Oral History of Grant Saviers

Interviewed by:
Tom Gardner

Recorded: May 17, 2011
Mountain View, California

CHM Reference number: X6162.2011

[0:00:00][1]

Gardner: I'm Tom Gardner, a volunteer with the Storage Special Interest Group at the Computer History Museum, and this morning I'm here with Grant Saviers, to mainly cover Grant's remarkable history at Digital Equipment Corporation. Probably a good time for the prisoner shot.

Saviers: Right, an original sign from DEC, when it was a 6,000 square foot operation at Building 12 of the Mill [DEC Maynard MA facility]. Of course, once you've been at DEC, you're always a DEC-ie, so I saved this when I left.

Gardner: Grant has a long career in the computer industry, long time at DEC, and then with Adaptec. We'll probably not get to Adaptec in this recording session, we'll leave that for later, and we probably will not cover all of Grant's career, focusing today mainly on the storage aspects of Grant's career. Grant, tell us about your early background, where you grew up.

Saviers: Well, I was born in Baltimore, Maryland, and my dad was learning to run a machine gun, so he could go off to Europe and fight in World War II, and I grew up as the only child in a family which was not a technical family, my dad was an accountant, my mother was a housewife, and at some early age, I got interested in technology, you know, the kids' chemistry set, how to blow up things, and things like that, and I got into electronics via ham radio, and that was when I was 13 years old. Built many Heathkits, and many electronic things, and decided electrical engineering was for me. Went off to college, 1962, and got a bachelors and masters degree in engineering. This was when my alma mater was Case Institute of Technology. It later became Case Western Reserve University, but there was no such thing as a computer science degree at that point. So I got an electrical engineering and kind of computer engineering degree, all wrapped up in one. Left there in 1968, with a masters degree, lots of interesting stories to tell about my career at Case, and maybe you want to get into that a bit, since there's some early computers there, but then joined DEC in 1968.

Gardner: So were you at Case, when they taught tubes first or transistors first?

Saviers: It was transition time. The most venerable professor was known for his lightning protection of oil fields, and taught transistors and tubes. He was Doc Hoover, a great instructor, and looked like the Colonel in Kentucky Fried Chicken. Anyway, I fell in love with computers at the end of my freshman year, and there was no computer course until the sophomore year, but I discovered that in the computing center was a Burroughs 220, which was a tube machine, and Case had a very different philosophy about students and computers. They allowed the students to touch the computers, so we loaded our own cards, we ripped our own line printer paper out, punched our own cards, and pushed the button on the front of

[1] Time stamp in form of H:MM:SS from start of file X6162.2011.Saviers-QuickTime 480x 270i H.264 CHM.mov

the computer, to start everything. So that was the Burroughs 220. I learned to program in Balgol, Burroughs Algol, a version of Algol 58, and shortly, then, Case bought a Univac 1107. Interestingly, it also had one of the last surviving Univac I's, the original Eckert-Mauchly marvelous computer, mercury delay line memories, and I got so enthralled I almost flunked out my sophomore year. I managed to convince the powers that be at the computing center to give me a job, and I worked there for the next six years, I guess, five years, maintaining things like the Univac I, the 1107, building specialized equipment, writing programs, a little bit of everything, this and that. As a graduate student, I had some graduate student administrative and course management duties, but most of the time I was hacking around with electronics and software.

Gardner: Any interesting Univac I stories?

[0:04:26]

Saviers: Well, you know, yes, I mean, it was a marvelous machine. One of my jobs was testing the tubes for it. It had something like 5,000 tubes, I think, and I think 4,000 of them were all 25L6s. If you're an audio buff, you know what a 6L6 is, it's an amplifier tube, a 25L6 is a 25 volt version of it. Anyway, we bought only from Sylvania, only case lots, all of one production run, and we'd buy enough for two years of maintenance. We'd test every tube, burn it in, whack it with a little mallet, to see if it had any microphonics or mechanical sensitivity, and then after all that was done, we'd plug it into the most sensitive part of the computer, which was the memory recirculating amplifiers on the delay line. If it failed, you knew immediately where it was, and you could just go in and pull the tube out and put a new one in. After that, those tubes migrated into the logic and control parts of the computer, and probably would last 10 years. The machine was turned on and off every day, it would only run from midnight to eight a.m. The Chesapeake and Ohio railroad had it as a spare computer to the one they had downtown Cleveland, and I think it probably broke once a month, maybe a several hundred hour MTBF. It was pretty impressive.

Gardner: How long did the tubes last, and how many tubes did you use?

Saviers: Oh, I think maybe 250 or something like that were qualified at a time. I think there may have been tubes in it from the original day it was installed, when it was decommissioned, and it had to be decommissioned with jackhammers, because the floor had special footings, and the cooling was a four inch copper water pipe with fins on it that ran cold water in it, air recirculated through the computer to cool all those tubes. I think the power drain was something like 75,000 watts. So it was a pretty impressive machine. It was a lot of fun. I learned a lot, and of course. An interesting side story, being affiliated with computer museum, much later, I had dinner with J. Presper Eckert, which was, wow, to meet the original guy who designed this computer was just fantastic.

Gardner: Any stories from that dinner?

Saviers: What a sharp guy. You know, he was, I think, in his early 80s, and I was very interested-- there's lots written about the computer itself, but I was very interested in the magnetic tape drives, because it was the first magnetic tape drive on a computer, metal tape. So I started asking him questions about how they did this, and how they did that, and why they did it, and he remembered it all. It was quite an interesting evening.

Gardner: That was a great segue into the Uniservo I, and I want to ask you about that.

Saviers: Oh, okay.

Gardner: You know, can you recall any of your experiences with it?

Saviers: Well it was probably the least reliable part of the computer, because, instead of vacuum columns, it had springs, kind of elastic cords, and pulleys to tension the tape over the head. It was a single capstan tape drive, it could run forward or reverse, read forward, read backwards, it was phase encoded, and used a Macintosh audio amplifier to drive the capstan so it could drive it at any frequency within the audio range, so rewind was a different speed than read or write. Had a little piece of Mylar film that went over the head, had a little clock mechanism to advance it slowly, so the metal tape, being very abrasive, wouldn't wear out the recording heads, so the recording heads lasted forever. It had ten tape drives on the machine that Case had, which I think was the standard. Most Univac I installations were like that. And The Univac I had overlapped I/O, so it could actually read and write a tape while the computer kept going. It had a separate delay line memory, mercury delay line memory for the tape I/O buffer. So-- and one instruction would take the 60 words out of that buffer, and swap them with a 60 words in memory in one instruction, so it was very, very clever design, way ahead of its time, in terms of building a computer that did data processing. It was the first time that data was stored on tape, as opposed to cards, a big transition in computing.

Gardner: And I think you've written a paper for the Computer History Museum on the subject?

Saviers: Yes. I did, on the Uniservo, but there's still, I think, a lot to be written. I haven't seen very much written, really, on what the internals of the Univac I was about, and why it was so brilliant and ahead of its time.

Gardner: Anything you care to put on tape right now?

Saviers: Well of course it was a fully redundant computer. Every arithmetic computation was done in parallel, and the results checked, bit by bit. They didn't know at the time, how reliable a gate of a flip flop was, they had no idea what the contribution of noise would be to the reliability of computing circuits. So they said, well, let's take the brute force approach and make it completely redundant. Lots of check circuits and margin checking, it was well engineered from a mechanical viewpoint, easy to service. I mean, you walked inside to do scope probing, and one of the fascinating things was the dead man cord overhead. There was a cord like in a bus, you know, I want to get off at this stop. When you pulled that cord, it dumped all the high voltage power supplies, assuming that you were in dire straits.

[0:10:22]
Gardner: And those were high voltage.

Saviers: Oh yes, those were hundreds of volts, right.

Gardner: Not like high voltage today, is five volts.

Saviers: That's right

Gardner: Any questions you asked Eckert that you got good answers for, you'd like to tell us about?

Saviers: I wish I had a tape recording of it, and could remember them, but it was one of those wonderful experiences of life.

Gardner: Your noise and redundancy comment is a great segue into something that we'll cover a lot later in this taping session. So you got an offer at DEC?

Saviers: Well, yes, I wasn't sure I was going to graduate school, so at the end of my bachelors, I interviewed everybody, Burroughs, Control Data, Univac, IBM, and got, I don't know, five or six job offers. Then decided, well I'm going to go to graduate school. There was this little thing going on, called the Vietnam war, at the time, that might have had some influence on my thinking, but anyway, I went to graduate school, and then, after graduate school, I said, there's only really two choices, DEC or IBM, and got an offer from both of them, and decided that, gee, if those crazy guys up there in that woolen mill, all these MIT refugees, could do all this neat stuff, I think that would be much more interesting, even though IBM was going to pay me something like 25 percent more money, it was DEC I chose, and it was a great decision. They were on a roll, and it was very-- it was kind of like, this is college as a business. You know, it was a university environment in terms of freedom, what you can do, what you contribute, what you can say, and that was very appealing to me.

Gardner: Which laboratory did IBM offer you a position in, do you recall?

Saviers: It was Gaithersburg, Maryland, it was their Federal Systems Division, and it was interesting, because I would have been the assistant to one of the PhDs who was one of the world's experts on data transmission. I did my Master's Thesis on data transmission, so that kind of fit. I also had a lot of family nearby.

Gardner: And what data rate was data transmission in those days?

Saviers: Well, for my Master’s thesis I had built a double speed acoustic modem, and on the fly error correcting system for connecting a terminal to a computer, a teletype to a computer, so--

Gardner: And double speed being how many baud?

Saviers: Six hundred.

Gardner: And today we get ten gigabit or something.

Saviers: Something like that, right, and of course, it was interesting joining DEC. Here's IBM, this huge company, multi-billion dollar company, and DEC, when I joined in 1968, was 57 million dollars in revenue, so it was this really tiny company.

Gardner: And what was it when you left, in revenue?

Saviers: Thirteen billion.

Gardner: And how much of that was due to you?

Saviers: Well, the storage business, which was disk, tape, memory, optical, all the subsystem work, was about four billion in hardware and service revenues.

Gardner: And a small part of that 57 million dollars in the very beginning.

Saviers: It, you know, might have been ten dollars. I should take that back. DECtape was there when I arrived at DEC, and there is an interesting story about DECtape, but you know, it was, I think, really the first personal storage media for computers that existed.

Gardner: Let's put DECtape aside, because it is definitely an area we'd like to talk about. You had an unusual first day at DEC

Saviers: Yes, DEC in those days, and for many years, was this “run the gauntlet” interviewing process. I think the day I interviewed at DEC, I interviewed with 12 people, and at the end of it, they said, well, we think we want to make you a job offer. We'll get back to you about which of those 12 people want to make you a job offer. Okay, so a week goes by, and the HR guy calls me up, and says, well, we want to make you two job offers. You can decide which one you want. One is to be a technical advisor to Ted Johnson the vice president of sales, and the other is to work for Ed de Castro, who is running the new PDP-X 16 bit computer project for DEC. Well I thought about it for a bit, and said, I think I want to do engineering work for a while, I don't want to get off now into sales. And so I said, Ed de Castro. And they said, kind of like, are you sure, because I didn't know it, but Ed had a pretty abrasive side to him, I think, that was--some people at DEC didn't like him much, but anyway, I said, yes Ed de Castro, so okay, come May, 1968, I finish up my Masters work, and go up to Maynard and go into the lobby and say, I'm here, new employee, first day of work.
[0:15:00]
They said, who are you working for? I said, Ed de Castro. And there was a silence. They said, Hmm, we'll get back to you, have a seat. So I sit in the lobby for about four hours, and finally this guy shows up, Don Vonada, and he says, "Well, we thought you weren't coming." I said, "Why is that?" He said, "Well, Ed and 12 other guys left two days ago to form Data General. We thought you were going with them." I thought, oh, that's really interesting. So that was my first day at DEC, and the place was in complete chaos, DEC had never had a defection before, Ken was apoplectic, and they weren't sure whether they were going to continue the PDP-X, do something else, or build another PDP-9. Anyway, Don found me a job, we're going to redo part of the PDP-9, make it cheaper. I completed that, it shipped, and then they decided that they were going to build another eighteen bit computer, which was the PDP-15, so that was, I guess the fifth 18 bit computer in the DEC line, the 1, 4, 7, 9, 15. I think that's the right sequence. Hard to keep them all straight. But anyway, at that point, John Jones the 9/15 Product Line manager said, we need a disk controller for the PDP-15, and there's this guy in the PDP-8 group designing a disk, why don't you go talk to him, see if you can make it work on a PDP-15? So I started on that project.

