of the key people and it was just a lot of fun. So basically when I finished business school I learned whatever I learned about computers at IBM. The first computer I ever saw was when I was at IBM.

Burniece: What was that?

Boucher: It was one of the 360s. This was in ’68-‘69 so by then the 360 was out and our computer lab was 360s and I’d walk by the lab. I’d never been in to actually see the machinery but I’d look in the window and see the computers, when I working on this e-beam. When we finished the e-beam, they sent me over to work on the 2305 and that’s when I started to learn computers, which is another story. But the point is when I finished my MBA I just went right back into engineering.

Burniece: So you finished the EE degree after …

Boucher: Yes. I got the MBA and then I went right back. I still had a little bit to finish the bachelor’s in EE at San Jose and just about finished with that. I had applied at Berkeley to start the master’s and Berkeley accepted me for the next fall, which I couldn’t do and also finish everything necessary to get the bachelor’s from San Jose. So with a few classes still missing, I went to Berkeley and got my master’s in EE. <laughs>

Burniece: Did you ever actually get the bachelor’s?

Boucher: Well, it didn’t seem like it made sense to go back and get a bachelor’s after I finished the master’s.

Burniece: So you skipped …

Boucher: Yes, <laughs> the bachelor’s

Burniece: That’s an interesting story.

Boucher: I mean I didn’t really skip it because I basically completed about everything necessary but nevertheless I mean yes, I didn’t get the actual degree.

Burniece: You ended up with an MBA and a master’s in EE.

Boucher: Right, and I figured that was probably not a bad idea, given that I decided to stay in engineering <laughs> at IBM.

Burniece: There’s an interesting transition. What time frame did you make that transition?

Boucher: That was probably ’68 I believe, because I was only on the e-beam project for a year before I left. What happened is that when I went over to the 2305 I didn’t know anything about computers--I’d only just seen one. The first job that I had was to modify the control unit. Everybody was now working on the control unit-- the second-generation control unit with a B-box, and what they needed was to be able to drive the B-box with a [active] controller as quickly as possible. So my job was to take the original code--microcode that was written for the controller - and rework it to be able to work with the B-box. They handed me the listing for the microcode and said, “Go to work.” The rest of the group that they had put me in was busily working on the B-box itself and …

Burniece: You’re saying this was about 1968, ’69?

Boucher: Yes. The guy that sat in the office next to me, Sterling Ho, was a relatively senior engineer and what my boss, Ken Beeman, told me was, “Leave him alone. He’s busy. He’s got other things to do. You can ask these other guys if you’ve got questions on how the machine works and what you need to do in order to change it.” Well, Sterling took me under his wing, despite what Ken had said. He decided that he was going to teach me, and if it had not been for Sterling things would have wound up a lot different, because he basically gave me the training that I needed to be able to figure out how to do this.

Burniece: Elaborate a little bit more on that. Tell me a little bit more about Sterling and how he trained you and what the things were that he taught you that are really unique.

Boucher: Well, he basically answered questions, as I picked up this code; this was horizontal microcode. Since I’d never done any coding of any sort, whether it was horizontal or vertical made little difference to me and this was actually much more logical because horizontal microcode is written like logic and logic is something that is more like an electrical circuit, which I knew really well. So this actually wasn’t all that difficult and I could sort of look at it and begin to figure out what was going on, but he answered a lot of questions, so I was in his office initially on a daily basis asking questions. Then later on, after I thought that I pretty much knew how to do the job, I’d come back in after I’d done some amount of it and say - “Okay, so this is what I did. What do you think?” and he’d take a look at it and basically just say, “Yes, that’ll work but you can do better.” He made a big difference in terms of the way I thought about—

Burniece: He was pushing you a little bit.

Boucher: Oh, yes he definitely pushed me and we developed a relationship, as a result. Later he left IBM went to work for another company and I almost left IBM, because, he tried to hire me to go to work at a startup that was going to build a disk drive for STC [Storage Technology Corporation].

Burniece: Are you talking about the one that Jon Kevill [developed]?

Boucher: Right. He basically went to work for Kevill and tried to hire me and that’s what convinced IBM to hire me permanently even though we were in a hiring freeze. In a way, Sterling hired me into IBM after he had left the company.

Burniece: Talk about the 2305 a little bit more in detail and particularly when you say it was a replacement for a drum. It was actually a fixed-head shoe. Right?

Boucher: Correct.

Burniece: [for a] disk?

Boucher: It was fixed-head disk drive that spun at 5000 rpm, which at the time was very fast. It was scary actually when you’re standing there working on it <laughs> but yes, we basically got that working and then started to attach it initially to the high end of the 360 line, so the mod 91. It was really designed for the 370s but the 91 was sort of the precursor, if you will, to the 370 series so the mod 85 and the mod 91 were the two 360 devices that we attached it to. So that was where I first started - I initially was learning control units and disk drives and then once we finished I had to learn channels and CPUs, since that was the next step. I wound up living in Kingston attaching it to the mod 85.

Burniece: You got to be pretty much a broad-based guy at that point about disk drives. You learned channels; you knew something about the heads and disks; you knew something about the controller.

Boucher: Yes. In order to work on the products at that time <laughs> at least everybody I knew in IBM wound up fairly broad based. It was a smart group of guys and you really had to work across the board. We never had to go back and flip switches but you occasionally had to write machine language for the 360 in order to debug problems on a channel and control unit. So yes, you wound up having to know a reasonable amount about the entire system.

Burniece: I’m going to segue for just for a second because, just before we turned the camera on, you and I were talking about Dal Allan’s interview and how he was one of the four guys, who developed the original BOS (Basic Operating System) for the 360. You mentioned that it lasted for a long time. Tell me a little about how long BOS lasted and why.

Boucher: Well, BOS basically didn’t last very long as a commercial operating system but IBM used it as an educational device. When you first joined IBM, if you were joining as a programmer, the first thing they did was send you to school and they taught you assembler. So one of the projects that you learned in the process of learning assembly language was to write one of three programs for a basic operating system, either a scheduler, an interrupt handler or a dispatcher. The teacher in the class would allocate one of these functions to each of three students, they’d go off and write them, key them in on punch cards, and then they’d take their sets of punched cards and they’d load them in to replace those three elements in BOS. So while BOS was a very basic operating system, which is what the mnemonic means, <laughs> it worked very well for this function. Then what would happen is you’d run a tape-to-tape transfer, which would force all three of their components of the OS to operate and see if it worked.

So, [after] I had been at IBM for I would say maybe six years, FS had gotten started and I looked at …

Burniece: FS is Future Systems?

Boucher: Right. I looked at Future Systems and I said, “I don’t want anything to do with this” but the only jobs that were available were in that. Evans had basically said, “This is what everyone is doing” and so I had gotten to the point where people would at least sort of be nice to me, so I said, “Look. I want to go learn assembler.” I had never really gotten a good shot at assembler and so I went to this class. What happened was the BOS system that our instructor was using for the class was running on a mod 30 360 and the IT guys needed the mod 30 so they put MVT on it. When they did that my instructor said, “Well, that’s okay we did a dump of the BOS system and we’ll reload it in order to run this.” When they did that, they’d upgraded the disks to 2314s and the 2314 didn’t run. I mean the BOS system had been generated on a 1301 and they thought they could just do a disk-to-disk copy and they could run it.

So I said, “No. You’re going to have to regen BOS”. This was in the mid 70’s and it had still been operating across all of IBM’s programming centers as their teaching device. So I went to the instructor and said, “You can no longer do this culminating class project”. My instructor was at a loss; he didn’t know how to do the final class job so I said, “I’ll tell you what. There’s still enough time. Let me out of the rest of the class and I will build you an emulator that you can run on top of MVT to emulate BOS so that your students can now run their programs, their elements of BOS, on top of MVT.” And he said, “Sure.” He didn’t believe it but in any case I got that done, which is one of the things that I really enjoyed doing if I look back over my career.

Burniece: Was that a pretty unique idea to build an emulator [at that time]?

Boucher: Yes, and to build it in the amount of time that we had to emulate the entire operating system on top of another operating system, where you have to handle and really treat basic machine level interrupts, is nontrivial. It’s not all that complex an operating system but it was a challenge to get it done in time for these students to be able to run their programs.

Burniece: Was that one of the first emulators ever done at IBM or had that been done before?

Boucher: Oh, no. They had built emulators to emulate the earlier systems. I forget the numbers now. It was before my time but the machines before the 360 they emulated on 360s so that you could still run programs for those systems and those were already in existence.

Burniece: So that idea was not totally unique but it was a unique project that you [personally] took on.

Boucher: Yes, and it kept BOS alive, so BOS lived all the way up to the late ‘70s, in reality within the educational system at IBM. God knows how long they used it afterwards, just to allow students to continue to be able to build those components.

Burniece: That’s fascinating. Were you kind of getting hooked on computer programming and operating systems at that point?

Boucher: I never got into programming. I enjoyed it but I never did it as a job. I had already done PL/I programming. Just about everybody in IBM I think had to use PL/I at one time or another, because it was

so simple. It was very inefficient but you could make something happen very quickly. So if you just needed to get something done at a high level on a CPU you could use PL/I and have it done in a day.

Burniece: That was a good segue; I’m glad we asked that question. Let’s go back to the 2305. How successful was the 2305 program and how long were you on it?

Boucher: I was on the 2305 until pretty much we finished. It was a very successful program. I don’t think there were many of the higher-end 370 systems that didn’t run a 2305 because MVS really ran a lot better if you could page memory.

Burniece: I know that some of the 3330s and 3330 mod 11s had a shoe in the bottom, which was kind of this idea of a head-per-track disk as the bottom disk in the system. I might be wrong about the model. It might be [the] 3340 …

Boucher: No, [there was no shoe in the 3330s]. That was the 3340, [but it was] a whole different thing than the 2305].

Burniece: Was that taken off the 2305 concept? Was that where that came from, that head per track?

Boucher: No. The head per track was purely a high-speed function to figure out how to get a disk drive to be fast enough to do what we want for paging main memory. As I say, it was a replacement for the 2301 drum, which was not paging but was used as a spooling device for the earlier systems.

Burniece: The 3340 was the “Winchester”, when IBM put [in-contact take-off heads and disks together into a fixed disk HDA (Head-Disk Assembly) for the first time]

Boucher: Yes, the Winchester drive, which was Ken Houghton’s baby. The name came from the fact that when they were originally thinking of it there was a huge fight about whether or not you could really have fixed disks or if you had to have removable disks. The 2305 was fixed because it was a very specific product, but the controller we were building was the same controller that was used for the 3330, [which had removable disks]. So the 2305 and the 3330 were developed at the same time on the same floor; the controller was the same controller, just different microcode and slightly different hardware.

Burniece: So you could mix those two drives in the system then. Right?

Boucher: No, not on the same controller because the front end of the controller was different, until we built the B-box, which was the second generation. In the mod 1 controllers there was no B-box. The controller literally went to the data separator on the drive so the 2305 and the 3330 microcode on the controllers was quite different.

Burniece: Did you work on those?

Boucher: Yes. After I finished the 2305 controller, I went to Iceberg, which was the double-density 3330

Burniece: [The 3330-]Mod 11

Boucher: … but before that I did the attachments of the 3330 to the 135 and the 125, so yes, I worked on the 3330 as well as the 2305.

Burniece: At this point in your career, you’ve been at IBM maybe eight years or something like that and you are now getting into the subsystems side of it again, [including] the interfaces, which became a career for you.

Boucher: Right, yes. For the attachments on the 125 and 135 we were writing all of the code in the CPU microprocessor in order to cut the cost of the overall system, so there was no external controller for the direct attachments of the 3330 to those devices. Since the channels were integrated [I was] basically deep in the guts of the channel and the CPU.

Burniece: How long did you stay on the Iceberg program then and where did you go from there?

Boucher: I was not on the Iceberg program for very long. I basically was in the front end of it because I was working for Jim Carothers who owned the project. Actually I was working for George Ahearn but sort of dotted line to Jim. At the time people were thinking there’s only going to be one more generation after this, since we’re going to run up against the wall from a physics standpoint, which is really interesting, if you think about it. <laughs>

Burniece: That was wrong in those days but people believed it.

Boucher: Yes, <laughs> when you look where we are now, but in any case that’s when FS was getting started and that’s when I went over and learned 360 programming, operating systems, etc. Then I got dragged out of that to do the tape attachment to the 125, which wound up being the last job I had within IBM as an engineer. From there I went into management at IBM and was responsible for the next generation of CAD software to develop what we called Emerald, which was a CMOS technology.

Burniece: When did you join Shugart--

Boucher: ’79.

Burniece: Seventy-nine. So you were there for 15 years?

Boucher: I was at IBM for somewhere around 12 or 13 years.

Burniece: Started in ’68?

Boucher: Yes

Burniece: So you left them somewhere around ’77 or …

Boucher: No. I think I left in either December of 78 or January of 79.

Burniece: Tell me a little bit more about some of the people at IBM. You’ve mentioned several. Who were some of the more interesting people you met and really influenced you along the way at IBM?

Boucher: Well, in general everyone that I worked with was an influence. I mean the engineering team at IBM was an amazing team; they were really impressive. I guess Wally Bass from an engineering standpoint.

Burniece: Tell me a little bit about him.