Gardner: You mentioned the PDP-X?

Saviers: Yes, the PDP-X was DEC's first serious attempt at a 16 bit computer, and I guess one of the reasons that de Castro had left, was there was a big debate about its future, and Ken decided not to do it, and that was the beginning of the DG computer, the Nova. The X and the Nova architecturally had nothing to do with each other, except Ken Olsen frequently claimed that they designed the Nova while they were still working for DEC. So who knows?

Gardner: So the PDP-11 was then a different project.

Saviers: Completely different project, yes. You know, people at DEC tended to be banished by Ken Olson, he never fired anybody. But after the PDP-6 absorbed so much of the company's resources and it was kind of a shaky business proposition, and Gordon Bell was the project leader on that, he had been banished, and he went off to Carnegie Mellon for a while. Well, then he came back, and was redeemed. It's like the gulag, I guess, you go off to the gulag for a while, you come back, you're okay. So anyway, Gordon comes back and says, here's the kind of computer we should design, and he and Harold McFarland, who was the chief architect, put together the PDP-11 design. I consulted a little bit with it, made some dubious contributions, but you know, we were way behind the curve, since there were other 16 bit computers out there, including Hewlett Packard. DG was very short in its time to market, and so it was a catch up game.

Gardner: You then go off to do a controller for DEC's first disk drive.

Saviers: Yes, let me back up. When I arrived, there was a product just going into production, called a DF-32, which was a fixed head, head per track disk drive. And designed by a group that had no degreed engineers in it. They did a very clever job, and it was 32,000 12 bit words for the PDP-8. Now you might say, why would you bother? But if you were a programmer, using a PDP-8 with only 4,000 words of memory, you'd have thought you'd died and gone to heaven having 32,000 words of reasonably fast memory, disk memory attached to it. That product was in production, but having problems, and the group went on to design a drive that they scaled up by about a factor of two to the seventh. You know, everything they changed was a factor of two, but not really thinking about it from an engineering perspective, and didn't realize that this was 100 times more complicated. And it didn't work, so I was designing this disk controller for a disk that was probably, likely, never going to see the light of day. The powers that be figured this out, and said, "Okay, Saviers, you're on the task force to fix the drive." It was Roger Cady, who was designing and managing the PDP-11 development, so he was kind of busy. Tom Stockebrand who had a history of magnetic recording, since he was kind of the father of DECtape, after LinkTape came to DEC, and me. Now I had never, ever seen a disk drive in my life before. Roger and Tom had full time jobs, and this was all on me, and a couple of technicians.
[0:20:03]
So I'm trying to figure out what a disk drive is, and how it worked. By the way, the day I started working on the disk part of this subsystem, we had put two startups out of business, had three lawsuits and were three months past the committed first revenue ship. So the company was a little bit concerned about its future. So anyway, I started to learn about disk drives, and there's a long path of learning and trial and error, and finally, after about nine months, and by cutting the capacity by a factor of four, to 262,000 12 bit or 16 bit words, actually, we did make an 18 bit version of it, too. We shipped something that kind of worked. And that was my first disk drive.

Gardner: This is the RS08 drive?

Saviers: Yes, the controller for the PDP-8 was the RF08, and there was an RF15 for the PDP-15, which was the original project I was supposed to be working on, and RF11, for the PDP-11. So it covered 12, 16 and 18 bit computers.

Gardner: A new disk drive in nine months, one person, two people, that's incredible.

Saviers: Well, I hired an engineer or two along the way, and there were some good technicians around, but, learning about, what is a flying head, the magnetics, etc. We were buying plated disks from a company out here on the west coast that the guy was a former car bumper plater, -- he liked to plate rhodium, and we had rhodium on the surface of the disk, and this caused all kinds of grief -- anyway, we couldn't make it work with the rhodium disks, they had very unusual chemical reactions going on inside of the HDA, something called frictional polymer, which Bell Labs can tell you all about. But anyway, we ended up finding a company out here in Mountain View called Data Memory, which was making a video disk recorder, one of the very-- if not the first stop action and replay video disk recorder, and selling them to CBS for the NFL. Now their disks wore out after maybe five or six games, so they were making disks and replacing them, but they made a disk which had a better protective coating on it. They actually thermally oxidized the magnetic plating to make cobalt oxide, beautiful deep blue disks, I mean, a gorgeous color. So that worked, and we bought the company, DEC's first acquisition, and started making disks in California and putting them in the products back in Massachusetts.

Gardner: I believe you told me the Rhodium problem manifested itself as brown crud.

Saviers: Brown crud, right. We'd take a disk and absolutely make it spotlessly clean, clean everything in the head disk assembly, and come back 24 hours later, and there would be a film of brown crud on the disk, because the mechanically excited low energy electrons polymerized organic vapors and turned them into polymers, solid material. Amazing. People don't believe it today, but it happened.

Gardner: I think, at this point, we're going to head down your experiences in buying disk drives, maybe do a segue into building the first cartridge drive, but I think you then were more of a purchaser than you were a developer.

Saviers: Yes. There were a number of drives that DEC had been purchasing. There was this Data Products monster that had, I don't know, 24 inch disks, that was connected to the PDP-6, which was kind of ancient. The PDP-10 group bought some Bryant, the hydraulic actuator, those big things that weighed four tons or something like that. And then, the IBM copycats got in the game, so the plug compatible market started to move, Century Data, Memorex, and I think Control Data were around at the time. But the PDP-10 folks wanted a disk pack, and so they were out looking at disk pack drives, and since they thought that I knew something about disk drives, asked me for my opinions about, what should we buy. There were Century Data drives and Memorex drives and so on, and after looking at it a bit, we said, well, let's go with the Memorex drive, and I think that was what we called the RP01, which was a slightly modified IBM compatible drive. The PDP- 10 folks had developed their own controller, and I remember that one of the real strains in getting that out the door, was trying to make the data separator work and, because everything had been designed for IBM formats and IBM systems technology, etcetera, and all had to be converted to the fixed block length approach that DEC used on everything.

[0:25:21]
Gardner: I believe that was probably John Green who actually did the controller, or Dale Ives?

Saviers: Yes, Dave-- both were certainly involved. I couldn't tell you who was the lead. I know Dave was also looking at, then we did some looking at head per track drives for the 10, because in that time, memory was extremely expensive. It was all core memory, of course, and it was extremely expensive. So it made sense to have something fast, head per track, as well as something big, the moving head drive, and I think, I remember going out to California to look at some companies with Dave, to evaluate who was doing what in head per track drives, and drums, also. DEC had bought some drums from Vermont Research over the years before I arrived.

Gardner: Is it true that the Bryant drives were installed once on the top floor of the old mill, and a truck backed into a loading dock and crashed all the heads?

Saviers: Mixed stories - not quite. The Bryant drive was so heavy we had to put it in the subterranean basement of Building Five, which frequently was underwater, because the mill was on a pond, but we fixed that problem. Anyway, so the Bryant drive had to go in the basement. The PDP-10 manufacturing was on the fifth floor, and this was a circa 1910 building, 500 feet long, 125 feet wide, five and a half stories, and when the trucks bumped into the loading dock, you could feel the whole building rock back and forth.

Gardner: Just let me interrupt. But as I recall, there was a hillside road on the uphill side of the building, and there were loading docks on several different floors, like the second floor, and the third floor.

Saviers: Yes, the main loading.

Gardner: And they were actually downhill from the road, so if you lost your brakes, you could go and hit the loading dock on the third floor of the building, or something like that.

Saviers: Yes, the main loading dock was on the fourth floor and the PDP-10 production was on the fifth floor, and most of the PDP-10 engineering was up there, too. Anyway, Alan Kotok and Alan Kent -- they were both gurus of PDP-10 land, but Alan Kent was keeping track of what was going on with the disk drives that they were buying, which was from Burroughs. Interestingly enough, the Burroughs model number was 9760, which reappears later in disk drive history. But anyway, it was a, I don't recall, 20 inch, no, 25 inch, 28 inch disk that sat vertically, and was actually the disk that Burroughs had parallel tracked for the Iliac computer, and was using on their own computers, and they plated them. They made the disks in Westlake, down in Southern California. Well, anyway, this is probably one of the most unreliable disk drives ever developed, and manufacturing had to line up all the spindles, so that when the dock was hit, it was along the plane of the disk, rather than perpendicular to the disk. Alan Kent kept track of every disk that we bought, and I think we bought more than 100 of them over the life of the project. But he could show you the day we bought it, the day it shipped to the customer and the day it failed, and I don't think any of them survived more than eight months. So they were very motivated to get an IBM compatible disk pack drive on the PDP-10.

Gardner: So you were involved in the purchasing of the cartridge disk drives as they emerged in the late '60s.

Saviers: Yes. A soon as the trickle down of the guys building large IBM compatible disk drives, trickle down into smaller cartridge drives and the 2315, which IBM had first shipped, I think on the 1130 data processing system.

[0:29:52]
Gardner: I think it was announced first on the 1800.

Saviers: Yes, I think you're right. So that was a front loading cartridge, and I think IBM shipped it as a slightly more than a megabyte of capacity, 1.2 megabytes, so companies started to get formed, and there was probably six to eight of them, Iomec, Caelus, Diablo, Computer Memory Devices. There were a number of them I remember, so we started looking at those in terms of buying from them. Andy Gabor and George Comstock from Diablo came to Maynard and they brought their disk drive in a briefcase, they called it. They had just started, and so Andy had the spindle, which was a DC motor, which was a different idea altogether, and the actuator, which was another DC motor, and so he said, these are the critical parts, and we're going to build a Diablo model 30, which was going to be 1.2 megabytes, and a model 31, which was going to be two and a half megabytes. And you know, of the available choices, it looked like the way to go, and so we said, well, you guys don't really exist yet, but we're willing to work with you if you give us a manufacturing license. So back and forth we went, and we signed up a deal, where they were going to provide the Diablo 30 as engineering prototypes, because we really don't want to ship that size. And then, the Model 31 was going to be the mainstream product. So Tommy Orr designed the PDP-11 controller and started qualifying the drive, and it was a good drive, certainly was the best choice, and we're Diablo's first customer by far, and their biggest customer by far, later. So I got that company going, and we had this manufacturing license This thing is taking off, much more so than we expected, because basically the cartridge disk drive became the replacement for DECtape, it became the new personal storage media. Maybe just set some context here, DEC at the time, had 8s and 11s, and

15s, and we probably had two or three or four, or five operating systems for each of those computers, so a programmer wanting to use a computer, could have operating systems, their data, their software, all on a disk cartridge, and just walk up to any computer stick it in and load the operating system and do what they wanted to do. So it was a wonderful improvement over DECtape, because it was fast. It was big, I forget what a DECtape capacity is, maybe 300 kilobytes[2], so it's a huge improvement in capacity, a huge improvement in performance. So it was clearly going to help the mini computer business grow dramatically. So anyway, we were looking at this and decided that we better do something, either exercise this manufacturing license or get in the drive business ourselves. So in DEC's typical decision making process, which some people have characterized as, any decision worth making, is worth making five times, we got down to three o'clock on the day that the option to buy the manufacturing license was going to expire, in Ken's office, talking to Ken, "What are we going to do here?" And at the same time, Honeywell and GE had merged, and decided that all the peripheral operations in Billerica, Mass would move to Oklahoma City. Now, for Bostonians, this was going further into the wilds of the unknown, say than Siberia, so the Billerica team was in complete revolt, and we were getting resumes by the pile. So anyway, we say to Ken, "Look, we can exercise this license, or we can hire a bunch of engineers from Honeywell, that know something about disk drives, moving head disk drives. So, you know, like in typical DEC fashion, three minutes before the option expires, we call up Diablo and say, well we're not going to exercise the license. They probably thought that was a wonderful thing. So we then say, well let's hire some engineers from Honeywell, and we made about five or six job offers. It was right in the midst of one of those depressions in the electronics business, and DEC had a complete hiring freeze on, so Ken was having a hard time with that. But we got the job offers out, I think on a Friday evening, and they all accepted, and on Saturday morning, Ken calls me up and says, "Rescind all those offers, I don't want you to hire those people." I said, "Sorry, Ken, it's too late. They already all accepted."

Gardner: This was 1971?

Saviers: Yes, I think 1971.

Gardner: Okay, and how much was the Diablo license, do you recall?

Saviers: I think we bought enough to make it nominal, like 50,000 dollars or something like that, maybe even free. It wasn't much, just enough to pay their actual license transfer costs.

[0:35:00]
Gardner: Well those were 50,000 1970 dollars.

Saviers: Yes, they were real dollars.

[2] Editor's Note: DECtape capacity was either 184K 12-bit words or 144K 18-bit words.

Gardner: Right, and was it the license that was the determining factor? Did you make the same offer to, say, Iomec?

Saviers: No, we ruled out Iomec -- didn't like their design.

Gardner: Really?

Saviers: Yes. They were a fixed plus removable design, but we didn't like that from a systems perspective, because even though it would tend to give us fits sometimes in pricing, it got real messy to make a copy of a disk, especially if the operating system is sitting there on the fixed disk. So we said, no, two removables is the way to go, and that's the pattern that we followed from then on. We never built a fixed plus removable drive.

Gardner: Diablo did, though

Saviers: Fixed plus removable was very popular, as an entry point. But if you want to stack up three or four of them, it made no sense whatsoever.

Gardner: I think it may have been popular more intellectually than physically. People really bought two cartridge drives rather than a fixed plus removable.

Saviers: Right, right, I think that's correct.

Gardner: That was my perception in that time period.

Saviers: Yes, Hewlett Packard brought out a fixed plus removable. That was probably the most serious competitive issue we had to deal with.

Gardner: So you have these guys from Billerica?

Saviers: Billerica. Honeywell Billerica.

Gardner: And again, it's remarkably small group. We talk about five or six, compared to what development is later, or even today, in 25, 50, 100 people.

Saviers: Well I said Ken got even with me about a week later. This was October.

Gardner: 1971.

Saviers: Yes. Maybe the license expired, like, September, the end of September. So a week later, it's the first week of October, Ken calls me up on a Saturday, and says, "Since you've got all these guys now that know something about disk drives, why don't you take the disk drive-- make a disk drive and take it to the Fall Joint Computer Conference?" I said, "Ken, that's five weeks from now." It was in Houston. He said, "Yes, I know you can do it." So okay, so I get all the guys in that Saturday, and said, "Ken wants us to have a disk drive at Fall Joint Computer Conference in Houston." So we sat down and figured out, what could we cobble together? Now fortunately, a couple of the guys from Honeywell, were very good engineers, in fact, one of the best was a guy named Elmer Simmons, who went on to be a top ranked consulting engineer at DEC. He knew the read-write and the logic code, so he could put that stuff together real quick, and another guy named Dale Jensen, who was the servo guy, and he had built servos before, so he could build that pretty quickly, we had hired a mechanical engineer, Gunter Schneider, from-- a really good German mechanical engineer, so we said to those guys, look, turn over what you've got to some of our other less experienced disk drive guys, Pete McLean and Peter Svendsen made big contributions, and they'll put it together, and you try and keep on the track of designing what became the RK05, , Five weeks later, we take a working disk drive to Fall Joint Computer Conference, and showed it on the floor, The intent at the time, because DEC was still behind Data General in 16 bit computers, was, “we're changing the name of the game.” That was the banner for our show, and the first DEC writer, the LA30, was put on the same crash program to be there, a tape drive, the TU10 was put on the same program, and the disk drive. So the idea was, peripherals are going to define the future of mini computers, not being able to put random 7400 series TI MSI IC’s on the circuit boards, because I think, at the time, there were, I don't know, upwards of 50 mini computer companies that had been started, so somehow, DEC had to say, yes those guys can make a computer, anybody can make a computer now, but only DEC can build you a system.

Gardner: Tell me more about this five week wonder.

Saviers: Five week wonder. Well, I mean, it was 24 hours a day, for five weeks, and so we took a piece of aluminum jig plate, and cut holes in it, and made an actuator, using linear bearings that we knew were completely ridiculous to build a disk drive out of. But we could do it in five weeks. And a voice coil motor, and put this thing together and it would read and write. We didn't have a controller for it, we just had a little exerciser box, so it could do random seeks, and it worked the whole Fall Joint Computer Conference, and before the conference opened, I invited the Diablo guys over to say, look, you know, we didn't do this license and we're still going to be a customer. It's still going to be a while before we really get into the business ourselves, but-- and their jaws just dropped, and they said, oops.

[0:40:15]
Gardner: DC motor, or AC motor?

Saviers: Well, another interesting story there. Ken said to me, when you design the DEC disk drive, I want you to survey every disk drive being made, tell me, in the major components, how it's done, and come to the Operations Committee, the top committee in the company, and explain how each of these is done, and which one you're going to pick. So I did that, and said, you know, there's probably intelligence out there of why people have done it this way. So an AC motor, and a voice coil, AC drive motors, spindle, separate spindle, belt, voice coil motor, optical encoder for the feedback, and that’s how it turned out that. There's another whole story there, but it turned out that the 2315 IBM cartridge was an abomination, an awful choice, even though it was designed by IBM. We didn't know any better. We thought IBM knew what they were doing, but IBM didn't, so we all followed them down that path and paid the price for it.

Gardner: Tell us some more about going down that path from this curiosity at NCC, to Fall Joint Computer Conference.

Saviers: Yes, -- there were four engineers, and three technicians on the project, and we got the product into production about the fall of the following year, I think it was, 1972. We thought we might sell 50 a month. And as soon as we introduced it at a new price point, below the Diablo, we were selling the product from Diablo at a higher price point, it just took off. So we were chasing the whole thing for two or three years, and it eventually got to almost a couple of thousand a month production, but manufacturing wasn't ready, we completely underestimated what the system software needed to do, you know, basically the software guy said, hey, I've got a disk error, crash the operating system, we're done. So a transient read error, boom - that was the end. And we said, well, we probably ought not do it that way, guys. The service people had never seen a disk drive before, and DEC was particularly weak in servicing electromechanical peripherals, they had only serviced digital logic modules and electronic computers in the past. And manufacturing was, put it together and shotgun module swap to make computers work, and now we had a complex electromechanical device, and they just weren't ready, and the engineering, because of the 2315 cartridge had problems, and because, basically it was a very thin team and we rushed it to market, there were a lot of gaps in the engineering.

One interesting story was the first instance of what I call, The Ken Olsen Door Test. Ken would show up in the lab on regular bases, in fact, there's another little story there I'll tell, but, you know, he would come by, and he would we try and open the door on the peripherals. This is the first instance of the opening the door test. So Ken was a big guy, you know, over six feet, probably weighed 280. I remember seeing him wrestle around a snowmobile in one meeting when I was with him at his lake house, and he goes up to the RK05, and it's running, and he pulls on the door, rips all the stops, pops the cartridge off, while it's running. Of course that wasn't so good, the heads were out. So the instructions to all mechanical engineers after that was, the doors have to be Ken-proof. And they were. But anyway, we were bailing the ship, you know, success is the father, sometimes, of real problems, and we were bailing the ship for a couple of years afterwards, making an improved version of the RK05, which we called the RK05J, making a fixed version, double capacity, which we called the RK05F, but it was a hugely successful product for DEC, and it made a lot of money, and certainly put mini computers in a whole new space in a market than they've been before that product had been available.

Gardner: And by this time, you've stopped being a contributing engineer, you're now in management?

Saviers: Well, probably the RK05, I was kind of the project engineer, I was coordinating things, rather than-- I didn't do any design work on the RK05.

Gardner: But no other broader responsibilities.

[0:44:49]
Saviers: Yes, we were doing a second and a third generation of a head per track disk, so we did get this larger capacity, faster head per track disk, called the RS03, and RS04 out, and it was used as a swapping memory for 11s, 15s, and DEC 10s, so we shipped a number of those. And that was the end of head per track drives for DEC. The guy who was one of the technical gurus of DEC, Bill Strecker, had done a lot of work on his PhD thesis about memory hierarchies and he had done that at Carnegie Mellon, and so using the tools he developed, it became very clear that some intermediate-- what was being called fast access memory, at the time, bubbles maybe, or head per track disks or something else, surface acoustic wave devices, the whole panoply of choices. There were electron beam memories, were trying to fit in between core and moving head disks. Well, it turns out-- forget it, the price performance requirements are so extreme that they never happened, and they never will happen.

Gardner: You don't think it's happening today with flash?

Saviers: Well, I think flash is really a replacement for lowest capacity disk drives, as opposed to an intermediate storage level. I mean, when you really look at the performance and cost.

Gardner: You mentioned deficiencies in the 2315 cartridge.

Saviers: Oh yes, I mean, it was wonderful, you could stick your hand-- to carry it, you could stick your hand in the air inlet, put your fingers on the disk, carry it around that way. You could set it down on your desk and the magnetic-- residual magnetism in the hub would pick up all the metal filings, or on your ashtray and the lid that was kind of a cantilevered piece of plastic, wouldn't close properly, so it was open to the environment. So there were all kinds of issues and the other problem was, there were no specs. You know, everybody had their own version of what a 2315 was. IBM didn't release any of the design drawings, so every manufacturer of a 2315 cartridge had slightly different ideas of what the plastic should look like. So, interchange, was ugly. So, IBM recognized the error <laughs> of their ways, and the 5440 was a completely sealed design. Once you took it out of the drive, it was locked closed, and you couldn't stick your fingers in, or dump your ashtray in it, or get metal filings in it, or anything like that.

Gardner: At this point, I think we'll like to segue into your large-disk experience, which, I believe, at least for many years, was mainly as a buyer of large-disk drives.

Saviers: Right.

Gardner: And you mentioned earlier your first experience was Memorex RP01.