Boucher: Wally was the key designer of the microcode and micro-engine of the 2314 and he had very early on proposed an architecture called DABS [Direct Access Block Storage], and it seemed like a much better idea to me than count-key-data (CKD), and in fact <laughs> I don’t know if I ever came up with any ideas in general. Other smart people came up with the ideas, I just incorporated them, but in any case Wally was definitely an influence because …

Burniece: When would that have been? You’re talking about the original idea of block addressing. When was that?

Boucher: I think I probably ran across DABS in the very late ‘60s, maybe ’69, ’70, ’71, somewhere in there …

Burniece: Okay.

Boucher: … but nothing ever came of it. IBM stuck with count-key-data, which made sense initially just because of the fact that real estate on the disk drive was so precious and you wanted to use as much as possible, so track length records were not atypical. In any case other people were, George Ahearn who I worked for quite a while, and John Adler was another guy I worked for …

Burniece: Do you want to talk a little bit about Ahearn, what he did?

Boucher: George was a second-level manager. When I initially went to work on the 2305 he was my second-level manager and Ken Beeman was my direct manager. I worked for Ken for quite a while and then went on staff to George, who taught me a fair amount about the way IBM worked.

Steve Behman was another guy, who was a significant influence. Steve was one of the architects at IBM and a black sheep for sure but was the guy that you usually called, if you wanted to figure out why something wouldn’t work. He’d figure out all the things wrong with it, which usually meant once you’d covered everything you’d figure it out and you could make it work. So <laughs> it was …

Burniece: You say he was an architect. On the systems side or on the disk side?

Boucher: Systems side, yes, a systems-level architect. His background I think was math but in any case a real sharp guy. Many of the people, who are household names in our industry were there and were both a direct and indirect influence on me, as a result of the great environment that they built.

Burniece: Did you know Al Shugart when he was still at IBM?

Boucher: No. Shugart had just left when I …

Burniece: You never really knew him at all.

Boucher: No, I didn’t get to know Al until I was trying to sell chips to him at Adaptec.

Burniece: How about some of the guys like Jack Harker? Did you get to know Jack?

Boucher: Yes. Jack is one of the “Jacks”: Jack Harker, Jack Kuehler, and Jack Hildebrand. I worked directly for Jack Hildebrand for a while. [They were some of] the guys who originally built the RAMAC, that built the industry in San Jose. All were impressive, really sharp guys. They weren’t as much of a direct influence because they were very senior to me, and while I knew them and in general they knew me, it was mainly through meetings that I would be involved with them.

Burniece: Anything else you want to talk about on IBM before we move on to the next step?

Boucher: I think that’s probably it for IBM.

Burniece: When you left IBM you went where?

Boucher: Shugart Associates.

Burniece: To Shugart Associates directly and who brought you there?

Boucher: Hank Zauderer. Hank was the VP of engineering at Shugart Associates and he brought me in to work for Yoshi Narahara,

Boucher: [Yoshi] was head of research for Shugart Associates so I worked on staff to him for six months and then became director of design services working for Hank Zauderer and …

Burniece: Was the original job to work on Al’s dream about a computer or was there already something else?

Boucher: No, it was already a floppy disk company.

Burniece: by that time?

Boucher: Right, and the job that I had at design services was basically the engineering [development] computer so I had to specify it, get it up and manage it, along with all the CAD/CAM/CAE software. In addition, I was responsible for the disk controllers. When I first got started, I inherited the job of building a disk controller for our first eight-inch [hard] disk drive, the Shugart SA1000. At the time they had already built a team to do the low-level hardware design for the SA1000. Richard Steele was the guy managing that and he had already chosen SASI as the name of his device, which was “Shugart Associates Storage Interface.”

Burniece: So he picked the [acronym] “SASI,” S-A-S-I.

Boucher: Right.

Burniece: I thought you did.

Boucher: No.

Burniece: That was just the name? You had to come up with the interface?

Boucher: Well, no. That was the name for his interface, which was at the serializer / deserializer level, so he basically was doing a data separator and had named the read channel interface [SASI]. At that point in time, as I’m sure you’re aware, all of these disk drives were being sold with the read channel as the standard interface and the customer had to build their own higher-level controller from that point. We’d give them a higher-level controller [reference] design but, because of the fact that they were building a system, usually they’d build the controller into their system, and this was a huge problem as far as I was concerned.

I wanted to build a [higher-level controller] but we were building a low-level controller that would be just like all the controllers that had been built up to that point, which was basically at the read channel or a bit higher at the SMD level, if you included the data separator. The project was already under way but I looked at it and I said, “What we really need to do is build a different sort of a controller” and so that’s the point where Wally’s influence came into play. I wanted to build an equivalent of DABS, so that we could then sell a disk drive with an integrated controller.

With an integrated controller, a new disk drive could be implemented much more quickly by the user [or] end customer, since we could update just the disk drive and he could [immediately] make use of it. At the time, if you wanted to sell a customer a new disk drive it would be on his next system that he’d integrate it because if you just changed any characteristic of the disk - more cylinders, more tracks, more heads, anything - he couldn’t do it because the existing controller was absolutely locked to a physical disk. I wanted to basically isolate the physical qualities of the disk so that we could upgrade the customer on his present system.

I figured we’re going to have to do this, so I put together an architecture for the next-generation controller to be used for the SA1000. I felt Shugart Associates System Interface was the perfect name, so I wanted the SASI name. I had to figure out a name for Richard so I came up with SADI for Shugart Associates’ Device Interface and I went to Richard and I said, “Hey, Richard, would you mind if we switch so I can have SASI and you take SADI?” and he signed up, which I really still appreciate. <laughs>

Burniece: That’s fascinating. So basically he had come up with the name but in a certain sense had misnamed what he was doing from an architectural viewpoint and you wanted to swap them. You said, “You’re really a device interface so why don’t you be SADI and I’ll be SASI, which is more at the system level, and I’m going to start working from at that point, with block addressing and all that stuff” …

Boucher: Right.

Burniece: So he came with the name but you actually ended up [developing SASI].

Boucher: Right. <laughs>

Burniece: And of course S-A-S-I became SCSI [Small Computer System Interconnect] at some point, which is one of the most important storage interfaces of all time. Talk a little bit about your team and how you did that, plus what were some of the concepts that you built into SASI.

Boucher: Well, that’s really interesting. I couldn’t [actually] build the integrated controller because I was committed to build this low-level controller for marketing. Marketing said, “No. We have to have that low-level controller.” My team wasn’t so large that the management job was everything, so I had]some personal] time …

Burniece: A half dozen guys or something?

Boucher: Well, in the controller group I had maybe a half dozen guys but I also had the computer, the software and everything, plus all the administration associated with that. I had good managers in those areas, so my overall job wasn’t that demanding …

Burniece: So you still had some time to put into the actual design …

Boucher: From an architectural standpoint. I split the interface into three pieces, a logical interface, a physical interface and then a third small sliver because I knew that our physical interface at the upper level had to be parallel then but I knew that at IBM, FS was going serial. I figured serial ultimately is the way things are going to go, so we needed to switch from parallel to serial without any trouble. I gave Bernie Nieman, one of the managers who worked for me, the job of doing the architecture of the physical interface, which was sort of a standard bus and tag style interface. Bernie gave most of it to one of his guys, Jim Korpi, who did the actual logic design from an architectural standpoint for the spec. Then I did the logical spec and we got all that done, so now I had a complete spec for a [high-level] controller at the SASI interface but I had nobody to build it, because the rest of the team was designing [the low-level] controller, including Bernie and Jim. So I had to give it to an outside company; I gave it to a company by the name of DTC.

Burniece: You outsourced it.

Boucher: Yes. David Tsang had started this company and I needed to get this controller built and couldn’t build it myself, so I got Tom Mackman, the marketing guy who was on this project, to agree that once I wrote the spec and wrote a document to justify building this controller, if marketing was convinced of the justification they would support it.

Burniece: How did you find this company you outsourced to? Did you know these people?

Boucher: Actually, Tom Mackman knew them and I knew David because I had met him when he was still at MSC.

Burniece: What was MSC?

Boucher: Microcomputer Systems Corporation. It changed to Xebec over time.

Burniece: Yes, Jim Toreson’s company.

Boucher: Right.

Burniece: So he was at Jim Toreson’s?

Boucher: Yes, Dave worked for Jim but that’s another story. So I gave the project to DTC and DTC built the controller. Then, just to make a long story short, we started to use the controller and I think Tom recognized that this was ultimately the way we wanted to go. So I wrote a business plan to build a controller company inside Shugart. By now, Jim Bochnowski [the CEO of Shugart Associates, when I started] had left and Jim Campbell had taken his place. Jim [Campbell] was relatively new to the company; he’d only been there for a few months. When I finished the business plan and took it to Jim, Shugart Associates was having some real problems.

Burniece: Al was long gone back then.

Boucher: Oh, Al had in the meantime started Seagate—

Burniece: So Al had already started Seagate.

Boucher: In fact, before I left Shugart, Seagate had already come out with a five-and-a-quarter-inch [hard drive] ,,,

Burniece: So we’re now talking about 1980, ’81, something like that.

Boucher: This was ’81. I left Shugart in ’81. I had proposed to Jim Campbell that we build a division inside the company to do controllers.

Burniece: This was essentially then the germ of an idea that eventually became Adaptec.

Boucher: Well, yes, because …

Burniece: But at that point it would be a division of Shugart.

Boucher: Right, and that was because Shugart was having serious morale problems. Bochnowski had tried to clear things up by going to four ten-hour work weeks to see if that would make people feel better and it didn’t, so he brought Campbell in to see if he could get things rolling. One of the things Campbell had done is said, “We want to do internal startups” and that’s why I had …

Burniece: So he actually had proposed having a little bit of an in-house entrepreneurial kind of a system.

Boucher: And this made all the sense in the world, so I brought the business plan to him and he gave it to George Sollman, who was our marketing VP. George was not just a little bit against it; he was vehemently against it. And the reason for that was, in the meantime, some of our key guys had left to start Quantum and George was not at all happy about that, so his statement was, “If we build a controller and sell it, it’s going to help Quantum get going.”

Burniece: So he could already see how that could help Quantum. That’s interesting.

Boucher: And I said, “If we build a controller and we are able to sell it to Quantum customers, we are also selling the SA1000. If there are any problems between the controller and the Quantum eight-inch drive,” the first product Quantum built, who do you think is going to win?” It made sense to me, because why wouldn’t the customer try our disk and see if it works, but George didn’t sign up for it.

Burniece: Had Shugart at that point already started thinking about and working on hard drives, as well as floppies?

Boucher: You mean Shugart Associates?

Burniece: Shugart Associates.

Boucher: Oh, no. That was my job was to build controllers for the SA1000, which was an eight-inch hard disk.

Burniece: Oh, that was a hard disk. I’m sorry.

Boucher: And then part of the team that built that disk, some of the members of it, left to start Quantum to build a competing product, so George was saying, “Hey, I don’t want to help the competition.” If I were the competition, I’d hate it if somebody bought a Shugart controller” <laughs> but I lost the argument, so I left to start Adaptec.

Burniece: Before we leave and get on that point, what are some of the other interesting things that happened during your tenure at Shugart, which sounds like maybe three, four years. Right?

Boucher: Two years. I guess the other interesting thing is while I had DTC building the initial controller, which was all TTL, it was obvious that was not going to work over the long run. The controller was much more expensive than the disk drive and that made no sense from my standpoint. At the time, the way you built a disk controller was you bought an 8X300 from Signetics or you designed a bit slice and you put a lot of bipolar memory behind it. You spent a thousand dollars in parts to build the controller, plus you needed a stupid amount of power in order to drive the thing. The entire data flow went right through this very high-speed bipolar processor. As a result it took two revolutions to read a track; you could only read every other block. So at the same time as DTC was building this sort of a disk controller, I was designing a CMOS controller in house for the next-generation SASI disk controller, which was going to cost less than the parts cost of a bipolar processor.

Burniece: In a certain sense, what you outsourced was almost more like a prototype than your actual design?

Boucher: Absolutely. It was designed as a prototype but it was going to be a prototype that we’d sell while ….

Burniece: The bottom line is it was not your end product. You needed somebody to go and do this and prove it works while you go out and design a CMOS version

Boucher: Yes. At the time nobody thought you could build a CMOS disk controller. All of the disk controllers were still bipolar.

Burniece: They were 2901s and that kind of thing. Right?

Boucher: Correct, and that’s what I wanted to do - basically build a CMOS controller. We hadn’t built anything like this inside IBM but we had already built the B-box and that was the thing that made it clear that you could do this because the B-box integrated an awful lot more of the data flow in hardware. So it wasn’t too hard to figure out how to go one step further and put enough of the data flow in hardware that you could use CMOS technology for the high level processor, rather than having to go with bipolar. The CMOS was then fast enough to manage the disk data flow.

Burniece: So you were already working on a CMOS version of your controller for SASI, when you came up with the business plan to create a division?

Boucher: Right, which is why it just seemed to me …

Burniece: Was [Jim Campbell] okay with what you were doing; he just didn’t want it to be a separate division?

Boucher: Correct, but I felt like this was too big a business. In other words we could make a lot of money selling controllers not just for our disks, but for lots of disks. I didn’t want to do it just for our disks and I might have never done it, if Jim hadn’t instituted this policy of in-house entrepreneurship, but once he’d done that …<laughs>

Burniece: Like opening Pandora’s Box …

Boucher: That’s right. <laughs>

Burniece: You were now going to be an entrepreneur one way or another.