Saviers: Correct. We started with Memorex. I was an advisor; it was really a PDP-10 land decision. They decided to go with the Memorex 630 — Memorex had 2311 and 2314 capacities. DEC bought the 2311 to get things rolling, because Memorex wasn't ready yet with the 2314-compatible. But the intent was to ship only the 2314 version. So, they did that, and I got involved, then, with basically buying the disk drives for the company, because I was the guy who was supposed to know something about disk drives, and it was clear that multiple business units, multiple product lines-- we called the "product lines"; not "business units," which were organized by architecture. So, there was a PDP-8 product line, a PDP-11 product line, a PDP-15 product line, a PDP-10 product line. But, more than one product line was going to be using disk drives in the future. So, part of that was regular trips to California, the Mecca of disk drives, of course. There was nothing on the East Coast, except DEC. So, many, many trips to the various suppliers: the alphabet soup of half a dozen or more large drive suppliers. These were either spun out of IBM, or spun out of each other. And one of the first trips to Memorex, I remember, I met Al Shugart for the first time. So, you know, here we are, buying disk drives from him. But I think, DEC was kind of a small player at the time, small company at the time: only a couple of hundred million dollars. And Al probably had no idea who we were. So, I show up in his office, and Al's got his golf clubs propped up against his desk. And he says, "Well, good to meet you. Here's the guy that can tell you everything about what Memorex is doing; I've got an important golf game." So, my first meeting with Al Shugart <laughs>.

[0:50:00]
Gardner: Any idea when that was?

Saviers: Well, it was probably late '71, early '72, maybe, I think, is probably the time frame. I don't know when exactly he went to Memorex, but I remember later, when he lectured at the computer museum, asking him the question. I said, "You know, Al, IBM sued you when you went to Memorex, 'cause you took a bunch of engineers." So, I said, "I've never seen it anywhere in print about how many engineers you took." And, his answer, I think in typical Shugart fashion, was, "Oh, about 300," <laughs> which I think was a little bit of an exaggeration, but was indicative of the motivations at the time to be in the copy-IBM business.

Gardner: In the interests of full disclosure, I have to say that I was at Memorex before Shugart came in April of '69, and I may have been the person that Al handed you off to; I'm not sure, because at the time, I was one of the folks in Memorex' drive side. I know I met you sometime in that time period, but I'm not sure it's then. And, by my count, it's something on the order of 35 engineers came from IBM to Memorex. Although the rumor of two to three hundred is well-promulgated in the industry, the facts support a much lower number.

Saviers: Right, I believe that; right.

Gardner: And they were some really good guys. So, you were buying the RP02. Any recollections of the RP02? By the way, this RP02 was the Memorex 660 Model 1.

Saviers: 660 Model 1; okay. 630 was the 2311 version.

Gardner: And, DEC did announce the RP01. Whether they ever shipped one or not, that was the Memorex 630 model 1.

Saviers: They all went into the DEC data centers, instead.

Gardner: And I believe the engineer on the other side was John Green. I think John Green and I worked together to get that product on the PDP-10. But, that's a whole different story. So, any RP02 recollections?

Saviers: Well, there's the RP02 one, and then there's the RP02 version two. You know, it came back at RPR02. That would jump forward several years, but maybe it's worth telling, because it was an interesting context to Memorex. We met Larry Spitters for the first time. And, "Larry," I said, "we're going to buy these drives. We're going to be a big customer. Please come and talk to DEC senior management about why you're a credible supplier." So, Larry gave his marvelous song and dance. I mean, anybody, at the end of the meeting of an hour of Larry Spitters' talking was reaching for his wallet: "How do I buy Memorex stock?" you know. So, he was quite a salesman. So, anyway, I think we had reasonable success with the RP02. I don't remember any disaster stories coming out of it, anyway, but it did get the PDP-10 problem solved. And it was sufficiently interesting that there were some initial attempts to build a controller for the PDP-11, but they went into hibernation for a while, until the RPR02 happened, which was the same product, reconstituted about six years later, I think it was; probably late '70s, '77, '78 timeframe. The small-business product line at DEC needed bigger capacities. We had nice

PDP-11 computers; but, for the data-processing market, not enough storage capacity. And there really wasn't anything out in the market that fit their price point. We knew that Memorex had a warehouse-full of off-lease 660-1s, and they were in serious financial difficult at the time. There was a real question as to whether the company was going to survive. And, if I remember right, this was one of the Arnold Cooley sales interactions: Arnold had an interesting history with DEC, selling us large disk drives. But, anyway, we went to Memorex and said, "Look, you've got these things," and, at this point, the controller for the PDP-11 had been resurrected and made to work for the IBM interface, or the modified IBM interface that was on the RP02. So, we said, "Look, you know, you've got these things in your warehouse. You've fully depreciated them. Why don't you sell them to us, and we'll pay you about two thousand dollars, each."
[0:55:00]
So, they gulped, but, hey; it was free money, in a way, for them. So, they refurbished them, and we probably bought several thousand of them for the small-business market. So, the RP02 had two separate lives: its original and its reincarnated life, something like six or seven years later.

Gardner: I think I found one RPR02 out there, still sort of working. I'm trying to acquire it for the museum. I actually ran Manufacturing for Memorex, at the time; I was building them for you, and we were refurbishing them, and making them into DEC drives, an RP02. But you didn't buy from Memorex for the next project.

Saviers: Well, yes. I mean, if we go back now to the RP02, version one, this new company, called ISS, was started. And I guess-- was it the Dirty Dozen? Was that what they were called? Was that ISS?

Gardner: Yes. The Dirty Dozen founded ISS.

Saviers: So, there's another generation of ex-IBMers starts another company[3], called ISS, Information Storage Systems, I think it stood for. Anyway, these guys had a really good-looking RP02 equivalent, as well as a double-density version RP03. So, I'm not sure where Memorex or other people were in this game, but it was real clear that ISS had the product, had the best product. So, we decided to go with ISS. And, at the same time, we were on another track internal to DEC, designing a storage-independent, storage-technology independent interface, called the MASSBUS. But those soon converged with the RP family in the next generation. So, the RP02, for Memorex, because the RP03 and 4, from ISS, single- and double-density, 2314. And then, we moved to-- back up; I got a couple of nomenclatures mixed up there.

Gardner: Do you recall anything about the RP03? How did it work out for you guys?

[3] Editor's Note: ISS was the first spin out from IBM San Jose; it was formed before Memorex started its disk drive development.

Saviers: I think it worked out very well; I mean, I think it was a good-quality product. And that led us to the first of the 3330-generation products from ISS, which was the RP04.

Gardner: As an aside, the other players for that RP03 generation were, I believe, CalComp or Century Data, and possibly CDC by then had a servo disk drive, and there was at least one other player who could do the double-density 2314s. Memorex couldn't compete because the product had an open loop positioning mechanism. Everybody else used some sort of closed-loop external transducer system. So, Memorex was technically blocked.

Saviers: But, a generation behind. Yes. That's good that you reminded me. Memorex had a mechanical detenting system for positioning the tracks, and ISS came along with an optical, which we had been doing on the RK05s. We said really, this was a better way to go, in terms of reliability and durability, maintainability of the product.

Gardner: So, that then comes to the RP04, which is an IBM 3330-class, 100-megabyte disk drive. But, it had a very different interface.

Saviers: Yes. Well, we, with the last generation of head-per-track drives-- because, at DEC, we had, as I mentioned earlier, 12-bit, 18-bit, 16-bit, and 36-bit computers. And every one of those computers had more than one operating systems; in some cases, four or five operating systems. There's a huge combinatorial problem for us to manage, like, 35 different software product releases when you wanted to ship a disk drive. Well, that was kind of a total mess. So, we had to simplify it somehow, and simplifying at architecture cut-points-- you know, you take the entire system print set, and you take a scissors, and go through, and say, "There's an interface." For the software, you do the same thing. So, we decided that a standardized storage interface that was more or less technology-independent, so that we didn't have to worry about the particular vagaries of the disk technology or the tape technology, or, even in the case of memory technology, we created something we called a MASSBUS, which had an 16/18-bit data path, so it could either handle 16-bit, 18-bit computers, or 36-bit computers, and then a separate control path of 16 bits.
[1:00:08]
So, it was a big, fat cable. And we wanted to solve the problem of being able to string these things across the computer-room floors, so we needed to have 100 feet of daisy chained cable. And, so the MASSBUS was invented, and was decided to be the standardized interface for computer-storage peripherals. So, we built a tape drive that way, built a head per track disk drive that way, and we built many generations after that of large disk drives that way. And it was very simple on the computer end, because the parallel interface is a very simple computer interface, which was another thing: prior to MASSBUS every computer that came out had to have another adapter for this or that particular disk drive. And, in the case of a MASSBUS, it was a very simple, one-board computer interface, a one-board MASSBUS controller. So, we did that. And, some other business objectives at the same time: when we looked at 3330-generation technology, we said, "This is an opportunity to bring large-drive technology to

the minicomputer market." It had been there on the PDP-10, with the RP02 and 03. But it had not been there on the PDP-11, or the PDP-8, or the PDP-15. So, the interesting story was: we weren't sure about doing this product, and we had a customer in Holland that said, "Hey"-- the 11/45 had just hit the market, which was the first super minicomputer. It really had lots of processing capability, a bigger memory-- "and I can get the contract for automating all of the Dutch Social Security System. There'll be a dozen installations; each installation will take three or four large disk drives of the 3330 class. Would you build that product for me?" So, we all flew over to Holland; there was about five of us: Demetrios Lignos, who became the RP04 product engineering manager; Dick Clayton, who was the project manager for the PDP-11/45; myself; and one or two other folks, I think. We spent a week with this guy, testing the waters and seeing what he wanted, in the way of a solution. He was deep technically, modifying the operating systems, etc. And, we thought maybe we'd sell 80 of these things. So, we said, "Okay, it'll cost us about a million bucks to do the interface to the ISS drive. We'll get our money back, anyway. This is a shot in the dark, in terms of the future of large disks on super-minis." So, we did it, and it was a huge, runaway success. And set the stage for some really large purchases of big disk drives by DEC in the future. But that was the beginning of the MASSBUS family of large, IBM-compatible disk drives.

Gardner: But your business target was the low end of the IBM product range?

Saviers: Yes. What we looked at immediately was, you know, here's a 3380 controller sold by IBM for 100 thousand dollars, plus. So, we're going to sell this on a 50 thousand-dollar PDP-11/45? Well, that didn't quite fly. So, our objective was to be able to sell one through four drives cheaper than IBM. And the way the MASSBUS was developed and designed, we could do that. When you get about four drives, we started to cost more than IBM on a cost-per-megabyte basis. So, it was an interesting set of trade-offs, optimized for where we were in the systems market at the time.

Gardner: But at the same time IBM went to what they called "integrated file attachments," which had a similar, arguably, bus to the MASSBUS. IBM called it the "drive-controller interface."

Saviers: Don't forget that IBM had one family of computers, called 370s. And, only on one or two models of the 370 did they do that. And they didn't have four different computer product lines, with new models coming out every few months and the numerous operating systems for each. So, it was a very different environment.

Gardner: But, if you're targeting the low end of the IBM 370 line with this product, where they were competing not with a big control unit, but with a small controller, and they were sharing the controller with multiple drives, it seems like you might not have been competitive, even at the one- or two-drive level.

Saviers: Well, the question in the customer's mind was dollars per megabyte and purchase price. And, in that time frame, we never thought we were competing with IBM. I mean, who was IBM?

[1:05:00]
They were making these big computers. We make minicomputers and super-minis. So, it wasn't until much later, into the mid-1980s, almost 10-plus years later, that DEC started to think more seriously about competing with IBM.

Gardner: So, it was more likely the IBM System/3, or one of what IBM would have called their small systems, but most people would call minicomputers, that you were targeting.

Saviers: You know, trying to reflect back to that conversation with the customer, it was, like, "I want to buy a DEC computer; it just doesn't do everything I want. So will you do it?" It wasn't a question of, "I'm going to buy somebody else's computer. I might have to buy somebody else's computer, but I really don't want to buy anybody else's computer. I'm a PDP-11 guy." So, architectural mindshare, what we used to call "architectural mindshare," was gaining a lot of traction because of the position that DEC had, and you have to think back, maybe another environmental issue here. DEC created independent software vendors: ISV's, they were called for many years. But, before DEC came on the marketplace, you didn't buy your software from Oracle; there was no Oracle. You didn't buy something from SAP; there was no SAP. You bought it from IBM. Now, DEC sold minicomputers with huge discount structures. We used to say the Maynard list price, which was the published price, was the price that every customer was guaranteed not to pay, because the average discount was 25 percent, across the company. And, for really big OEM's, it could get to 35 percent. And what would happen was: the software suppliers would sell at MLP <Maynard List Price>, and the spread was how people paid for software. There was no software list price. You bought an MRP system-- a friend of mine had an MRP company. If you bought the MRP system, you paid list (MLP) price. He didn't have a price list for software. He just paid for the DEC computer, and the discount that he got made him the money. Of course, that later changed dramatically, as you got to 1980 and beyond. Real software companies started to appear, and separately DEC also blew it with the ISV's. And, another terrible story about DEC. But, that whole minicomputer discount structure is what caused independent software companies to exist.

Gardner: That's actually very interesting.

Saviers: Not very much to do with storage, though, is it <laughs>?

Gardner: Well, it had a lot to do with storage, because it sort of drove the market.

Saviers: Drove the market; that's right. So, this guy in Holland: he probably-- you know, he was going to make it up in service, too, so he had a big discount from DEC. And he had, obviously, an applications specialization knowledge that he was going to do for the Dutch government, and an on-going service revenue. But he was prototypical of an ISV at that point.

Gardner: Also fairly interesting at that time about MASSBUS was the idea that-- device independence was one thing; that was happening in the IBM world, too. But the idea of sharing tape and disk on the same interface was unusual. Today, of course, SCSI does everything.

Saviers: Yes. At the time, we built a head per track drive, RS03/4. We had the RP04. We had two different tape drives, TU16 and TU77. And the tape drives: we did them with the master/slave, so there was a master MASSBUS, and then a string of tape drives off of that that had some tape-drive unique interface on it, for cost-effectiveness. But the drives: all the disk drives had the MASSBUS embedded, one per drive.

Gardner: If it was economical with tape, why wasn't it economical with disk?

Saviers: Well, you could have done it that way, but we decided not to, because we thought we'd sell a lot of one- and two-drive systems. And, for tape, you almost always need more than one. And I think the tape MASSBUS controller was enormously more complicated, more expensive, than the disk MASSBUS controller.

Gardner: I think you always needed at least three tape drives.

Saviers: Yes, something like that.

Gardner: It's hard to do anything without three, and most people have more. I'm, as a student of the industry's history interested in different directions taken by IBM and DEC.

[1:10:00]
Saviers: In retrospect, maybe there were better alternatives, but it worked. And it also provided something else that we wanted, which was-- and became a main theme, over many years-- which was to beat up on the guys that copied us. So, I was in federal court many times, suing for intellectual property protection. And, probably the MASSBUS was one of the first instances of those. And that was one of our strategies, and we were the small guy; we didn't have to worry too much about anti-trust, like IBM did. So, we patented everything we could to keep the plug-compatible guys at bay.

Gardner: That's another interesting aspect of the industry. IBM, because of prior issues, had to license everything. You guys didn't, but you got big enough, and then you started to see people enter the compatible business. I think it's a thread I'd like to push off 'til later. So, you come out with the MASSBUS. RP04 is the first instance. Business plan to sell eighty, and you sold thousands.

Saviers: I don't remember the number for it, but it was certainly well over 1000, and really set the stage for going at this in a much more serious way. The MASSBUS interface was pretty well engineered. Demetrios Lignos became the project manager for subsequent generations, until he went off to a different job in storage. But, the next step was the 3330. And, of course, this was probably the really first big shootout we had, as a big buyer of disk drives, in fact, outside of the plug-compatible guys selling them on IBM systems, we were the largest disk purchaser in the world. So, these contracts became multi-hundred million-dollar deals, when you looked at the life cycle. One of my guys was a product manager for large drives, Kevin Smith, was a wizard at putting the numbers together, and extraordinarily accurate at predicting how many of these things we'd need over the product life. So, I think when we got to the RP05 and the RP06, we knew that this was somewhere between 300 million- and 500 million-dollar contract value, when you add it up.

Gardner: So, just a minor correction: I think the RP04 was a 3330.

Saviers: I'm sorry; right. Second generation.

Gardner: And the double-density 3330s, the 200-megabyte disk pack, disk drive is the big contract.

Saviers: Right. So, we had this big success with ISS with the first generation. They were behind the curve. With Memorex; sorry. They were behind the curve, in terms of the second generation, the double-density 3330s. So, we're out in the market, and go with ISS, and have this big shootout among all the suppliers who are building double-density 3330s. And we're extraordinarily impressed with the Memorex design. I mean, over the years that we built moving-head drives; in fact, over the whole generation of RP disk drives, it was by far the best product we ever bought: the (RP0) five and the six. And I was just kind of blown away by the capability of the engineering team at Memorex, too. I think it was Marco Padalino. Is that right name?

Gardner: Marco Padalino was the head of engineering for a long time.

Saviers: Yes. So, when he explained what they did and why they did it, and how the product was engineered, the proof was in the pudding, too. And the proof of the pudding was in the eating, because we had great success with the product. I should tell, also, a second generation of Ken Olsen storage stories. I mentioned earlier that this big guy would come around, and now and then, we'd know he'd been in the lab, 'cause Monday morning, we'd come-- and he was a wanderer around the mill on the weekend. You'd be in the mill on a Saturday or Sunday, and run into Ken. So, anyway, we'd come in on Monday morning, and the door, the sliding door on the RP06 is ripped off. <laughs> Epoxied on some heavy-duty slides. Well, we knew that Ken had been in the lab, so there's not much we could do about that one. We said, "Well, we'll just live with the design as it is."

Gardner: That was safety glass, so if he ripped it off, he had to rip the glass off the door.

Saviers: Oh, yes. He broke the adhesive right off the rails <laughs>.

Gardner: Yes, an impressive product-test person.

Saviers: So, that was kind of the high-water mark, in terms of quality of the product. And really was a big deal, in terms of the purchasing and the contract negotiation. And, sometimes, they'd get to be pretty intense.

[1:15:10]
Gardner: Again, in the interest of full disclosure, I'm pretty sure I was running manufacturing at that time, making those RP06's. We called them the 667, I remember. I do recall visiting Colorado Springs, and giving your CFO there a gold Cross pen, because he was signing checks every two weeks, because Memorex needed the cash. We had to be a little careful about spiff'ing, but give him a gold Cross pen so he could sign the checks.

Saviers: Yes, I think it was in one of those generations-- I'm not sure which product it was, but it was, I'm pretty sure, ISS. It might have been for the RP02 and 3. DEC was growing so fast, I'm going to have to look at my notes, here. Let's see: '71 was 150 million, '72 was 190 million, '73 was 270 million, '74 was 425 million, '75 was 535 million. So, the company was the darling of Wall Street of the era. And we completely lost control of accounts payable. It was a total disaster. My wife was working in accounting, and they called up all of the accountants from all over the company, and put them into accounts payable, and said, "Call up the supplier, ask them what we owe them, and write a check. Don't even ask for documentation; we need the parts." And, at this point, it was ISS, I'm fairly sure, we said, "We're beating them. We need more drives, we need more drives." And they said, "Well, you haven't paid." <laughs> Well, that's a problem. "How much do we owe you?" "Well, you owe us 11 million dollars right now." Oh. So, I go over the CFO, and I said, "I need 11 million dollars." And he said, "Why?" I said, "Well, production's going to stop shipping computer systems unless I get this check to them." "Okay." So, after dealing with the bureaucracy-- the guy actually cutting the check refused to give me the check, and back and forth to the CFO, I have an 11 million-dollar check in my hand. I fly out to ISS, and give them the check so they would keep shipping disk drives to us. It was the first time, and the only time, by the way, that I had an 11 million-dollar check in my hand.

Gardner: Those are the fun problems. So, you're shipping lots of RP06's. It comes to the end of its life. And you're in another procurement cycle.

Saviers: Yes, well, that whole large-disk drive thing was moving up, so there was a demand for something not quite as big, and a demand for something bigger. So, at this point, CDC had really been aggravated with us, because I think there had been three or four "buy-stuff" shootouts, and they' lost every time. So, they said, "You guys don't like us, or something."

Gardner: I have to ask the question now: is that because they'd never hired Arnold Cooley?

Saviers: No, of course not, but he was good. Well, I'm not sure if he was a salesman at Memorex, but he was the salesman at ISS, I think.

Gardner: Our observation from the outside was that the DEC contract seemed to move with Arnold.

Saviers: It did.

Gardner: So, we hired Arnold from ISS.

Saviers: Okay. Yes, well, Arnold was just a good salesman. He did a good job of representing us back to the factory and the marketing people, and the executives making the business decision. I mean, there was nothing unique about Arnold, other than just being a top-notch guy who listened, and worked hard.

Gardner: I think he communicated well back to his development and manufacturing folks. Arnold was a very good advocate for the person he was trying to sell to.

Saviers: And he was successful in that regard. Of course, we wanted what we wanted, and it was a big number. And we thought we should get it for a $500 million contract. And, of course, the suppliers agreed <laughs>.

Gardner: So, CDC was unhappy with you.

Saviers: Oh, extremely unhappy. They were at the point where it was, like, "Don't even bother coming here any more." Because we'd been up to Normandale (MN) many times, and they'd give us the dog-and-pony show, and give us the whole menu of stuff, and yada-yada. So, anyway, time comes that we want larger and smaller drives, and we got into a contract with CDC for the 9762, the 80-megabyte SMD; and for the bigger one, the RM05, the 9766, 300 megabytes, both removables.
[1:20:05]
So, those got adapted. Those got MASSBUS interfaces on them, and were shipped as large drives. They were reasonably successful. They had massaged the head-disk interface, and it turned out to be a little bit of a problem, so we had lots of head-crashes with those products; but, otherwise, they were pretty decent.

Gardner: By the way, just to add to the record, Marco Padalino probably sold you and then left Memorex. Because the guy who actually executed the 677 program was Tony LaPine - he had a really good team of folks, and was coming off of a pretty good base.

Saviers: Yes. Sometimes, you'd get the names mixed up as to when they were where. I certainly knew Tony better afterwards, in the small-disk business <laughs>. What a thrill <laughs>.

Gardner: So, was the 80-megabyte and the 300-megabyte basically, actually the 3330 mod-11 technology?

Saviers: Closer flying height, different head.

Gardner: Yes, but it was basically the 200-megabyte disk pack, 50 percent higher capacity for a short-stack version of the same technology to get the price point down, and reasonably successful products. But the industry is moving on.