Boucher: <laughs> Right.

Burniece: That’s fascinating. When you finally decided you had to leave to go and do what you wanted to do, how did you raise the money? I assume you’d never been involved in raising money for a company before.

Boucher: Correct.

Burniece: How did you do all that?

Boucher: Well …

Burniece: … and with who?

Boucher: I knew Jim Bochnowski, because he was the CEO of Shugart Associates for most of the time that I was there. Jim had gone back to TVI, so once George had shot me down I told Jim what happened and Jim gave me a couple of good pieces of advice. He said, “First of all, go see a lawyer” and he gave me the name of a guy at Wilson Sonsini. Then he said, “Go see an accountant”, because I’d shown him the basic business plan and he said, “You need some help with your business plan.” I mean I had an MBA but that didn’t qualify me to … <laughs>

Burniece: He’s now saying, “Start thinking about how an MBA operates …”

Boucher: Yes, because I had now been an engineer for long enough that’s right. <laughs>

Burniece: That’s all good advice.

Boucher: <laughs> Right, and those were very good pieces of advice. He also gave me the sense that it was unlikely I would get funded directly right then, because I needed to have a team. Jim introduced me to Dave Marquardt to broaden TVI’s knowledge of us.

Burniece: He was early in Seagate, wasn’t he at one point?

Boucher: He was an early investor in Seagate.

Burniece: So by that time he probably had already made some [other] investments.

Boucher: He’d made a couple. He’d made one in a controller company as well - Spectra Logic. His words to me, which I’ve never forgotten, were, “What the world needs is a good five-cent cigar and another controller company” so that … <laughs>

Burniece: <inaudible>

Boucher: He was not all that excited about controllers, because of the company that he had invested in, Spectra Logic. Not the company that exists today, but the one that built DEC-compatible controllers in the ‘70s.

Burniece: It was not Toreson’s company, it was somebody else.

Boucher: No, it was a competitor to Toreson. They were doing SMD drives and DEC attachments. Toreson initially did DEC attachments and then got [sued].

Boucher: In any case, it became clear to me, since they were the only venture guys I knew and they weren’t all that excited, that if I was going to do it we were going to have to do it on our own. There were two guys who both worked for me at Shugart that had been involved in the planning process and really wanted to do it, Wayne Higashi and Bernie Nieman, plus a guy that I had met - I’m not sure how I met him now but outside Shugart - who knew the IC industry reasonably well.

Burniece: And who was that?

Boucher: John Hulme. I had gotten to know a little bit about the industry, because I was in the process of negotiating the generation of the mask set and building of the chip. We were doing the design - at the time you’d draw every transistor - and then we would generate a tape, in order to build the mask set. The mask set and the building of the chip itself were getting farmed out. So I needed somebody who knew that industry well enough that he could take care of that end of it, so I didn’t have to worry about that in addition to everything else, because we had to do all the hardware design. So Bernie, Wayne and I all took out maximum loans on our houses and then …

Burniece: This was the proverbial mortgage your house to start a business.

Boucher: Right, and then we figured we’ve got six months that we can afford to live, so we each got a hundred thousand dollar loan. We put fifty thousand in the company, so we had a hundred and fifty grand in the company, and we had fifty grand each to live. We figured we can probably live for at least six months on the fifty grand, maybe longer but we need to start looking for jobs after six months.

Burniece: So the three of you were the founders.

Boucher: Well, the four because of John Hume …

Burniece: But did John invest?

Boucher: John did also but I don’t think he had to mortgage his house …

Burniece: The four of you then founded the company and put some money into the kitty, so you would have six months to either raise some money or go do something else.

Boucher: Right, and ….

Burniece: When was that?

Boucher: That was in 1981.

Burniece: Early?

Boucher: Mid 1981.

Burniece: Mid 1981, so that was the founding of Adaptec.

Boucher: Right.

Burniece: At that point had you left Shugart long before, so you had some time in the middle or did you …

Boucher: No. You can’t get a loan if you’re not working and as soon as we had the loan we needed to do something <laughs> so yes, we got the loans and left.

Burniece: Did you have some difficulty with Shugart on leaving and going off and starting this thing? Had you actually started [before left ...]

Boucher: No. When I …

Burniece: … Shugart?

Boucher: No. When I left they were worried about my doing what Quantum had done because Quantum had hired a team out of Shugart, but I went to Jim when I was doing this and I said, “Look. You guys shot me down. I really want to do this and I have no interest in raiding you guys. These two guys were a part of building this business plan that we brought to you and they desperately want to do it too, but what I’ll do is I will tell you that if anyone comes to me and says, ‘I want to come work for you’ from Shugart I will tell them, ‘That’s fine. I’d love to talk with you but I won’t talk to you until you’ve told your manager at Shugart that you want to talk to us.’” Jim felt that was going way above and beyond. He felt very good about it. He said, “Hey, I’ve done a startup. I know just how you feel” so we left on really good terms and Jim felt we were doing the right thing, given that he hadn’t let us do it internally.

Burniece: What happened next? Did you immediately start walking up and down Sand Hill Road for money or …

Boucher: No.

Burniece: This is pretty early for Sand Hill but …

Boucher: I basically had felt that that was not going to be in the cards, which is why we did what we did. So we found a cheap place that was an old converted apartment building. They’d covered up the stalls for parking and turned it into a cabinet-maker shop and then the apartments had been turned into offices, so we rented one of those because it was really cheap.

Burniece: Where was that?

Boucher: That was in Campbell, right off of Campbell Avenue. I forget the name of the street but it was just before you get to the creek, so the place was backing up on the creek on the side away from Bascom.

Burniece: Okay.

Boucher: And at the time Wayne and I bought a Cromemco system, which was still using eight-inch floppy disks for the disk drive, and that was our word-processing system. We used WordStar as the word processor, in order to do the business plan, because we had to start the business plan from scratch. We had to do everything from scratch but we’d already done an internal business plan, so we knew basically what we were doing, At the same time we had Arthur Andersen helping us from a consulting standpoint,

so they were looking over the business plan to let us know if it made any sense. That was a huge help, because, while I could do the engineering half without too much trouble, what made real sense from a business, I only had the academic view of how to put financial plans together, with numbers that made any sense at all --

Burniece: That advice you got about get a lawyer and an accountant [was right on] …

Boucher: Oh, it was absolutely key. Without that I don’t think it would have come to pass. By the time we had a business plan that I felt was pretty good, we took it back to Jim, who felt that he would be slightly conflicted, so he gave us to Dave. I thought, that’s not good, given what Dave had already told me about <laughs> his feelings about controller companies. I’d also gone to Weiss, Peck & Greer, plus a few other companies. I have forgotten who else I had talked to but the other key one was Weiss, Peck & Greer and the two of them ultimately funded the company.

Burniece: So they were your two initial investors?

Boucher: Right.

Burniece: TVI and …

Boucher: … Weiss, Peck & Greer.

Burniece: And then Marquardt decided he’d go ahead and try this five-cent cigar, huh?

Boucher: Yes. He basically finally decided this was enough different from what he’d seen in the past because we were talking about building a controller, moving the decimal point on cost, and one that would literally bolt onto the side of a disk drive. Our initial controller was going to interface the ST-506 and just bolt onto the side of the ST-506, so you’d have an intelligent drive.”

Burniece: Your initial controller was an ST-506 controller rather than a SASI controller.

Boucher: Well, a SASI controller for an ST-506. By now, Seagate had announced the ST-506 and we were trying to go for low cost. That was a key component and we could see that we could fit the form factor. Within a month of when we started Adaptec, IBM announced the PC and Apple took out their full-page ads welcoming IBM into the market. It became clear that the five-and-a-quarter-inch drive was ultimately where things were going to go. The PC came out with floppies, two five-and-a-quarter-inch floppy drives, but it seemed obvious to me that it wasn’t going to be that long before there was a five-and-a-quarter-inch hard drive in a PC. So making an intelligent one was a good idea plus Wang, CPT, Altos and everybody else was already switching to five-and-a-quarter. The five-and-a-quarter-inch drive just made all sorts of sense from their standpoint and our initial market from our standpoint was going to be those guys. Those were the guys that were the big market for hard disk drives at the time, all the word-processing guys. Which was Interesting, because we weren’t doing our word processing on a Wang.

Burniece: A number of the pieces that you’ve just talked about in the last hour and a half or are starting to come together - your initial fascination with word processing, with microcontrollers, with a higher-level

interface like a B-box. Using some of those principles, you end up starting Adaptec right about in the boiling point of the five-and a quarter inch hard disk world starting to take off with Seagate, Quantum and everybody else that entered at that point in time in ’81?

Boucher: Right.

Burniece: You’ve got to Dave Marquardt to back you. This had to be some fun time.

Boucher: Yes. I mean and one of the …

Burniece: Probably a little scary but fun and exciting.

Boucher: Yes-- <laughs> no. It was scary. There were scary times of course but ultimately it was a lot of fun and the industry was really accelerating.

Burniece: You were just talking about when you formed Adaptec, you had the four founders including yourself, put some money in the bank and bet on six months to raise more money. You started talking to a couple of people like Dave Marquardt at TVI and Weiss, Peck & Greer. How did you get it funded and how long did it take?

Boucher: So Bob Loarie at Weiss, Peck & Greer took me to talk to Stu Mabon and after that he …

Burniece: Stu was at Micropolis.

Boucher: Correct, and I believe that Weiss, Peck & Greer had an investment in Micropolis, so they used Stu as a consultant and Stu and I got on very well. He’s a sharp guy and I enjoyed him. I guess he gave the go-ahead from his standpoint, because Weiss, Peck & Greer decided to invest and then TVI decided to invest, so Bob and Dave were the two guys that started—

Burniece: They were the initial board members probably with you then. Right?

Boucher: Correct.

Burniece: Did you have just the three-man board or did you have somebody else on it also?

Boucher: They asked for an outside board member, so we brought Ferrell Sanders on.

Burniece: Ferrell came in?

Boucher: Yes, that was a really smart thing because Ferrell understood sales, which none of the rest of us did, and he was somebody who was a known quantity to all of us because I had worked with Ferrell at Shugart Associates. Ferrell had left and was doing some investing on his own, I think in real estate at the time. It was before he’d gotten into venture capital, so ….

Burniece: So he was truly then an outsider

Boucher: Oh, he was pure outside …

Burniece: He was not an investor at that point; he was an outside guy?

Boucher: Correct.

Burniece: How long did it all take where you actually got to the bank and could stop worrying about running out of your own money?

Boucher: It was close to the six months. We spent the first couple of months getting the business plan to what we really believed in, which Arthur Andersen helped a lot with, and then I took it to TVI and I think TVI actually is the one that introduced me to Weiss, Peck & Greer. I spent the rest of the time with Bob Loarie and Dave Marquardt until they decided that they were going to invest.

Burniece: Your six-month plan worked.

Boucher: Yes, barely.

Burniece: How fast did you build the team and what was the first year or so like?

Boucher: Well, Bernie and I designed the controller and John put us together with a company that could then take our logic design and build our ICs. However, the first thing that we did was we built a TTL emulation of the full design.

Burniece: Is that the same one you were working on when you were at Shugart or did you restart the TTL?

Boucher: Oh, no, no. This was neither ...

Burniece: It’s also for an ST-506. Right?

Boucher: This is for an ST-506 although it would have worked for the SA1000 —because they both interface to the data separator. What we did was a TTL emulation of what we wanted to do in CMOS, so it was bigger and much more cumbersome than what you would have for just a TTL controller, because we were actually trying to prove that you could do this in CMOS. So we designed it as it would be done in CMOS. We used an 8085 and had an 8085 on our emulator that was running the code that was driving the TTL that ultimately would be a piece of CMOS. It actually was three pieces of CMOS. Our design was a data separator, a SERDES and a sequencer, in order to let the 8085 do the data-flow control.

Burniece: You put the three chips together with the 8085 …

Boucher: Oh, I …

Burniece: … system?

Boucher: Yes. In fact, the sequencer was the chip that got us from the SERDES to the data bus, a standard SCSI data bus, so it was the SCSI chip. So we had a SCSI sequencer chip, a SERDES and a data separator.

Burniece: At this point you were already designing a SCSI design?

Boucher: Oh, this was a pure SCSI disk controller.

Burniece: How long did it take to do all this, [including] the emulation to get it to the point where it was ready to go, get the CMOS design done, taped out and …

Boucher: We basically had the TTL version working at about the same time we got funded.

Burniece: That was pretty quick!

Boucher: Oh, yes.

Burniece: Six months or less.

Boucher: Yes, yes, we were not working eight-hour days and five-day weeks. <laughs> This was relatively important, since our houses were on the line. <laughs>

Burniece: I guess you did have a little incentive here.

Boucher: Oh, yes, a lot of incentive. <laughs>. We basically were able to show a working prototype in TTL of the three chips we were designing to the venture community before we got funded and that, plus the business plan, is what allowed us to get funded. Then we took those three parts, had the logic converted, taped them out, built mask sets and built the chips.

Burniece: When did you get the first chips back?

Boucher: It was about I would guess 14 months from when we started the company.