Saviers: Yes, unfortunately. The next one was the worst drive ever bought, the RP07, Winchester. So, I'm not sure what the exact IBM sequence was, but I think the 300 and 600 were the two single-density, double-density Winchesters, or something like that, anyway.

Gardner: I think the IBM 3350 was a 317-megabyte, fixed-disk disk drive. IBM never did double density. The rest of the industry-- Memorex, ISS as a part of Univac, Storage Tech-- did double-density 3350s, and they were all 635 megabytes; and, in DEC's format, somewhat less. I think what we're talking about now is the double-density, 3350 class of drive.

Saviers: Correct, right. And we weren't interested in a 300; we already had 300. So, we were interested in the 600. So, we engaged with ISS on a little bit different basis. We said, "Look, let's put the MASSBUS stuff inside the drive." So, we said, "We'll provide all the drawings and how to do it, and you just have to package it so it goes inside. There's enough space inside the ISS drive -- well, I shouldn't sully the name of ISS with this product. You know, the Unisys <laughs> RP07. Well, it turned out to be a real mess, because they were very late, which is a little bit understandable, in terms of building a new interface. But, the product just wasn't reliable: it had head-crashes, and it had manufacturing problems, and they just couldn't make it in any kind of quantity reliably. And, to do the deal, we had to contract for 5,000 of them. And, it was a very painful 5,000, and we had a big party the day that number 5,000 arrived.

Gardner: But you did buy 5,000.

Saviers: We did buy 5,000; hm-hmm. And that was, what? $11,000 a pop?

Gardner: Probably. I think the OEM price of a high end disk drive was between eight and twelve thousand for many years.

Saviers: Right. On all the other ones, we were paying seven, eight thousand dollars; but on this one, it went up to eleven, with the interface built in. But it was also a more expensive drive to build, I think, too.

Gardner: I seem to recall that was almost the standard price for a high end machine was $8,000, and you just got more and more capacity for the same price. And Digital--

Saviers: Probably got the best price.

Gardner: Always got the best price. And the rest of the guys got the same price three months later. But DEC was always the favorite customer. Any recollections of what the problems were that Unisys had with the RP07?

Saviers: Unreliable HDA's. I mean, I think there were a number of them in that. Every few months, there's a new one.

Gardner: That's actually the beginning in the large-disk era of serious unreliability problems.
[1:25:00]
Storage Tech famously had a recall for their double-density 3350, so much so that it made the 10K. They had to disclose the financial impact of the recall, and I was told, in another interview, that at one time, they had more disk drives in trucks coming back to the factory for repair - it was mainly head-crash, HDI problems. And then, of course, the next generation, which was the last generation of 14-inch. This was the end of 14-inch. I think you guys were maybe lucky, because you never did buy a 3380-class.

Saviers: No, we didn't. No, we looked at the 3380, and said "whew". This was the beginning, also, of looking at IBM stuff, and saying, "I guess they can justify that on the basis of system performance and system availability; but boy, that doesn't look like the right way to build a disk drive." So, we shied away from that completely. And, at the same time, we knew these products were generating more and more revenue for DEC, and we also saw that the plug compatible suppliers were generally unstable. They would come and go. They'd get merged, acquisitioned, go out of business, et cetera. And there was sufficient quality problems that we said, "Well, we should look at getting into this business, ourselves." So, that's what we decided to do. It was the beginnings of an advanced development group. I hired a guy named Mike Riggle, who was former Control Data executive. And he started down the path of a clean-sheet approach to the whole problem: the subsystem problem, partitioning, the architecture, and what the drive families ought to look like. The objective was: we'd been buying products from guys following IBM, and now DEC was getting to the point-- we're now into the late '70s with this. DEC said, "Hey, we have to get up there on IBM's aerial-density curve if we're going to really be in this business long-term, and sell larger minicomputers, et cetera." The VAX had come on the market in 1978. And with that, came a lot more processing power, and memory capacity. And our customers were saying, "Well, we want to take these into the edge of the glass house; not really into the glass house yet, but we want to take them into the edge of the glass house." We were the largest computer supplier to Boeing, largest computer supplier to ATT & Schlumberger. So, we had some pretty significant customers.

Gardner: Do you recall when you hired Mike?

Saviers: It must have been '75 or '76, in that time.

Gardner: Let me get my crib sheet here.

Saviers: Because by '77, we'd laid out some of the architectural elements pretty carefully. And actually started some prototyping of some of the subsystem concepts that we wanted to do.

Gardner: This sort of takes us back, now, to the small track, which then goes into the big track.

Saviers: That's right. So, the genealogy of being in the disk business, of making our own disk drives, starts with these head-per-track files in the '68, '69, '70 timeframe; moves into the cartridge drives in the '71-'72 timeframe; gets a larger cartridge drive, called the RK06, 07 in the mid-seventies timeframe. And, at the same time, we've got all these buy-outs going on: Large drives, and then maybe the smallest large drive being the RM03, the 9762. So, we decide that we want to have of a clean-sheet architecture, it would multi-generational. The VMS guys, Dave Cutler and VMS guys, were working on a technology called clustering, so maybe let me back up now, and talk about Gordon Bell and the overall corporate strategy that Gordon had. We called it the "Big E." Basically, it was three tiers of computing. The top tier was the work-station, desktop. This was before the personal computer hits the scene, but it was the personal computer, if you will, which could be a PDP-11, or small VAX, or a smart terminal. And then, the mid-tier was the VAX servers. So, this is circa '78-'79. And the biggest tier was very fast networking the biggest VAX's together into clusters.
[1:30:00]
And the dream with VAX is becoming the 370 of DEC — could we build a computer that had a 1000-to-one cost range and the appropriate performance range with it? The big question was: could you take an architecture as complicated as VAX, and put it into silicon someday? You know, it's a nonsense question today, but back then, it wasn't such a nonsense question-- and then, build big computers. And the VAX-11/780 was in development, shipped in 1978. We had the then high-end model in place, and started thinking about how to build MicroVAX. They had built an LSI-11, so the 11 architecture had been

reduced to single-chip silicon, or multi-chip silicon. So, we knew that was possible. So, the idea was if that happened, then PDP eights, and elevens, and tens would all go away. And DEC would be the VAX company. One computer line, and get rid of this huge combinatorial problem, and there'd be one operating system, called VMS. Get rid of this huge combinatorial problem of operating systems. So, it was a great strategy. And Gordon was very explicit: we're not going to build mainframe computers. That turned out to change after Gordon left, because Ken wanted to build mainframes, which was a terrible idea. But, anyway, we weren't going to build mainframes; we were going to solve the problem of mainframes with clustered computers, clustered VAX's. And the two normal technology curves; you know, there would be increasing performance, and there would be the decreasing price curve. So, we'd go both ways, getting a bigger and bigger slice of the market at the same time. So, anyway, the new Digital Storage Architecture or DSA had to span all of that: had to be able to handle these clustered computers all the way down to a single PDP-11, which was still going to be around for a while, or superseded by a small VAX, a MicroVAX. So, that's what we designed: called the Digital Storage Architecture. It encompassed Q-bus controller, or Unibus controller, a BI controller, and a cluster controller, called HSC. It had serial interconnect between the storage device and the controller. That could be a tape drive, it could be electronic memory, it could be a disk drive, it could be a head-per-track disk drive. And it was all three, over time. And, on the other end, in terms of the Q-bus, Unibus, BI bus had the computer bus connection. And, in the case of the HSC, had the cluster interconnect, which allowed-- initially, it allowed, I think, it was all a matter of VMS software capability-- allowed, like, six things, six entities in a cluster. Any mix of any VAX and HSC's. And then later, I think it got to 24, over time. You could have one VAX, and twenty-three HSC's, or twenty-three VAX's and one HSC, or any mix in between.

Gardner: I'm a little confused; I'm not a student of DEC. Why wasn't MASSBUS promulgated into this space?

Saviers: Too slow.

Gardner: Too slow, even though it was parallel?

Saviers: Yes, it was limited to about-- with a daisy-chain, long busses, 100-foot busses, it was limited to about 3 megabytes per second. Too expensive. Cables were way too expensive. Not ESD bombproof. You know, we're talking about IBM, fire hose-sized cables <laughs> -- 20 bucks a foot and $80 connectors.

Gardner: Actually, there was an industry at one time of going into those glass houses, ripping the floors up, pulling the old cables out, and then reselling them. People would offer customized cables for your glass house, in return for being able to take all of the cables dropped below the floor, because typically when you replace a peripheral, you just dropped its cables.

Saviers: Right. Then there were a whole bunch of RAS issues: Reliability, Availability, Serviceability. We had not paid, to that point in DEC, much attention to... let's call it IBM system-level RAS features. But we said, "Hmm. We're going to have people out there with 100 spindles on our systems, and four VAX's, and so we need to pay some attention to that problem," because it was-- we'd taken the full 3380 documentation and looked through it and said, "This is really smart from that point of view."

Gardner: RAS is Reliability, Availability, Serviceability, at least in IBM's terms?

[1:35:00]
Saviers: Right. Same thing in DEC's terms. Right.

Gardner: This is a revolution in architecture.

Saviers: Right.

Gardner: Massbus is pretty interesting, but not adequate. I'm still a little confused. DSA is Digital Storage Architecture?

Saviers: Well, the Digital Storage Architecture had several major elements to it: something called Mass Storage Control Protocol, which is a SCSI-like-- probably SCSI borrowed some things from it, actually-- SCSI-like in today's SCSI, but was pretty revolutionary in terms of a packet control interface for how to control peripherals of all kinds, okay-- tapes, disks, electronic storage, CCDs or bubbles if they appeared, a standard set of wires called the "Storage Interconnect," which was the first time a serial data path to and from the storage device was developed, and there's a couple reasons for that which I can go into, and then an on-disk format for disk drives that said, "Here's how we do bad block revectoring, how we do error control, where the spares are, et cetera." It had a lot of flexibility in it, but it-- you could look in a certain place on the disk drive and find out all the information that an operating system driver needed to understand how that disk drive was built. So people weren't doing that at that time. That was a completely new idea. And now every drive from the factory comes with all that information written on it. But-- and part of it was, how do you build the zap-proof-- you know, lightning strikes a drive out there on the string, so it was a radial interconnect that was transformer-coupled that had to meet all the new FCC requirements, yada yada yada. There was a really long list of current requirements and looking-forward requirements that the Digital Storage Architecture encompassed, so...

Gardner: But it's a lot more than storage, right? Doesn't it become the method of hierarchically interconnecting computers in storage, in almost anything?

Saviers: Well, clustering was the beginning of that. Yes, DEC clustering was really the first of block-level sharing of storage across multiple computers. IBM had multiplexing capability, but not at the level of what DEC was going to do, and also not at the level of modularity. You could mix an 11/730, an 11/750, an 11/780 on a cluster. You could have multiple kinds of VAX's on a cluster.

Gardner: Which came first, the chicken or the egg? I'm quite interested and quite confused. Did the clustering come first, and then the storage folks adopted that as its architecture?