Burniece: We’re now somewhere in the second half of ’82.

Boucher: Right.

Burniece: Now you’ve got a chip set and …

Boucher: Yes, fairly late in ’82. Yes.

Burniece: You’re now starting to show this to people?

Boucher: Well ….

Burniece: Was the reaction good or was it …?

Boucher: No. By now we’d done the second round and we’d built enough of a company to be able to start to market the product but we wound up with what Jeff Miller, our marketing VP, has since then termed as Black Friday. We got the chips back, got everything together and tested it out and discovered we had a major problem that initially looked like it was going to take a turn, which would mean we didn’t have the money and time waits for no man. By now Xebec had a great, big monster thing that they called their equivalent to what we were doing and ….

Burniece: Did they have it in the market or was it at that point just talking?

Boucher: I can’t remember if it was physically in the market, either that they were just beginning to sell it. I don’t think they had customers for it yet but this was another thing. It still had the same problems as every bipolar controller; it could only read every other sector, which means it took two revolutions to read a track, and we had a controller, if we could make it work, that would read a track in a revolution, which means that we could provide a significant performance [advantage], compared to anyone else. I’m not sure how important that was to most people at the time because anything faster than a floppy was in general what people were looking for and a hard disk was so much faster than a floppy that twice as fast maybe wasn’t that big a deal, but nevertheless ….

Burniece: It’s always a problem when you think you’ve got something that’s a lot faster than anybody else but somebody else has got something that’s good enough.

Boucher: Right, exactly, <laughs> and that’s where we were with the Xebec product.

Burniece: Black Friday was when you thought you had a big problem?

Boucher: And by Monday we had figured out how to solve it.

Burniece: You had it figured out by Monday.

Boucher: Right, <laughs> so ….

Burniece: Sweet Monday. What was the solution?

Boucher: I no longer remember. It was something not that difficult to get the data separator operational. It was just a basic data separator issue that initially looked like there was nothing to be done except turn the chip, but then I think Bernie figured out how to externally do something that ….

Burniece: You didn’t even have to do a chip turn?

Boucher: No.

Burniece: You basically figured out how to patch something.

Boucher: Yes, we patched it.

Burniece: And it worked.

Boucher: Yes.

Burniece: Did it work almost completely at that point?

Boucher: No. It did work completely. That was good enough.

Burniece: That was almost “miracle Monday”.

Boucher: Yes, it was a huge relief. I have still not forgotten the feeling, when we discovered the problem, because initially it looked like there wasn’t a solution, but I have to give Bernie credit; he came up with a creative solution. Sorry. I can’t remember now exactly what the problem was.

Burniece: What was his background? What was he good at?

Boucher: He was a logic designer. He went to Arizona State or University of Arizona, one of the two, and had been in logic design for his entire career.

Burniece: How quickly then did the market begin to develop for you or did it take a really long time to start getting people to adapt this?

Boucher: It developed slowly. We competed for the IBM business and lost to Xebec.

Burniece: They did beat you out at IBM.

Boucher: Xebec beat us out in the first round at IBM and …

Burniece: That would have been the first hard disk into an IBM PC?

Boucher: Correct. Yes, Xebec got the first hard disk controller into an IBM PC and then Western Digital got the second one. The second major scare for us at Adaptec was we didn’t get IBM, so we had to go with some of the other startup PC clones and our single biggest customer was Eagle Computer.

Burniece: Eagle Computer. I don’t even remember them.

Boucher: Oh, Eagle <laughs> …

Burniece: How successful were they and how long were they ….?

Boucher: Eagle was incredibly successful initially. They were ramping like mad, they were the industry darling, and so we actually raised our third round on the basis of Eagle. But backing up for a minute, our first controller was a five-and-a-quarter-inch form factor SCSI controller that would bolt onto a five-and-a-quarter-inch disk, and our second controller was the same basic controller but running a run-length-limited code to give you higher density. This was at the time IBM was just getting started.

Burniece: They announced in ’81. I had just joined DEC and the announcement was made within months of the time I joined DEC in August of ’81.

Boucher: Okay, so you joined DEC at the same time I started [Adaptec]. <laughs> So initially we were trying to sell to and we got business with Cromemco and other small mini-,/ micro-computer companies that were doing word processing, so we were doing okay. Then when IBM started asking for quotes for a controller for the PC it was clear that the world was going to go that way, so we designed a controller for the IBM PC and that’s when we lost to Xebec.

Burniece: Did you get a second opportunity later to win their business or did you never get that?

Boucher: Yes, we competed with Xebec and Western Digital for the second generation and Western Digital won it, but we won the business with Eagle and a couple of other of the smaller computer clones, Eagle being the largest. As I said we really raised our major round, our third round, which was the biggest, on the basis of the Eagle business, because we needed it in order to ramp. We needed working capital, we needed large amounts of inventory, and we had to withstand large amounts of finished goods, plus the time from shipping until we got paid.

Burniece: What time frame would that have been when you got to the third round? Would that be like ’82?

Boucher: I would have guessed late ’83 probably.

Burniece: At that point your revenue was still pretty small, because you hadn’t really ramped yet, right?

Boucher: Right.

Burniece: Were you at a ten million run rate or so?

Boucher: Oh, probably. I would guess it was around ten million or so, maybe even less ….

Burniece: Do you remember how big the round was?

Boucher: Pardon me?

Burniece: Do you remember how big the round was?

Boucher: I think the round was five million, which at the time was actually a lot of money.

Burniece: Times had changed.

Boucher: Oh, yes.

Burniece: That was a pretty good-sized round.

Boucher: Yes. Our initial round was one and a half and the second round was maybe two and a half so five million was a lot of money.

Burniece: TVI and Weiss, Peck obviously went through those three rounds. Right?

Boucher: Right. The second round we brought in Merrill, Pickard, Anderson & Eyre and that was it. Then in the fourth round we brought on a few more investors, what you call the late-round investors.

Burniece: How did it go at that point? Did it start to take off? Did you start to get …?

Boucher: Well, we started really ramping and then IBM sent a letter to Eagle that basically told Eagle, “Cease and desist. You copied our BIOS.”

Burniece: Oh.

Boucher: Well, Eagle had. They sent the letter to Eagle, just as they were going public and …

Burniece: Eagle was going public?

Boucher: Eagle was going public and maybe that’s why IBM sent them the letter at that point ….

Burniece: They probably filed an S-1.

Boucher: Oh, no. They’d filed the S-1. They were going public the next day and were already celebrating that fact …

Burniece: They were literally on the …?

Boucher: The next morning it was going to be funded. I guess they had a liquid lunch and then the CEO went out and test drove a Ferrari into a lake, off the wall about a block or two from their building, and killed himself, so they pulled the IPO

Burniece: Was this before or after he got his letter from IBM?

Boucher: It was just before so …

Burniece: So he never saw the letter from IBM?

Boucher: No. They had to pull the public offering, because IBM sent them a letter that meant they had to basically stop shipping. So what happened is this was the second major crisis for [Adaptec]. We just about went out of business, because we had spent a lot of working capital, had huge amounts of inventory and accounts payable of I forget how much, but it was like we’d split the five-million-dollar round into accounts payable, working capital and inventory and we’d used it. So the five million was in play and we were going to get none of it so the company went seriously into a mode of selling all of our inventory.

Burniece: You actually started by selling all of the inventory components?

Boucher: Oh.

Burniece: [including] at the component level?

Boucher: Yes. However, 8085s at the time were like hen’s teeth. I mean we’d had to pay a small fortune for them and fortunately were able to get even a slightly larger fortune for them, because Intel was decommitting to everybody, so you couldn’t get an 8085.

Burniece: So you’re basically saying you had a little bit of luck on your side here, because you had to dump the components and they were starting to [go up in price]?

Boucher: Yes. Well, we had some good luck and we had some bad luck, because we’d paid through the nose for them and Hamilton Avnet had forced us to take all sorts of garbage that we couldn’t sell, in order to get 8085s, after they’d decommitted our orders. So the 8085 commitment that we had they would only honor if we bought all these components that you couldn’t sell, caps and resistors, etc., passives that nobody wanted. So it was not a great deal but we certainly had at least some very valuable inventory that we could sell and we were selling like mad.

We also went back to our custom semiconductor supplier. He was about to go public, so we cut a deal with him, where by taking more inventory of our chips he would give us 90 days. We had to have 90 days or we were going to default on what we owed him but he gave us 90 days in exchange for taking something that we didn’t need <laughs> , so this was how desperate we were; we were in desperate shape. In addition, we had bought an IC tester, a Sentry 7, and at the time IC testers were unavailable because the IC industry was really screaming, so we were able to rent our IC tester to AMD for a significant amount of money on a per-month basis. We did all of these things in order to just keep our head above water, but what we didn’t do, which I had been told to do by the board and they were not happy with me, was default on any payments that we did have. We had little guys that if we didn’t pay were going to be in trouble and I figured we’re going to need them in the future as long as we make it through this.

Burniece: I assume you stretched the payables out a little bit as much as you could but ….

Boucher: We negotiated everything we could. I mean that’s how we got the deal that we got for our ICs but yes, it was touch and go and …

Burniece: What turned it around? What you got you out of what looked like a .

Boucher: Well, we …

Burniece: … death rattle at that point?

Boucher: We basically got what inventory Eagle hadn’t actually put in machines back and were selling all of our inventory - we had a lot …

Burniece: … of inventory?

Boucher: Right, I mean we were going to garage shops, we were going to the Homebrew clubs and selling stuff for cost, if we had to sell it for cost, whatever. We did that until we could get back on our feet and we kept trying to sell to other customers, who hadn’t just gone into the tank. We were still selling our

five-and-a-quarter-inch cards; the problem was the PC cards that plugged into a PC. We still had revenue coming in for our five-and-a-quarter-inch cards. We cut back on everything we could without missing payroll or not making payment to anybody that had to get paid and yes, we pulled through. We got through it and got back rolling. Then we did really well until the next crisis, which was when Compaq decided that they wanted to integrate as much of the controller as they could and came up with the controller interface that they created with Conner Peripherals.

Burniece: The IDE.

Boucher: Right, and so IDE of course killed the standard disk controller and it caused us to have to come up with what we called the host bus adapter. We were lucky in that at the same time this little company in Salt Lake City by the name of Novell had given up on trying to build a microcomputer and had decided to build file servers. What they really needed was something that would allow them to put as many disks as they could into a single file server, a single processor, because the regular PC with an IDE interface could really handle no more than four disks and usually just two. You had these slave disks you could put in but in general people didn’t want to do that, so you really only had two disk drives and Novell wanted to be able to cram lots of disks into a box and create a file server. So between ourselves and Novell we made the HBA an important commodity and thank God for Novell, because our IBM controller business just started auguring into the tank.

Burniece: You ended up missing the IBM contract, getting Eagle, who was hot and then the CEO died in a car crash and they got sued by IBM. You’re now in trouble and how did you find Novell? What was it that got you to Novell and what got you to do an HBA?

Boucher: Novell found us. We built the …

Burniece: They found you?

Boucher: Yes. We built the HBA because it was the only way to connect the controllers that we were still selling. We were still selling these five-and-a-quarter-inch controllers that bolted onto the side of a five-and-a-quarter-inch drive and what you could do--

Burniece: They were still ST-506 at that point. Right?

Boucher: Actually, I think by then they were all 512s but yes—

Burniece: Next generation--

Boucher: Right. And so what we needed to do was build a controller that you could plug seven of those in ….

Burniece: You actually developed that on your own; Novell needed it and found you?.

Boucher: Yes, we sold it to anybody that wanted it and certainly ultimately the PC companies bought them, as well, but initially Novell was our big customer.

Burniece: When did that happen?

Boucher: That happened - actually I don’t think we developed Novell until right about the time I left. We had built the HBA and were starting to try to sell the HBA, along with continuing to sell our SCSI controllers. The other thing that we were doing, of course, was trying to sell our chips to the drive guys. so we were trying to convince Seagate but not Western Digital, because they wanted to do their own, but anybody else in the drive business--

Burniece: The Maxtors and Quantums of the world?

Boucher: To buy our controller to make intelligent drives. Initially, the disk drive companies didn’t think very much of the idea of intelligent drives but ultimately they started to think hey, that makes a lot of sense, so we started to sell disk controllers to the drive guys. That plus our five-and-a-quarter-inch boards were our two main businesses for quite a while, until the HBA business really took off.

Burniece: Did you actually end up getting a disk drive company to use your chips in the intelligent drive?

Boucher: Oh, yes. We …

Burniece: What are some of the ones that did that?

Boucher: Seagate did it; that was one of our good customers. I can’t remember who else.

Burniece: Several did it.

Boucher: Oh, yes. Mainly the smaller guys who were startups and then ramped from there. Some of them were successful and some of them that weren’t ultimately, but yes, I would guess that Seagate was our most serious customer.

Burniece: I want to segue back. We’re now up to somewhere around ’87--

Boucher: Eighty-six. I left in eighty-six.

Burniece: You left Adaptec. You didn’t leave the company’s board though; you left as the CEO, right?

Boucher: Correct.

Burniece: … to found Auspex.

Boucher: Right.

Burniece: I want to go back and fill in the blanks between how the SASI interface that you ended up developing at Shugart eventually became the SCSI standard and what your role was in that. Did you participate in making that happen? Were you on the standards committee?

Boucher: <laughs> Yes, that’s an interesting story. Basically, I’d already started Adaptec and NCR had fallen in love with SCSI and was a Shugart customer. Shugart finished the design that I was doing at Shugart.

Burniece: You mean they’d fallen in love with SASI. Right?

Boucher: Correct.

Burniece: NCR liked SASI.

Boucher: Right, and they were a Shugart Associates customer, so they decided that they wanted to make SASI a standard. Interestingly enough, that was not on either Shugart’s horizon or ours at Adaptec but ….

Burniece: That by the way is an old story in the standards world.

Boucher: It is.

Burniece: CDC didn’t want SMD to be a standard either. After it became a standard, they thought it was great but they originally didn’t like that idea.

Boucher: Yes. You can sort of imagine that that would have been the case because George didn’t even want to sell them separate from the drive. Okay. So …

Burniece: Did NCR take it to the standards body and then knock on your door?

Boucher: Sort of. I mean basically John Lohmeyer at NCR had been a part of ANSI ….

Burniece: He was a key guy at ANSI at that point. Right?

Boucher: He had been a part of ANSI for quite a while. Because he was both part of ANSI and he was NCR he sort of looked at himself as the guy that takes things that we’re doing at NCR and gets them to become standards because NCR was, unlike us, seriously about wanting standards. So John went to both Shugart and us to conscript us into this and so we …

Burniece: When did that happen? Before you left Shugart?

Boucher: No, no, no. That must have been sometime-- it must have been in ’87- late ’87 or ’88.

Burniece: So it’s well within that timeframe.

Boucher: See if I can keep my ….

Burniece: Well, actually …

Boucher: … right now. <laughs>

Burniece: Those things don’t click with me because by ’87 SCSI was already …

Boucher: Oh, wait a minute.

Burniece: You’ve got to be talking about in the early days of Adaptec.

Boucher: Yes, it was the early days. Oh, pardon me - it was late ’81, early ’82.

Burniece: I’m going to pull some pieces together to try to keep these dates [correct].

Boucher: Yes …

Burniece: I left CDC in ’81 and I had started the ISI project [in ~1978], which was on the table [at the ANSI standards committee in 1981] to replace IPI or compete with IPI …

Boucher: Wait a minute.

Burniece: … and SCSI came in the same time frame.

Boucher: Wait a minute. I thought ISI was CDC.

Burniece: Yes, I [started] that at CDC [mistakenly said DEC].

Boucher: Okay, you did that ….

Burniece: And by the time I went to DEC in ’81, [CDC] had already told the standards committee, “We’re going to give it to you.”

Boucher: Oh, because Gene [Milligan] was pushing ISI.

Burniece: [Gene was the standards rep from CDC / MPI in Oklahoma City, Oklahoma] but the bottom line is that in the time frame from ’78 to ’81 we had developed ISI [internally at CDC / MPI in Normandale (Edina. Minnesota)]. IPI was being developed by the standards committee and SCSI was on the horizon, [so Gene threw ISI into the ring] …

Boucher: Now basically …

Burniece: … in that time frame.

Boucher: The only reason you’re saying IPI was being developed by the committee was because Dal Allan is a personal force - I mean “personality force” - and Dal Allan really wanted IPI.

Burniece: That’s true and he’ll tell you that. In fact, it’s in his [oral history] interview[1] .

Boucher: So Dal basically went like this with Gene. I mean they were friends but basically the way this worked is Dal was running the committee, which sort of gave him an unfair advantage over Gene. <laughs>, who was such a gentleman but anyways …

Burniece: Gene was a wonderful guy.

Boucher: Dal’s a bulldozer and he was running that committee. John [Lohmeyer] was having a problem, because he figured how the hell are we going to get in there. I mean he said, “Politically, we’re not going to be able to do this. I’ve got the committee because this other XT committee (X3T committee) doesn’t have anything to do and is looking for something to do but Dal will shoot us down immediately because he wants this IPI thing.” And I said, “Well, look. We just need to differentiate it. He wants to emulate the FIPS60 channel in one mode of IPI.” That’s a different world. So we’ll just say, “No, no, no, no, Dal. We don’t have anything to do with you guys. We just want to be the low end”-- now this is, pardon my French, “bullshit” but we needed some political way into ANSI, so by saying, “No, we’re just the very low end then he’s not going to be able to argue” and …

Burniece: That clicks. SCSI was definitely positioned as the low end and IPI as the high end.

Boucher: The ironic thing is …

Burniece: … it turns out you ended up being the high end.

Boucher: Oh, no. I mean there’s no question about that - I knew it was going to be the high end. We were doing what we wanted to do back at IBM in terms of interface, but there was no way into ANSI if we said that. Plus who wants to compete with Dal and Gene, so by differentiating ourselves sufficiently, we could get in and we’d be okay, so …

Burniece: That was all happening in the late ‘70s to early ‘80s – in that time frame, not the late ‘80s--

Boucher: [I agree] This was early ‘80s.

1 – Oral History of Dal Allan, March 12, 2013, CHM Reference number: X6787.2013

Burniece: About the time you were forming Adaptec this was all going on hot and heavy.

Boucher: Oh, yes. I think the first time that I spoke with John, I’m not even sure if we were funded but it was close. We may have been funded.

Burniece: Time-wise that makes sense. We’re talking about right in the ’81 time frame.

Boucher: Yes. It was that timeframe. It might have even been late ’81, early ’82. Anyway, that allowed us to go through our own committee at ANSI **…**

Burniece: You are now an active participant making that happen?

Boucher: Yes. John had asked that I and Daniel Loski from Shugart. He's a French guy. I'll think of his name.

Burniece: Okay.

Boucher: John had wanted both Shugart and me to attend, so he invited me onto the committee and I agreed to join it.

Burniece: Why?

Boucher: And so the first meeting-- the first thing-- the first order of business was to figure out what we were going to call the standard. We were told it can't be SASI, which was really a shame because Shugart Associates System interface, I mean SASI, was such a cute word [pronounced “Sassy”]

Burniece: Plus Shugart [Associates had sued Al Shugart over re-using the name “Shugart” in Shugart Technology, forcing Al to rename his new company Seagate Technology].

Boucher: Exactly from Steele. So, basically I felt that we should keep it as close to SASI as we could and there was a lot of discussion. There were all different opinions on the committee. But we came up ultimately with SCSI [“Scuzzy”] because “Small Computer Storage Interface” would help us continue to differentiate from those guys that were trying to build this monster [IPI]

Burniece: That was a political piece there.

Boucher: Exactly. So I figured well okay, we can call it “sexy”. I was hoping that it would go from SASI to sexy but when the committee meeting broke at the same time as Dal’s committee, Dal walks in and says, "Okay. So what's happening?" I said, "Well, we've got a name. It’s SCSI." And he said, "Oh scuzzy." I said, "No. No. No. [it is “sexy”]

Burniece: So he's the one that pronounced SCSI scuzzy rather than sexy?

Boucher: Yes. I think he claims he's not, but he is. He's the one …

Burniece: Wow.

Boucher: … and because he wanted in the worst way to kill us. So he laughed when he said, "Oh, so you guys are scuzzy." But, yes, that was ….

Burniece: So how long did it take before SCSI became a standard? Did that take years?

Boucher: Well, it was already a standard.

Burniece: I mean by the time it actually became an official documented standard.

Boucher: Yes, what we were doing at ANSI was rubber stamping it. Now, in the process, however, there were a lot of concessions that were made. So it wasn't a rubber stamp of what was done originally and before SASI became SCSI it wound up with two major changes. One, NCR had already decided even before that we had to increase the addressing. I felt that the way the original architecture was designed there was enough to address a disk today. It was set in terms of number of blocks and block size. Well, the block size was simple. You could increase the block size so large that the number of blocks just didn't matter. But NCR said, "No, we need to be able to keep the block size." And they were right. I hadn't

understood, because at the time magnetic disks were the major application, and hadn't really understood optical disks.]

Optical disks were coming on like a freight train and they still wanted very small block sizes but, because of the capacity of an optical disk, the number of blocks, therefore, was huge. In the original SASI architecture I hadn't allocated enough blocks. NCR wanted to increase the block size drastically, so they did allocate another four bytes to the block size. That plus the block number gave us a huge capacity. So that change NCR wanted even before ANSI and that's one difference between the original SASI and SCSI.

The second one was Bill Roberts at Emulex felt that we weren't going fast enough. I didn't disagree with that, but he wanted something faster on day one, because Emulex up until then was really a DEC emulation shop. Roberts had decided that he wanted to be on this committee, because he was going to get into SCSI. In fact, I think he may have already built a SASI controller. For sure he was going to build a SCSI controller. In fact, he was in the process right then. As a result, he wanted to go faster than the spec could presently go. So he put in what we called offset interlock, which was a synchronous data transfer.

So those are the two major changes that took place between SASI and SCSI. One other thing that was really interesting was there was a field in the original SASI architecture that as the time had not been used. I put in because the [IBM] 2314 had both a physical and a logical capability. You could take a physical disk drive, and change its logical address by taking the plug out and popping another plug in. That was felt to be important. It was no longer important, once we started believing in non-removable disks, but Disk 0 was a very important disk for the OS because that's where it booted and if you wanted to be able to boot from any physical location, those packs were movable. So IBM had this logical versus physical concept.

By the time we were defining the SCSI architecture, fixed disks were the only thing that people were thinking of. People didn't believe in removable disks packs any longer - at least in our market. But I had not really figured out if we were going to want that in the future. So SASI had this LUN [Logical Unit Number] that for the longest time was just a spare field that finally got used. It was one of the things that ANSI was going to take out but we didn't take it out, which turned out to be very fortunate, because in the meantime the logical unit number has become a very useful field. Those were about the only real changes that did and did not get made. Oh, there were a couple of other changes that got made but I don't remember exactly what were. They were for single commands because DTC, who was also on the committee, couldn't do them with their TTL controller and wanted the changes made.

Burniece: I don't remember the details but there was also a common command set that got created.

Boucher: That was …

Burniece: But I think it even bridged IDE as well as SCSI, right?

Boucher: Yes. The common command set happened later, because Shugart and OMTI, who didn't even exist when we started the ANSI committee, wanted to simplify things. They didn't want to have to build a

controller that did some of the search capabilities that I’d put in. I think those have been completely wiped out. I don't think anyone has ever [used them]. In fact, I haven't looked at a recent spec for SCSI, but it may be that the search functions are no longer there. I originally wanted them because they could be useful, when you look at count key data the way IBM did - you did real time searches of key fields in order to be able to enable some set of functions, like database, to run faster. That seemed to me to be a useful function going into the future that could be incorporated, especially if you could search whole data fields and do it in real time, rather than having to read everything in the memory - waste memory, the channel space and memory bandwidth. Do a search with the CPU and then, once you got a hit, do something. So I felt that those search commands would be useful. I think most people felt that the search commands were too complex, because what you were doing was two-way, simultaneous data transfers. You transfer the search argument that you want to the controller. The controller then stores it, does a search until it finds the data and comes back with search complete status.

Burniece: I would be surprised if there isn't something that does that for metadata, whether it uses your command set or not, but don't know. That's a good question …

Boucher: Well, the …

Burniece: … because search is definitely important right now.

Boucher: The key is was that in the original …

Burniece: SASI.

Boucher: It never was removed from the original SCSI spec but OMTI didn't want to have to incorporate it, because they felt it was too complex. So they pushed and this was now our guy - I'm having a senior moment forgetting his name, who had been at Shugart and then moved to OMTI that drove the common command set. It deleted all of the search commands and may have also deleted some of the special tape commands that we had. I had originally assumed SASI was going to be an interface for tape and printers, as well as disks, so assumed that we'd have both printers and tapes on SCSI. We had some special commands for both of them. I think those basically stayed in SCSI, as well.

Burniece: I think they may have. I'm going to just ask one more question on this, and maybe we'll go a little further. SCSI actually ended up having two facets to it. One was the interface itself that was used on so-called SCSI Drives that went from parallel to serial and it's still out there. It never became the majority of the market but it's always been a significant minority of the market. The other is the storage protocol that became universal almost across all platforms, [including Fibre Channel and IDE].

Boucher: Well, Fibre Channel is …

Burniece: How did that division get made between the physical interface and the protocol that's so pervasive?

Boucher: Well, that's why, when I originally designed the spec, I designed it in three pieces.

Burniece: So that was the idea from the beginning that you …

Boucher: Oh, yes. I knew we were going to go to serial and so it had to be done in a fashion where the logical half was the important half because that's the half that allows the system to, irrespective of anything in hardware, continue to talk storage - well, not just storage because printers, as well. I considered storage being tapes, disks and printers to be the key. The system needed to be able to talk to those devices, without any knowledge of the physical interface. So whether it was serial or parallel shouldn't make any difference and certainly how many tracks or cylinders should not be important, other than to be able to find out what the capacity was. So the logical interface had to be one where you could format a drive and once it was formatted be able to use it no matter what. So, yes, it was literally designed that way. When IDE was designed they basically took the logical interface of SCSI.

Burniece: Yes, they did.

Boucher: And Fibre Channel did the same thing. They didn't reinvent their logical interface they just did the physical part.

Burniece: So fundamentally, the logical part (protocol) is used by everybody.