Saviers: Well, no, we knew that we needed to build a box, a very high-performance box like the 3880 IBM controller for the 3380. We had to build a very high-performance box that could handle multiple parallel data transfers simultaneously, and fire-hose one data stream to a computer. We knew that we could not achieve the performance levels we wanted to achieve with what existed. And then there were other requirements, too, like it would be nice if this thing kept track of bad blocks. It would nice if this thing did RAID, which it did. It'd be nice if this thing could copy disks to tape, which it did. So it had a lot of functionality in the performance domain, as well as the pure functionality that didn't exist in any other controller we built. It was very contentious. Very, very contentious. I had some really big fights with Gordon Bell about this, but we prevailed, and we prototyped it. He said, "Okay, if you're going to go build that, show me that you know what you're doing." So we actually built it out of a bunch of PDP-11s and buses and so on, all connected together, and prototyped it, and said, "Hey, this'll do it." And we showed the architecture's right, because it had a control processor which was a fast single-chip PDP-11, and then the HSC had blades that you could add so you could start with four drives and add more drives over time. You could add cache memory in it, and we never quite ever did that, but it was designed to be able to do that. And so it had kind of all those things which we think about today that storage controllers ought to have, and this was architected in 1978 and '79, and then we moved it to Colorado Springs to build it. The design team migrated to Colorado Springs.

[1:40:05]

Gardner: Before we broke for lunch, I was asking for help in understanding the relationship between the Clustering Architecture and the Digital Storage Architecture. Perhaps at this point you can help me understand.

Saviers: Well, yes. Of course, the basic problem in building larger and larger computers is where do you cut things and put interfaces between this and that so that they scale over performance, capability and physical space. You end up with computer systems that take up more than one floor in a building. So the idea with the CI interconnect was to be able to share computers across multiple generations of computers, across multiple types of computers. Of course, the computer in this case was always a VAX, but it could be a different performance/size VAX, and storage and perhaps other things like complicated networking things. So the Cluster Interconnect was-- the idea was, how can we make the investment for a customer scalable? In other words, they don't need to stop buying VAX 11/780s and buy DEC's bigger thing. They could just add to. And this way, you can scale performance very incrementally, and not have a big switching cost for the customer. And that was Gordon’s strategy of being able to cover at least the bottom end of the mainframe market with midsized super-minis, or even bigger than super-minis, as they really became over the years. So CI was the computer-to-computer and computer-to-storage interconnect, and it hooked together -- <laughs> the bus was actually wires coming into a resistor that was the central node, and they all came in radially from various things, and the HSC-- Hierarchical Storage Controller-- was the box that controlled all the storage and talked to all the computers, and sent the data to whichever computer it was appropriate to send it to. And there was file-level interlocking and block-level interlocking done by VMS, the operating system, so that you could share things.

Gardner: So the CI and HSC co-developed with DSA-- the relationship between Digital Storage Architecture and its set of interfaces and HSC. Now, in time, you implemented DSA directly connected to computers through the SI-- Storage Interconnect.

Saviers: Well, the Storage Interconnect-- so if you have this-- the controller and-- the controllers-- the big one, the HSC, talked via the Cluster Interconnect. The smaller ones-- the Unibus, the Q-bus, and the BI bus-- were one- or two-board controllers that plugged right into the computer, and all those controllers talked Storage Interconnect, which was the serial bus that talked to the storage device itself. And the storage device could be any storage device. We had, in fact, in the very first iterations, we had five-and-a-quarter drives with the SI interface. We had 14-inch Winchester drives with the interface. We had a 14-inch removable drive with the interface. Later, we added a solid-state disk with the same SI interface. We had tape drives with the SI interface. We even ended up OEMing IBM’s biggest, fastest tape drive at one point-- put an SI interface on it. So the HSC then became the storage optimizer, and it could handle multiple data transfers simultaneously. It was a blade architecture. They had a PDP-11 microprocessor to do all the policy work, but then it had blades, and each blade could have a data transfer going on simultaneously on one of its multiple ports. So at one point, I think we had the capability of putting 24 disk drives across 6 blades, so we could have 6 data transfers operating simultaneously.

Gardner: Let’s set aside CI and the HSC, and go back earlier in time to DSA and the first implementations of DSA, and how that came. Again, Massbus is sort of the prior generation to what became DSA and the Storage Interconnect. Did I get that correctly?

Saviers: Correct

[1:45:00]
Gardner: Where did that come from?

Saviers: Digital Storage Architecture was a clean sheet of paper, and so the question was, how do we scale, how do we provide the RAS features that we knew we needed for the IBM market? How do we provide compliance to all the new standards that were coming into the computer industry-- FCC and yada yada? How did we handle all the separation distances of bigger systems, and how did we build an architecture that could handle multiple generations of multiple kinds of technology-- disks, tapes, et cetera... electronic storage? So that was the overall arch of DSA. So there was-- part of it was called Mass Storage Control Protocol, which talked back and forth to the peripheral device; there were the serial wires, the Storage Interconnect; and then for disk drives, there was the on-disk format. How could a smart controller or an operating system figure out this disk drive you plugged in it had never seen before? Where were the tracks? How did the bad block vectoring work, et cetera? That was all actually done in the controllers, so there wasn't much the operating system needed to do anymore in there. But-- so it was an encompassing architecture laid out with prospects of a 10-year life, in terms of where we thought data rates and access times and things like that would go.

Gardner: Was there an architect?

Saviers: It was a group. Barry Rubinson was one of the prime leaders. He recruited Richie Lary, who was one of the brilliant guys of DEC; had more patents than anybody else in DEC. Bob Bean was one of the key software guys. Bob Supnik-- I hired Bob Supnik, who later became Mr. Alpha. Bob was involved in the early days of it, with storage. The guy who ran the project in advanced development was a guy I hired from CDC, Mike Riggle, who went on to become a fellow of the IEEE and was a vice president of DEC, and by far and away ran the best research advanced development shop in the company.

Gardner: So it started in?

Saviers: Seventy-six, I think, was probably the beginnings of it, and I think we prototyped the idea, because Gordon Bell was very dubious about this thing, and he was even more angry when we recruited Richie Lary to go to Colorado... <laughs> because he was one of his prime go-to guys. Barry Rubinson was a master of recruiting the best technical talent. So we built a prototype because Gordon wanted to say, "Show to me you know what you're doing-- if the buses are the right performance, that the architecture is right." So we did that and laid out some of the software concepts, and then moved it to Colorado for the product implementation.

Gardner: Part of the driving factor, also, in Massbus was to shield the technology so that you could go through many generations, and I believe you said part of the motivation was to allow you to get a step ahead of IBM, or at least get on IBM's curve.

Saviers: Well, yes.

Gardner: It's not as good a comparison.

Saviers: Yes. I would say Mike Riggle, his-- the agenda that we had was we knew that after this many successive generations of RP0x that we needed to be on the IBM areal density line by hook or by crook. Somehow we had to get there, and we knew that it would be extremely unlikely that in the near term we're going to be able to make disks or make heads as good as IBM, considering the hundreds of millions of dollars they're putting into that every year. So Mike Riggle said, "Look, I think I can do this, but with some systems architectural work, as opposed to raw component excellence." So the objective was, how do we push densities up to where IBM is, but when you got basically not-so-good heads and disks? Well, the error rate goes to heck, so the idea was we got a 10-to-the-minus-8 error rate plan for what we're going to be able to get, in terms of heads and media and channel. Now, how do we make that 10 to the 13^{th} for storage? —DEC was an early user of Reed-Solomon codes in storage anywhere[4], including optical, and so it was a key part of it. And then the channel itself, what kind of coding would we use? Mike had some mathematicians working for him, and developed new codes that were more efficient. And then the next question is, how do you get the errors down? So we had embedded servos. We developed that way back with the RL01-- first embedded servo drive-- so embedded servos were a key part of it. So we had basically taken a systems approach to get there, to the same areal density, but by different means. And that worked.

[1:50:15]
Gardner: And the first embodiment, then, was in the DSA and the RA80, 81?

Saviers: We were scared by building Winchester disk drives. Probably a good thing to be scared about, especially ones on 14-inch platters. So we said, "Well, the real objective is to build something on the frontier," and that was the RA81, but to get there we'll-- because we don't have all the DSA parts implemented yet, there was some LSI-- to do the error correcting, you need some special silicon. So that was coming along, but it wasn't going to be ready in time, so we built another Massbus with our own Winchester HDA assembly, called the RM80. So we called that the "practice product," and that was build a Winchester drive, start to get the factories learning how to build Winchester drives, debug the product. And then, as soon as DSA was available, complete with all the controllers, we introduced the RA80 on the Unibus controller. That was the first DSA controller that hit the marketplace. And then that was followed by the Q-bus, the QDA, the BI bus, the BDA, and HSC.

Gardner: And then the RA81, which was introduced at a higher capacity?

Saviers: Three times the capacity of the RM80 – the RA81 was 460 megabytes, I think. Everything came together when we had all the controllers done, we had the HSC done, and so then we-- and we had the RA81, we had the RA60, which was a removable drive. There were still customers who said, "I've got to

[4] Although Reed Solomon codes were first disclosed in 1960 [Reed, Irving S.; Solomon, Gustave (1960), "Polynomial Codes over Certain Finite Fields", Journal of the Society for Industrial and Applied Mathematics (SIAM) 8 (2): 300–304, doi:10.1137/0108018], the IBM 3370 announced January 30, 1979 is the earliest known implementation of Reed Solomon codes in a storage device.

have removability,” and I think that was 240 megabytes, or something like that, of removable. And we had an RA70, a five-and-a-quarter-inch drive, which was the high-performance fewer megabytes per spindle. So then we started to put together the differing drives which later became Storage Works, so you could buy a cabinet with 14-inch disks and 5-inch disks, and so on, and removable disks.

Gardner: Before we get back, I’d like to back up and talk about the RL01, but maybe go all the way back, turning now into the small disk track, your early floppy experience. We can talk about the RL01 or floppies.

Saviers: Well, yes. In the-- let’s see. Before 1979-- the tape operation became my responsibility in ’79, but before that, the tape group and I worked for one guy, Bob Puffer, and I said, “Floppy is flexible media. Let the tape guys worry about that.” I had enough on my plate. So it wasn’t a disk drive, it was a tape drive. But anyway, I was involved, obviously, because I knew something about disks, and the early days were trying to buy a reliable eight-inch floppy. Well, I guess the second time I met Al Shugart was-- he came to Maynard, and I guess we go back and look at the dates, it turns out he was still at Memorex at the time, but he wanted to talk to me about floppy disk drives. Okay.

Gardner: We’re reasonably certain this conversation occurred in November 1972.

Saviers: Yes. I think we can pin the date down exactly.

Gardner: Al left Memorex in February of 1973.

Saviers: Right. So in retrospect, after you and I put all this research together, we concluded that Al’s on a fishing expedition, because he was ready to leave Memorex and do something else. So he came in, and he said, “We got this floppy drive at Memorex. We’re using it on the IBM disk controller,” and I guess that even prototyped one that was a little bit more like a real product.

Gardner: Actually, Memorex shipped the first read-write floppy; IBM at that time had only a read-only device. Three companies in 1972 shipped read/write floppies. Memorex, CalComp, and Potter all announced they had shipped read/write floppies that didn’t look anything like what IBM ultimately came out with later in the 3740.

[1:55:00]
Saviers: In that timeframe. So Al comes to the mill at Maynard, and we sit down. We talk for a couple hours, and he doesn’t know much of anything about the minicomputer business or the desktop. We also need to set the stage that there were word processing systems out there using tape drives or cassettes. DEC had a word processing business based on a PDP-8 implementation. So Al says, “Well, here’s what

I'm thinking about doing," and, "What do you think?" And the size, et cetera. Well, I said, "Fine. It's a great idea. We'll sell a lot of them." So he said, "How many?" And I said, "Well, about 10,000." Well, Al thought that was a nice number per year, and I said, "No, that's about 10,000 a month. It's going to be a really big market because everybody's going to want this thing as personal data storage in word processing, for sure, and then on small computers." So I guess about three months later he decides to leave and create Shugart Associates.