Boucher: Right.

Burniece: And that was part of your original concept.

Boucher: That was …

Burniece: Did you patent that at the time, or think about patenting that at the time?

Boucher: No. I didn't really understand patents and Shugart Associates really didn't either. So I didn't understand them at Adaptec. We didn't really patent much of anything until later on at Adaptec. I couldn't have patented any of this anyway, because I developed it at Shugart, so I was happy Shugart didn't either, as it turns out.

Burniece: Plus the good news is it became pervasive.

Boucher: Yes. Yes, it did.

Burniece: One of the longest term [interface specifications] ever.

Boucher: So far, so good.

Burniece: It's an awful long time. So let's talk a little bit more about Adaptec, before we move onto Auspex. You started Adaptec in '81, and you were the initial CEO.

Boucher: Right.

Burniece: At some point you brought a CEO in and then you left in '86. So what was that transition and what did you do during that timeframe? Or was that done after you left?

Boucher: Well, after we stepped heavily in it with Eagle, the board said, "Hey, you need to get some more experience on board." And I thought - you're absolutely right. They didn't tell me this, until after we had gotten through the problem. When we were in the middle of the problem we couldn't have attracted anybody. So once we had gotten through the problem and were ramping again, I could start to attract people that the board would be interested in considering. So I got a search firm involved and I started bringing people in and having the board meet them. I must have brought at least three or four people to the board and each time they turned me down. So it became clear to me that, in the meantime, we were doing well enough that their criteria was going up and up and up because the guys I was bringing in were much more experienced than me in terms of general management.

This is my first general management experience, so I would be the first to agree I'm not the most experienced general manager in the world. I was learning on the job without any question. So finally before we went public but were doing extremely well, it became clear to me that they were just going to keep turning me down, because I was never going to find somebody they'd be satisfied with. So I waited until the search firm had found two guys that looked really good to me. Then I added a guy that I'd worked for at IBM, who then left IBM to go sailing for a few years and I had kept his dog while he was sailing. When he came back, he joined Amdahl and we took back up our relationship, plus I gave him back his dog. So I brought him in as the third [candidate], took the three guys to the board and said, "Okay, here's three people - choose one but the answer is not none of the above. So choose one." Well, they knew that I would have been partial to John Adler, so they said after they had met all three, "Well, we like all three of them. We know that you would prefer John. So let's bring John on." So we brought John on as the president and chief operating officer.

Burniece: You had mentioned his name early on in this interview and I didn't think to ask if that's the same John Adler.

Boucher: Oh, yes. I worked for John …

Burniece: Way back at IBM in the early 80's.

Boucher: Correct and we developed a close personal relationship. So it continued after I left.

Burniece: I also got to know John. He's a really good guy.

Boucher: He is a good guy.

Burniece: So he then eventually became the CEO, when you left. Is that how it worked?

Boucher: Right. So …

Burniece: So you brought him in as COO, under you.

Boucher: Well, basically he came on as COO and I immediately started working on taking the company public.

Burniece: So the company went public after he joined?

Boucher: Right. We went public and now the company was doing very well, plus we had what was now a stronger management team. So the street could see that just in case we've got a good backup.

Burniece: So when was it you went public?

Boucher: We went public I think June of '86 - June or July of '86.

Burniece: That must have been a really heavy time.

Boucher: Yes. It was a lot of work. We had a challenging time, but yes, it was very rewarding. It was nice to finally get the recognition of going public and we were running at around a $60 million run rate, which was …

Burniece: You about 60 at that point. Okay.

Boucher: So I stayed for probably another year but, after another six months or so, it became clear to me that John was doing the job of the CEO. So I sat him down and said, hey look, I will keep working for you on finance and sales but will give you the CEO because you're doing it.

Burniece: So did you become the executive chair at that point?

Boucher: Yes. I stayed and was chairman after that but I think John started feeling like people were still looking at me as the CEO. I wasn't but I think that because I’d been the CEO for the entire time, it was a little bit uncomfortable for John. I figured well, okay, I'll find something else to do. So I started thinking about what we could do. Initially I thought about doing Auspex, [under] the auspices of Adaptec, and just funding it as a [spin-out]. sort of what Cisco started doing later. But John looked at it when I gave him the proposal and said, "That's a computer company. I am scared silly of doing something like that." So I said, "Okay" and left. And for a while …

Burniece: Was that now '87 when you left the joint to found Auspex?

Boucher: Yes. Well, that was either late '86 or early '87. So I left and initially became what Merrill Pickard claimed was the original Entrepreneur in Residence. So they gave me an office.

Burniece: So you became an EIT for a while?

Boucher: Yes. And I …

Burniece: Or EIR (Entrepreneur in Residence).

Boucher: Right. My job was to write a business plan. They gave me a junior guy they just brought on to help with the business plan and they also occasionally asked me to review something they were looking at to give my opinion. But my major job was to write a business plan. Andy Radcliffe was the junior guy that they assigned to help me with the plan. So Andy and I finished the plan and Merrill Pickard, of course, was the initial investor, because they liked the plan …

Burniece: Was Andy then a co-founder?

Boucher: No. No. Andy wasn't a partner. The partner of Merrill Pickard that then joined the board of Auspex was Jim Anderson, who was basically the guy who had brought me in. Jim had invested in the second round of Adaptec.

Burniece: So he knew you from Adaptec?

Boucher: Right. So Jim was the primary first round investor for Auspex.

Burniece: So Adaptec was essentially a realization of the plan that you first proposed to Shugart [Associates] to create this interface that would be a universal high-level interface and became SCSI. Tell us about the basic concept that you founded Auspex on. What were you trying to do, and what was unique about that?

Boucher: Well, what I wanted to do was build what has become known since as Network Attached Storage (NAS). What a file server was at the time was a computer that was just repurposed as a file server, which is horribly inefficient, because a general purpose computer is designed to be a general purpose computer. It’s not designed to be a file server. This wasn't exactly a new concept, because Novel had built what's the equivalent of a NAS device. But they had done it for the microcomputer world and what I wanted to do was build a NAS product for the minicomputer world.

Burniece: So it's not fair then to say that you were the very first NAS in terms of the general concept. But you were the first one to actually build a company around the principles of NAS and network file systems?

Boucher: Actually, I would give that credit to Novel.

Burniece: Would you? Okay, [why]?.

Boucher: Because Novel actually sold a pure [file server] - in some sense that's not fair because it wasn't a design-from-scratch file server, since they were using a general purpose PC mother board. So in that sense, Auspex was the very first NAS device, because it was the first device that wasn't in any fashion a general purpose processor. The Novel machine [used] a general purpose processer, but it was sold [only] as a NAS device. If you look at what Sun was doing - Sun was also selling NAS devices [based on general purpose processors], so maybe I'm being a little bit too generous to Novel, because Novel was really the equivalent of Sun. But both Sun and Novel gave me the idea, because they were selling products that were very inefficient NAS devices. They'd give you a general purpose computer and tell you that it was a file server and it worked like that. But it wasn't great.

So what I wanted to do was build a competitive NAS product [that was really efficient]. But building a competitive product took competing with Sun Microsystems, who was already a monster. It seemed to me that Scott [McNealy]'s philosophy was to always have somebody to kill. First he was going to kill Apollo. Then he was going to kill DEC. So the mantra inside Sun was always killing someone and ultimately they had to kill Auspex. They had a thing called the Auspex kill pool to kill us.

Burniece: They actually had something they called the “Auspex kill pool”?

Boucher: Yes. Clint Eastwood had just come out with a movie, so [Sun] marketing named their project the Auspex kill pool.

Burniece: Don't forget we're talking about Scott McNealy here …

Boucher: Yes. Right.

Burniece: … CEO of Sun at the time. So did you actually then develop NFS or had that already been developed?

Boucher: No. No. That was the interface that …

Burniece: Sun had, right?

Boucher: Yes Bob Lyon developed it at Sun and we made use of it. I didn't want to try and boil the ocean so anything that we could use we did. We actually didn't want to design the processor either, so we bought a Spark processor that they sold as just a card.

Burniece: They already had the processor, okay.

Boucher: Right. But we didn't invent the backplane even though we totally misused it. In other words, we went with a VME backplane and then ultimately we went with Motorola for processors for everything else. We then went to Intel but not during my timeframe. Basically we took a VME backplane and modified it so that it was not a memory bus, but it was instead a communication bus. The idea was that we would build a very light-weight message based operating system that would run across all the processors on the device and the device would consist of storage processors, network processors and file processors. That's what the Spark card did. The initial Spark card was our file processor, because the file software really wanted to run on what would otherwise be a general purpose processor. But the Spark card was not a general purpose processor. It was a just a processor card for a big machine. So we ultimately put up to four network cards, a storage card and a file card in the system. That was our basic machine initially.

The way I knew we could compete was because at the time everybody in the mini computer business or the work station business, which was taking over the mini computer business, thought that they had to use high performance drives and the most prevalent one at the time was the Fujitsu Eagle. So by going with five and a quarter inch drives, which were considered way beneath the minicomputer and work station world, we could create much higher performance at a much, much lower price point. So the idea

was to stripe five and quarter inch drives. At the time RAID had not come out so we were just doing a parallel stripe, which got renamed RAID 3 later. But RAID 3 was around long before RAID [was formally named]. [The original] RAID, as I'm sure you're aware, [was later] designed as RAID 5 and they named everything from a mirror to striped drives with a parity disk that had been around for a long time as earlier generations of RAID.

Burniece: You’re talking about the Berkeley RAID paper now [“*A Case for Redundant Arrays of Inexpensive Disks”* by David Patterson, Garth Gibson and Randy Katz],

Boucher: Right.

Burniece: -- That came out in 1988

Boucher: Yes. Just when we started Auspex.

Burniece: It basically put together a taxonomy for what people were already doing [calling it RAID 1 to RAID 5].

Boucher: Except for RAID 5, which was Garth's - what he did for his thesis …

Burniece: One of the things they did miss was they didn't name striping, so we did that at DEC. We called it RAID 0 and it stuck. We were already doing RAID 1 mirroring, which we called “shadowing”, together with RAID 0 striping, on the VAXcluster, so that then became known as RAID 10.

Boucher: Yes.

Burniece: We wanted to be able to describe what we're doing, within Berkeley’s taxonomy. It was all done in software [on the VAXcluster HSC controller, with shadow sets going N deep].

Boucher: Yes. I did not realize that zero was DEC’s.

Burniece: I'm actually the one, who said let's name it RAID 0.

Boucher: I love it.

Burniece: It goes back to before I left DEC in late 80's.

Boucher: Yes. Basically …

Burniece: We did it to basically fill in the blank in their taxonomy

Boucher: You're absolutely right. Garth's paper was RAID one through five, and I’d forgotten that.

Burniece: Paul Massiglia, who worked for me at the time at DEC, institutionalized that for the RAID Standards group, when he put the first RAID Handbook together for them. He put RAID 0 in there to be able to use that name. So you were doing RAID 3?

Boucher: Yes. That was our original.

Burniece: It is an interesting story because essentially what the Berkley guys did was put a taxonomy together that pulled it all together, after the fact.

Boucher: Yes.

Burniece: They didn't actually invent RAID. They basically institutionalized the concept.

Boucher: Well, they invented RAID 5.

Burniece: They probably did do that.

Boucher: They did invent RAID 5. Up until Garth's paper, what became known as RAID 3, which was parallel striped drives, was known and that's what we started to do. When we formed the company, I hadn't seen the paper. I don't know if it had even been published, when we formed the company. So our original business plan had RAID 3 with five and quarter inch drives, so that we could outperform an Eagle, at a fraction of the price, and we could actually get higher reliability, because by striping you could lose a drive. So we could sell both reliability and performance. Then RAID 5 came out and, of course, we needed to incorporate that but initially we just sold RAID 3 because that was good enough. When the product initially shipped we didn't ship it with RAID 5. That was around a good six months later, I would guess.

Burniece: So if you look back on that and pick out the two or three things that were really unique that Auspex put in place, and eventually became part of what we now call NAS, what would they be? What were the key things that you added to the formula that nobody else had done before institutionalized?

Boucher: I think the biggest thing and I don't think it's even fully institutionalized now, was the fact that we recognized that the computer backplane was history, as it was known to date. In other words, up until that point the memory bus was considered the computer backplane. If you looked at the PC, the channel was integrated on the IC with the motherboard. The motherboard had the channel on it but the architecture was still basically the memory bus and then channels to peripherals. So you had both - you still had the fast channel and the slow channel just like we had at IBM - selector, block multiplex and byte multiplex channels. So that basic architecture hadn't changed.

What we recognized was that the new computer backplane was really this thing that had already been invented called the switch. So the switch backplane was really where the computer backplane was going and that's what the backplane of the Auspex was - a switch. We used a very light-weight message based operating system to manage the switch. We didn't need a special card, which all the original switches had to manage the switch. The switch was managed by anything that plugged into it, if you will. So the switch ran across the backplane. It was distributed. But ultimately that's what the computer is going to be, because the processor is now an IC. It's not a huge card any longer, so if you want to configure a computer or a file server or whatever device you want, you just need to configure different functions on a switch backplane and you've got a very powerful device. For instance, a super computer is just a whole group of processors that you want to group on as fast an interface as you can, because it's that

communication protocol that allows the super computer to run at whatever speed it's going to run. The fastest ones are going to be the ones that are on a very closely managed switch. As you have to move the switch further out it's going to get slower. If you want a super computer that is comprised of say 2000 processors, it would be grouped with very high speed local backplane switches linked to others of the same sort. So that I think was the single most important thing that we did.

Burniece: Did you patent that?

Boucher: Well we patented some of it but we did a horrible job. By then I'd learned that patents were important but at Adaptec it was too late to do anything about it. So we did write patents at Auspex, but we wrote ones there were poor enough that I don't think they really worked all that well. In some sense it didn't matter, because I don't think people started really seriously understanding it for years. In fact, I had a meeting with Cisco probably in the middle 90's where I described this and said, "You guys are the next computer company" and they understood extremely well. At the time they said, "Yes, but our best customer is IBM. As you're probably aware, they're building it now." So my guess is that ultimately you won't see a computer that's done any differently because it doesn't make any sense. I would guess that that was the single most important thing we did at Auspex, even though I don't think it's recognized. The thing that Auspex gets remembered for is the real beginning of the NAS business, because it was the first place that it happened in the larger market of minicomputers and work stations.

Burniece: So how successful was Auspex? And how far did you get with it?

Boucher: Auspex was quite successful. We took it public and …

Burniece: And when was that when it went public?

Boucher: In ’93.

Burniece: Okay.

Boucher: And yes, we continued to ramp.

Burniece: So about what kind of revenue stream were you at when you did that? Was that also around a hundred million dollars?

Boucher: Yes, I would think that we were running at around a run rate of a hundred million. I think our last quarter had been around twenty-five million.

Burniece: Okay.

Boucher: Last year was maybe sixty million. I think that’s about where we were. But don’t hold me to that.

Burniece: So you’ve now taken two companies public.

Boucher: Yes, and I stayed there for a little longer but again we ran into trouble. The board wanted me to bring a COO on board and I did that. I brought the COO onboard and about the time I did, we were doing extremely well again, but this time I didn’t do as well with the guy that I brought in. He really wanted my job badly, so he started finding ways to get rid of people that had anything to do with me. I didn’t really feel good about that so, rather than allow that to continue, I handed off the CEO and left.

Burniece: So when did you leave?

Boucher: I left in ’96.

Burniece: Okay. And was that to found Alacritech at that point?

Boucher: Well …

Burniece: Or did you just leave and decide what to do later?

Boucher: Yes. At that point I didn’t know what I was going to do, but I knew I just needed to leave, because if I didn’t he was going to tear the company apart …

Burniece: Okay.

Boucher: … getting rid of key people. It was a bad political situation. It’s really too bad, but once we brought him on, the Board felt we couldn’t just let him go. I went to the Board and said, “Look, this is what’s happening” and they said, “Look, he’s making these decisions. We need to let him have time to make them.” I said, “Okay, well, I need to give him the CEO then, because otherwise, he’s going to keep doing things that I think are going to hurt the company,” and they felt that was okay. So I did that and then I stayed on as executive chairman working sales and finance, but the company started going a way that I felt was really bad and I discussed that with the Board. They didn’t feel comfortable with the fact that things were bad enough that we could make a change. So I said, “Okay, well, I need to leave then” and I left.

Burniece: Now by that time, Larry, had Network Appliance …

Boucher: Oh, yes, Network Appliance …

Burniece: … started to take off and was it impacting Auspex or does it do that later?

Boucher: They hadn’t impacted Auspex at all, since they were in a completely different business. They were in the very low-end and Auspex was in the high end. But Auspex was unduly worried about them and this is one of the problems: they were managing in the rear-view mirror.

Burniece: Yes.

Boucher: They were so busy watching what NetApp was doing and the fact that NetApp was being successful in their market that they quit paying attention to their own market. They went after NetApp and, as a result lost their market and ultimately put themselves out of business.

Burniece: How would you differentiate the difference between the two companies? What was it that NetApp did different that Auspex did?

Boucher: Well NetApp was going after the very low-end of the market. They built a machine that was a fraction of the cost, but it was a fraction of the performance.

Burniece: Okay.

Boucher: As long as Auspex kept its performance and its capacity up, it would’ve had a good market. But there is no question that ultimately Auspex was going to have to be able to come down market. You just can’t forget your market. They needed to continue focusing on their market while they came down, maybe even buy NetApp. But instead they lost their focus, went after NetApp and ….

Burniece: Well, they basically lost.

Boucher: Yes, oh, lost big time.

Burniece: Okay, any other comments or things you want to talk about on Auspex or should we move on to Alacritech?

Boucher: I think that Auspex has been pretty well covered.

Burniece: Okay. All right. So how did Alacritech get started? What was the concept there and when and how did you put that together, do another business plan and create another company?

Boucher: At Auspex it was clear that the major bottleneck in what we were doing was the network card.

Burniece: Right.

Boucher: Each one of our cards cost around ten thousand dollars to build. These were 9U cards. The storage card had ten SCSI interfaces on it and each one was running five megabytes a second. Our initial network card had two 100Mb/s interfaces on it. So for ten grand, you could have ten 5MB/s interfaces or you could have two 100Mb/s interfaces. There’s a small difference there.

Burniece: Right.

Boucher: The reason was because the protocol for the network controller was still all software.

Burniece: Okay.

Boucher: But if you look at the network, the network protocol isn’t that much different from the disk protocol. It’s more complex and it has more content, because it has a bigger job to do.

Burniece: You’re talking particularly about TCP/IP at this point, right? The network protocol.

Boucher: Yes, but TCP/IP is a serial protocol. It’s more complex than the disk serial protocol, but nevertheless it’s a serial protocol.

Burniece: The big difference is a disk protocol is a dedicated channel that has to succeed every time and that is designed in. A network protocol is a best try and that will try again if it doesn’t succeed ….

Boucher: Well, not …

Burniece: Kind of a best effort versus a guarantee …

Boucher: Not at the TCP level. At the TCP level, it’s guaranteed.

Burniece: Okay.

Boucher: TCP guarantees you got your data. In UDP, you’re absolutely right. It’s best efforts. In fact, when we first started Auspex, UDP was the way that you talked.

Burniece: Okay.

Boucher: NFS had a built-in recovery re-try process, because TCP was so slow that it just was inadequate, and as long as you were on the machine room floor you didn’t have that many problems. So UDP would be good enough ninety percent of the time, and when it wasn’t, NFS could recover.

Burniece: Okay.

Boucher: But if you had used TCP, things would’ve gone so slowly it would’ve been a disaster, because TCP was designed to send the data from here to Podunk.

Burniece: Right.

Boucher: Whereas UDP was just not guaranteed, but would usually get across the machine room floor, it pretty much guaranteed to not get across the country. So if you really wanted to go outside the machine room you want TCP. But inside the machine room you wanted UDP.

Burniece: All right.

Boucher: That’s basically history today, in part because of what we did at Alacritech, which is not all that well understood. But, basically, what I recognized was that we needed to put TCP/IP in hardware and we did that just as we had done at Adaptec with disk protocol processing - remember, before Adaptec, the way you did disk protocol processing was with a processor.

Burniece: Right.

Boucher: A very expensive processor and very slow. I mean slow in terms of how fast it could read a track. It took a fast, expensive, power hungry processor to slowly read a track. Okay, so what we needed to do was the same thing for network protocol processing that Adaptec had done for disk protocol

processing. That was the concept. So I put together the team and made a couple of major mistakes. The first mistake was deciding that we would do another Adaptec. That was a major mistake, because I was very slow to learn this - about ten years slow.

<laughter>

Boucher: Because we kept being sort of successful. We sold enough product and we could raise money, so we kept going. But what we did is we built a controller card that looked like an Adaptec controller card plugged into a PC that put TCP/IP in hardware. What it would allow you to do was build a file server that was much faster. But the problem is that it looked to most of the world like it was just another network card and that world was owned by Intel and Broadcom. Well, those guys make Apollo and Sun look like saints. So they were more than able to keep us out of the market and our market was just the three sigma guys. I mean, the government, in terms of guys that just didn’t care because they needed speed no matter what. The second thing was for the last two years that I was at Auspex every single CIO I sat in front of said, “Hey, if you don’t get off of NFS and onto CIFS you’re history, because we’re moving from UNIX to Microsoft.” Well, I had been told that enough that I decided, “Okay, that hasn’t happened yet, but it’s going to happen. So when we do this card we should probably do it for Microsoft.”

<overlapping conversation>

Burniece: … only for CIFS or did you do it for both?

Boucher: It was a huge job to do it for one, so …

Burniece: Okay.

Boucher: So …

Burniece: So you went right for CIFS.

Boucher: Right, and …

Burniece: Just to calibrate on the tape here CIFS is the Microsoft version of NFS.

Boucher: Correct. Thank you. So that was the second significant strategic blunder that I made. We got the thing done and we took it to Microsoft to demonstrate, because I felt Microsoft was the perfect partner. We showed it to them and they had a thing they called TTCP. It’s a test program. TTCP may be “Test TCP” - I don’t really remember what the acronym stood for. They ran it against one of the cards that they were using - I forget whether it was Intel or Broadcom - but they ran it against, say, an Intel card and they looked at what they could get when the network saturated and what the bandwidth of their processor was. So they saturated the network with …

Burniece: We’ve lost the mic again.

Boucher: Oh, I’m sorry. This time I didn’t destroy it however.

Boucher: Oh, maybe I did.

<off-topic conversation>

Boucher: So we went to Microsoft and they ran this against their Intel card and I forget what their processor was. It was close to saturated when they saturated the network.

Burniece: Hmm ...

Boucher: Originally, we did four hundred-megabit interfaces on the card. So our first product was hundred-megabit. Gigabit was out and that was a strategic mistake on my part. We should have gone right to gigabit, but it was still so early that I felt that the hundred-megabit was a better play .

Burniece: You’re talking now, like, ’98- ’99, somewhere …

Boucher: ‘97. We started in ’97.

Burniece: Okay.

Boucher: So ….

Burniece: Yes, so a hundred-megabit was …

Boucher: Was huge and there was no market for gigabit …

<overlapping conversation>

Burniece: Absolutely. Yes. Actually, a hundred was just really starting to take off in ’97.

Boucher: Correct.

Burniece: Yes, anything over ten.

Boucher: Yes, so we basically took it to Microsoft and they tested it with four Intel hundred-megabit cards and our card. They looked at what it took to saturate the networks and I can’t remember if they could saturate all four networks with the Intel cards before their CPUs saturated or if it was about the same if the CPU didn’t quite saturate - I wasn’t physically there myself. But then they ran our card and when they saturated all four networks their CPU was idle. They couldn’t even measure the CPU utilization. They couldn’t believe it. They thought there was a problem or their test was bad. But when they finally verified it, they said, “Okay, we’ve got to do this.”

But what they wanted us to do was build a motherboard, using our technology to build a file server, and they would put their software on it and then we’d ship the thing. I felt that was starting over again, plus it was committing to them. So we didn’t do it. We decided to just sell our card without Microsoft supporting it. I mean, at the time we had built a driver for Microsoft, which they didn’t like at all, because we had to go into the Microsoft operating system, in order to build a component to allow our stack to run

simultaneously with theirs. It was basically a filter driver that would take everything that was TCP and run it through us, and everything that was not TCP and run it through their standard stack.

Burniece: Okay.

Boucher: Then it would come back to us and we’d handle it below TCP/IP, but if it was TCP we would do it. If it wasn’t, we’d let the Microsoft stack do it. Microsoft didn’t like that at all and ultimately we cut a deal with Microsoft to where we would use their software in our stack. That was one of the things that ultimately got us into real trouble, because their stack just didn’t work that well, so once we no longer could use our filter driver we started to slow down. Then we did a low-end single hundred-megabit card, which was a different IC that was much less expensive. We were getting business in both of those but not enough to seriously really be able to ramp. Then we did gigabit and all this took years, plus we did a couple of rounds of investment and basically could never get it going. In the meantime, the industry had started to copy us in ways that they felt that they could. We had written really good patents this time and we’ve now got something around sixty patents.

Burniece: Granted?

Boucher: Yes, but when we initially announced our product, and after we decided that we were not going to do what Microsoft suggested on building a competing product to what the guys selling file servers were building, we started selling the basic product As soon as we did that, everybody started saying, “Well, we can do that.” Up until then, the way you designed a software stack to do TCP was a religion. Every single layer of the ISO layer system had to be standalone. So that meant there was a copy between every single layer of the stack for everybody’s stack. For Microsoft, for UNIX, for LINUX, all of these operating systems had this religion that said, “We have to keep every layer of the stack inviolate.” Well, the programmers knew that they could go much faster, if they defined some global variables and didn’t do copies between every layer. Of course, one thing that we did in our stack was get rid of all those copies. So once we started shipping and showing how much faster we were than everyone else, the first thing that happened and as I said, this is one of the effects that we had that I don’t think anyone recognizes, we caused the religion to change. The programmers went back to the general managers and said, “Hey, they’re not doing anything special. We can fix that.” So within a year of our announcement every single one of the stacks in the industry had removed all of their copies, because it was a simple thing to do.

Burniece: Hmm.

Boucher: They put in some global variables and left the data in one place. There was one copy they couldn’t get rid of and that was the copy from the kernel to user space.

Burniece: Okay.

Boucher: While the entire kernel could work with global variables, you couldn’t give a global variable to a user space, so they got rid of all the copies except the last one. Immediately, within a year of our announcing our product, our performance was cut by close to fifty percent, maybe forty percent.

Burniece: Well, you mean, they closed the gap on your performance.

Boucher: Yes, our performance advantage …

Burniece: Your performance advantage had been cut.

Boucher: We were an order of magnitude faster now we were only about fifty percent faster.

Burniece: Ouch …

Boucher: That was not fun. <laughs> But, you know, our marketing guy kept making such a big thing of it, saying, “Hey, these guys are doing all those copies.” You know, it was obvious they were going to get rid of them. But it was a religion and so, until we did it, no one would do it. But then the problem, of course, was how do you get rid of that last copy? So everybody figured the only reason that we had any advantage is that last copy, which, of course, is crazy. It would have reduced us from sixty percent to fifty percent performance advantage maybe, getting rid of that last copy, but that last copy was really hard.

So the standards committees got together and a number of different ways to get rid of it - RDMA being the first of them - got proposed. These were ways of figuring out how to send a handle to memory from one system to the other across the Internet. This is what the RDMA-style interfaces would do. So now you’re sending over the interface, which could be halfway around the world, a handle to a memory location that is in the kernel space, which is actually a pretty scary thing. It’s something that we would never do. But it was a solution that would actually work. It used some of our technology nevertheless, which is sort of interesting.

In any case these things started to happen, in order to try and close the gap. Not too many people have done the other half of the job, which is the major half – to seriously put the protocol in hardware. Those that have, haven’t done that good a job so far.

Burniece: That’s what you call a TOE offload.

Boucher: Oh, yes, TCP offload is TOE. We didn’t invent the name TOE; somebody else came up with that. But we did TCP offload and, in fact, I told everybody in the world we were doing this, including IBM and Cisco, who got together to define iSCSI. When Cisco and IBM announced iSCSI, Andy Bechtolsheim got asked, “Well, why are you even thinking about this? TCP has so much overhead, there’s no way that iSCSI is going to work. You have to do Fiber Channel or it’s not competitive,” and Andy’s answer was, “Well, TCP’s going into hardware, so it’s not going to make any difference.”

We were really the only ones who were doing that at the time, but it is slowly going to go into hardware. There’s no question and, as it does, iSCSI’s going to become bigger and bigger. We can run iSCSI as fast as Fiber Channel. In fact, for a while we were faster, because Ethernet was ahead of Fiber Channel in terms of basic bandwidth. Fiber Channel tries to get out there before the next generation of Ethernet, even though Fiber Channel is piggybacking on Ethernet hardware technology.

Burniece: Now you did get some really strong intellectual property around that, including some patents, and did end up winning a couple of cases, correct?

Boucher: Yes.

Burniece: People were paying you license fees for your technology.

Boucher: In fact yes, but because we had been around for seventeen years and hadn’t grown fast enough, we became what the venture community knows as the living dead and couldn’t raise another round of financing. I had put as much as I felt comfortable of my own money into it, so we had to shut down operations. We had, in the meantime, finally given up on building cards and built an NFS accelerator you could put in front of any NSF file server and make it go up to around ten times faster.

Burniece: You and I talked about that a couple years ago, when you were first announcing that. It looked like a real winner.

Boucher: It was.

Burniece: What happened with that? It just didn’t take off or did it take off?

Boucher: No, it took off. It was ramping like mad, but we couldn’t …

Burniece: Okay.

Boucher: But once again, while it was ramping and we had customers, (for example, Qualcomm was an excellent customer), we were old enough that we couldn’t raise money and if you can’t raise money you can’t …

Burniece: You couldn’t fund [operations] - right?

Boucher: Right. The product was a killer product. In fact, I called Qualcomm when I was shutting down operations and said, “We’re going to shut down operations. We’ll continue to support our present product, but we’re not going to manufacture the new product.” They said, “How many more can we have?”

To give you an idea, a very high-end, over million dollar network appliance can do between two and three hundred thousand NSF operations a second. If you put us in front of it, we could do three million. That’s actually a significant difference.

Burniece: Absolutely.

Boucher: So …

Burniece: So you have actually stopped making those?

Boucher: Yep.

Burniece: And is there a demand that you could go and sell it to somebody that’d take over?

Boucher: Oh, Qualcomm would love it, if somebody started building these things. They’re still using them - all the ones that they could buy, they’re still using and we’re supporting. But I’m supporting them all on contract. Our engineers are all now working elsewhere but they’re on contract to Alacritech for support. So my job today is two things: It’s supporting the field on our present products and going back to all the people that are using our technology but not bothering to take out a license and asking them to take out a license or …

Burniece: Okay, so you’re trying to collect license fees …

Boucher: Which is not a lot of fun, but <laughs> yes.

Burniece: So if you look back on that, Larry, what were the two or three things that you wish you’d have done a little differently with Alacritech in particular?

Boucher: I really didn’t want to build a file server accelerator because I didn’t want to compete with Auspex. I made the same deal with Auspex that I made with …

Burniece: Adaptec when you left?

Boucher: [Yes] and the same deal that I’d made with Shugart when I left [to found] Adaptec. So, while I hadn’t hired anybody, there were people that left Auspex, not many - but the few that did, talked to their management and gave Auspex a real chance to make sure they didn’t leave, [then] they decided to leave and go to work for nothing. At the time that I started Auspex, the venture community was paying startups ridiculous amounts.

Burniece: Hmm.

Boucher: The venture community was putting so much money into startups in the late 90s that they could pay more in salaries than the major companies, which was just crazy …

Burniece: That’s remarkable to think about. Right.

Boucher: It was very remarkable, but we started Alacritech, based on the same principles that I started the last two, which is that there were no salaries until we had raised money and then we took very low salaries. Because of that we were able to survive for a long time but the engineers had worked with low salaries for seventeen years - I mean, low compared to what they could’ve gotten anywhere else. So it was a little on the too-bad side from their standpoint. A major investment for all the key individuals and I’m hoping that ultimately it’ll still pay off for them. That’s what I’m trying to do is get a payoff for the investors and employees.

Burniece: Fantastic. So when you look back again on the whole journey you’ve been on, who are some of the people that you really admire and maybe haven’t mentioned before, who really influenced you or just admired what they did?

Boucher: I think we’ve discussed a lot of them. I had very early influences. I guess the very earliest influences were reading science fiction and I think that’s what got me into being an engineer. Science fiction, I just really enjoyed it. I was done with science fiction by the time I was, I don’t know, a sophomore in high school, but by then I was committed to being an engineer and it was Robert Heinlein, Isaac Asimov, Arthur C. Clark and Ray Bradbury …

Burniece: These are writers?

Boucher: Yes. I never knew them, but they were a big influence in terms of what got me interested initially. Then, of course, I grew up in San Carlos and we had Lenkurt, EIMAC [Eitel McCullough, which was eventually bought by Varian], Dalmo Victor, iMac and Hewlett Packard, which had grown up through the 40s and 50s. That’s why I got into microwave initially, because it was so big in our area.

Burniece: Absolutely.

Boucher: I sort of moved in that direction and I mean it was just luck that the semi-conductor came along. Then after that IBM, which had been just a small card plant, decided to move into the development of the disk drive here and the rest is history.

Burniece: Yes.

Boucher: But, yes, that was early development and then there are a lot of people I haven’t mentioned, like Al Hoagland, who wasn’t an early influence but I got to know him later, Ken Houghton, who we already mentioned, and Ted Codd. I did my master’s thesis based on some of the stuff that Codd was doing …

<overlapping conversation>

Burniece: He’s the guy that basically invented the relational database.

Boucher: Right. These were really smart guys, who were still really down to earth. Hard to stay up with them ….

<laughter>

Burniece: Right.

Boucher: … if they started to talk about anything, it was challenging.

Burniece: So what are you most proud about in your career? What-- when you look back on, say, “Well that was really a prime moment-- premier moments”?

Boucher: Uhm - most proud - I think the single biggest thing that anyone does in their career is raise kids. So the thing I’m most proud about is …

Burniece: Okay.

Boucher: .. raising a couple of great kids.

Burniece: How many kids do you have?

Boucher: Two daughters. One of them teaches high school. She’s now the dean of admissions at Archbishop Mitty. The other is a clinical psychologist. She has her doctorate in psychology and has her shingle out. They’re both married; they’re both wonderful kids and I think we’re all responsible for bringing the next generation that’s going to take over for us.

Burniece: No doubt about it.

Boucher: And I’m really proud about that. After that, I guess, probably the fact that if you look at moving serial protocols into hardware I feel I made a real difference, even though I don’t think that it’s …

Burniece: People probably wouldn’t know that.

Boucher: I don’t think people recognize that but, from my own personal standpoint, I get recognized for SCSI, and SCSI was just sort of an obvious thing. At least, it seemed obvious to me, especially knowing Wally Bass and seeing what he’d done. What Wally was doing was just trying to make things simpler.

Burniece: Right.

Boucher: But really that’s all that we were doing was making a simple enough interface that it can survive multiple disk drives so I didn’t feel like I was doing anything all that different. I was just taking what Wally had done and making it work.

Burniece: But the key thing that you’ve talked about on SCSI is that you layed out a vision that ended up lasting for …

Boucher: Well, that’s true.

Burniece: … forty years or whatever it is and it’s still going. [SCSI is] still one of the key aspects of the storage side of computing.

Boucher: Yes, I agree with that and you’re right: I will take credit for the physical separation of the logical interface from the physical interface, and that is the reason that it could live that long.

Burniece: Right. And, of course, it all went serial. Everything’s gone serial.

Boucher: Right.

Burniece: And you’ve enabled that to happen, too. Fantastic. So, talk a little bit about the future. Where do you see the future of computing and storage going? What’s your vision on where it’s going to head?

Boucher: Well, I don’t see things slowing down and certainly artificial intelligence is becoming more and more real. We’re still a long way from “I, Robot” but, if you look at just in the last ten years, where speech

recognition has gone, it’s pretty amazing. And you look at the fact that we now have cars that you still have to keep in the same lane but, other than that, it’ll take care of you.

Burniece: And they actually now have some they don’t even put steering wheel on them, because …

<overlapping conversation>

Boucher: And that’s around the corner at this point, which is scary. But if you can actually get in a car and be able to maybe only secondarily pay attention to what it’s doing, and ultimately not pay attention to it …

Burniece: Right.

Boucher: … so it will be the equivalent to getting on a bus …

Burniece: Right.

Boucher: … and we’re not that far from that.

Burniece: No, we’re not.

Boucher: That’s one of the areas where we’re applying artificial intelligence, today. Still, you think about the definition of artificial intelligence and it’s just the next generation of computing. Computing is about what it’s going to do for us and, what can we use it for? Ultimately, what I was reading about from Asimov in the 50s is going to happen.

Burniece: Right.

Boucher: And it’s not going to happen that far from now, because the pace of development is accelerating - it’s an exponential, not a linear, function. So …

Burniece: So if you were to look back and say, “Wow, wished I’d done that,” or “what would I do differently?” Is going to artificial intelligence maybe something that you would’ve liked to have done? Or don’t you feel being on that bus was an important thing, although it’s certainly important?

Boucher: Well, I live most of my life in the world between pure hardware and pure software.

Burniece: Right.

Boucher: What gets named middleware, which is both hardware and software in some sense. But I really think while I just fell into it when I went to IBM, I loved it, still love it and wouldn’t mind still working in it. I’m probably spending thirty percent of my time on support and trying to make the intellectual property of Alacritech successful for the investors. So I have another seventy percent of my time that I could be doing something [else] and I would love to be back doing an engineering job. And if I was, I wouldn’t want to go into pure software or pure hardware. I really like living in this sort of nether land that I’ve been in for the

last forty years. So, I don’t think I’d do too much differently and I feel like I’ve been doing a big piece of what’s required for artificial intelligence, because that’s what it takes …

Burniece: Right.

Boucher: … [it] is moving computing to the next generation to get where we need to go. When you say, “Where are we going?” I feel like there is no destination. It’s all a journey, but I think that what I see on the horizon is a much higher level of artificial intelligence.

Burniece: So if you were to give advice to a young person starting now would you say, kind of like in “The Graduate” when he said, “plastics”, - artificial intelligence?

Boucher: Yes.

Burniece: Like that.

Boucher: Absolutely.

Burniece: Go work in that field.

Boucher: Yes, it’s still software, if you will …

Burniece: Right.

Boucher: The particular field of software that I would be going into right now would be robots.

Burniece: Right.

Boucher: And another is speech recognition. Yes, that would be an area of interest that will be growing very rapidly.

Burniece: And a fun place to be.

Boucher: Oh, yes. No question.

Burniece: All right, well, we’ve gone almost three and a half hours here, believe it or not. Any final thoughts that you’d like to wrap up with? Or [other] things you’d like to talk about for a few minutes?

Boucher: No, we probably could run through about a thousand, as we were going, that we didn’t hit when we were going down different avenues <laughs> …

Burniece: We could have …

Boucher: God help us if we tried. <laughs> So, I think you’ve probably hit the high spots.

Burniece: Larry, it’s been fantastic. We appreciate it and that’s a wrap.

Boucher: Thank you.

Burniece: Take care.

END OF INTERVIEW