Gardner: That's another interesting fact. IBM had not announced the 3740 at that point, November 1972. The IBM key to floppy disk system-- a 3740 system with a 3741 work station was announced in early 1973.

Saviers: Wang was in the market as was Mohawk Data Systems. I don't remember all the key to tape companies. They were the big-- and DEC, a few others. Sperry might've had something. But I don't think IBM had their floppy word processor out.

Gardner: So I don't think Al believed you.

Saviers: <laughs>

Gardner: Because I'm trying to get the number from his original Shugart Associates business plan, and I don't think he saw the total market; not just the DEC piece-- the total key entry market-- as big as hundreds of thousands per year[5].

Saviers: Well, the anecdote to the story was after that he never sold us one because we could never make a Shugart one work. <laughs> But I think we bought everything from CalComp for several years at eight-inch, and then I forget who we bought at five-and-a-quarter. To continue the floppy thread for a moment, when we got to the five-and-a-quarter, we're buying them, and then the double-sided came in, and everybody was kind of, "Does the double-sided really work or not? Does it destroy the media?" A big problem there. But anyway, some guys came and knocked on our door and said, "We have two floppies in a single five-and-a-quarter." It was called T&E Engineering, and the official name was Trial & Error Engineering. <laughs> And so we said, "Okay, we'll give you a hundred grand or something for this design, and you help us put in production," and we did. But it was never a great product. And then, course, the volumes of floppies went up to the point where making it yourself made no sense whatsoever.

[5] Editor's Note: Shugart Associates' 1973 Business Plan, Figure 1, gives annual domestic key entry market size as less than 100 thousand units thru 1977. The total domestic market increases from 327 thousand units in 1973 to 633 thousand units in 1977.

Gardner: I guess by this time, you're now running all of storage, so you're buying a lot of floppies as the world's biggest customer of floppies?

Saviers: I don't know. I think the word processing guys probably owned that honor. And, of course, I was in a separate organization until '79, and by that time-- I have to go back and look and see who was who in the floppy business.

Gardner: Shortly after that, I joined Shugart Associates, and the big players were, of course, Shugart Associates and Tandon. Micropolis had some. IBM was making a lot of their eight-inch, but they bought the five-and-a-quarters.

Saviers: Yes. Who could forget their meetings with Jugi? <laughs>

Gardner: Jugi Tandon-- an interesting guy.

Saviers: Right.

Gardner: So you don't have any interesting stories or experiences you'd like to share buying floppies?

Saviers: The problem was getting a reliable supplier of a reliable product, you know? That was the problem, and so it was the normal shootout (competitive evaluation) process. When anything changed, you went to double density or you changed the form factor. We had the auction again, and, see, everybody would knock on our door, and somebody would get the order. We thought about the kind of relationship that Honeywell, and NCR had set up with Control Data, and didn't think that was a very good idea.

Gardner: That's MPI?

Saviers: MPI-- right. So we thought it much better to be an independent free player for everything, and let the competition sort it out.

Gardner: That's also the same time period you had a conversation with Al Shugart about the ST506. [2:00:00]

Saviers: Well, that was 1980 so we're jumping forward a bit.

Gardner: Shugart (Seagate) was set up in ’79. They showed their first product in May 1980, and started shipping in late 1980.

Saviers: Okay. So it must’ve been ’80, then. Well, I guess the IBM personal computer wasn’t exactly the first personal computer in the market, but DEC had three competing personal computers, which is well chronicled in the annals of bad ideas: the Rainbow, the Pro, and the DECMate. But the Rainbow was almost IBM compatible -- almost compatible was kind of worthless. The Pro was PDP-11-based, and then there was the DECMate word processor —it was PDP-8-based. So somehow Ken had this idea that we should have different architectures for different problems, and I have no idea where one of the all-time worst ideas of the 20th century came from, but that was his idea. And so we had three competing personal computers with different internals. So John Rose was the lead for the Rainbow, and he was very plugged into Microsoft and Digital Research and the guys at, I guess, Lotus, or even the precursor of Lotus was, and the first word processor stuff, and so on. So we were talking to Al Shugart, and he and Finis Conner came by and said, “Here’s what we want to do in a five-and-quarter-inch hard drive,” and it kind of immediately clicked with us. We said, “Wow, that’s the thing that takes a product like the PDP-11 Pro and makes it real.” The PRO group wasn’t quick on the uptake. John was fast on the uptake for the Rainbow, and so we said, “Yes, let’s get engaged on the ST506.” And so seeing that this was really going to change the computer industry, and we had an all senior engineering conference coming up, where we took about 200 of the top people in engineering and took them off to a woods meeting, we called it, up to Stratton, Vermont; rented the ski facility. It wasn’t ski season, but we rented the ski facility. It was cheap. And I invited Al and Finis to come up and show the ST506 before they had an ST506. <laughs> So that was an eye-opener for a lot of folks. Like, “Boy, this is going to change things a lot-- a cheap hard disk in a small form factor.” So John Rose jumped on it with the Rainbow, and actually managed to ship a hard drive on a, quote, “PC” before IBM did. And, of course, the Rainbow was ill-fated. It faded away. But it was the first generation and an interesting relationship and a really great idea. It also led three years, four years later with the five-and-a-quarters really coming up the capacity curve very, very rapidly. We recognized “What this needs now is a companion 5 1/4 tape drive,” and that led to the origin of DLT, the TK50. So we thought, “Look, you got wonderful five-and-a-quarter inch hard drive, but how do you get the software there?” Tape is a distribution media. “And how do you back it up?” 3M had this funny rubber-band tape drive, and there were maybe some helical scan things floating around, but, boy, they didn’t look very good from a reliability viewpoint, handling computer data.

Gardner: Actually, you probably don’t want to start down the tape base now, so if I can cut you off?

Saviers: Okay. We should go back to the RL01.

Gardner: Tell me about the RL01.

Saviers: Well, the RK05 was such a business success, even though there were some real teething problems in getting there. We knew that we needed another cartridge drive for personal computer storage-- high-performance personal computer storage-- and we had a relationship with a guy named Fred Hertrich[6], who had a lab in Boulder and was in charge of IBM's recording technology at one point, before he went out on his own.

[2:05:10]

And so we said, "Look, let's think about what we want, and what we want is to build a cartridge disk drive which is bulletproof. It will last five years with no service calls, it'll never need a head alignment, never need a filter change, and will be office-quiet." It'll sit in my office and not annoy me, like all the disk drives at that point had done. <laughs> So we went to Fred and said, "Do you have some ideas here?" And we knew that Fred had developed an embedded servo idea, a little different reference track idea, for a startup that he had been in after he left IBM, but it cratered. So we said, "Can you develop something for us-- low-cost actuator, quiet, et cetera," and gave him that project as a consultant. Now, at the time, all engineering hiring had been frozen and all budgets had been frozen, so I went around and found a friendly vice president, Andy Knowles, who was in charge of the PDP-11, and said, "Andy, give me a hundred grand to get this project started." <laughs> And he did, and so that became the beginnings of the RL01s. Fred designed it and Pete McLean designed a controller internally for it, and married that all up in Boulder to see if it worked. As kind of a parallel track, DEC was growing very rapidly, and it's time that we need more factories to build disk drives. So let me tell the story about going to Colorado Springs. Fred is in Boulder. So the question is, where do you find a place that you can hire engineers out of California? Can't move many of them to Boston, that's for sure. So we thought the Colorado Front Range was pretty interesting. So at the time, DEC is popping up plants all over the world, and there was this map in Ken Olsen's conference room that had little pins in it with prospective plant sites, so one day the guy who was my manufacturing compatriot, Dave Brown who ran the Westfield, Massachusetts, plant where we made all the disk drives, and I were in the conference room, and said, "Let's stick an extra pin in. Let's put it in Denver." <laughs> So we stuck this pin in there, didn't say anything. So it was a way to plant the seed. So Fred is there. We start talking about moving to Colorado. We start talking about moving engineering to Colorado, and, man, talk about touching the third rail. This is not good with Ken Olsen. If Ken Olsen cannot drive in two hours to the engineering location, it's not going to be any further away. Then I said, "Hey, look, making disk drives, engineers in Boston and manufacturing in Colorado-- this is going to be hard." So I started lobbying, and my boss started lobbying, and over a period of time we convinced the company to have an experiment in Colorado of having some engineering in Colorado. So Dave starts up a manufacturing plant in Pueblo, south of Colorado Springs, and buys 360 acres just south of the Air Force Academy.. That was part of that deal-- buying 360 acres in Colorado Springs of the most beautiful manufacturing site in the country. It's gorgeous. Looks right onto Pike's Peak. And we had to build a railroad bridge—a new experience in the disk drive business. So we buy this site and start building what ultimately became a million square feet of engineering and manufacturing. So all the large drives came out of there, and all the large subsystem stuff came out of there-- HSC. But anyway, so we designed this RL01, and part of the desire of making that five-year, no-maintenance disk drive-- and it had to cost less than $1,000, by the way-- was inventing the embedded servo. So it was the first full embedded servo. No optical secondary track reference. It got everything it needed off the platters. And we used the IBM 5440 cartridge, which, after our 2315 experience, we knew that that was a much more reliable product, much

[6] See obituary of Dr. Fred Hertrich - Father Of DLT, Computer Technology Review, July 1999

more user-robust than the 2315. And it turned out to be a very successful product. Came out at five megabytes. We doubled the capacity to 10 megabytes shortly thereafter, and shipped a bunch of them.

Gardner: And you started up manufacturing in Colorado Springs?

Saviers: We started the manufacturing in Colorado Springs, and I think we eventually made that product in Germany, and may, in the end, have transferred the manufacturing to Springfield, Massachusetts.

[2:10:00]
Gardner: Any interesting stories of servo-writing experiences that you'd rather forget?

Saviers: No. It was a good product. It really met the objectives. We were pretty happy with it, and I think the customers were pretty happy with it. So it was the last of the removable, small-- well, not the last. We built from that, and then we built the RK06 and RK07, which were scaled-up 5440 cartridges. We just stuck more disks inside the basic, similar plastic enclosure design, and those met the needs for upscale small data processing installation, small-- and larger departmental computing. And those were reasonably successful products.

Gardner: Actually, I think DEC holds the honors for being the last disk pack disk drive designing company.

Saviers: Well, that was the RA60. When we laid out the HSC and the large-- the RA80 and the RA81, we still had a large number of customers who said, "I've got to have removability. I can't run my applications without it." And so we said, "Okay, we'll build one."

Gardner: And now we're jumping around in time, but--

Saviers: That's about 1982.

Gardner: And the product number was...?

Saviers: RA60. It was a couple hundred megabytes-- 240, something like that.

Gardner: Short stack of disks?

Saviers: Yes.

Gardner: Would you know how many cartridges you actually sold per disk drive?

Saviers: Haven't got an idea on that.

Gardner: It's one of the great mysteries in our industry. I've asked several people, because IBM shot removability in '74 on the basis of few disk packs per spindle, but of course SMD was very successful, and apparently so was Digital, but no one seems to know how many disk packs per spindle were sold.

Saviers: Well, I think the customer need was for backup, and so you could imagine where they're going to cycle a week's worth of packs through on a daily basis or a-- four on a monthly basis, or something like that. So I think those are kind of the reasonable numbers. It wasn't that I-- some other application where you needed more offline storage. I think the vast majority of removable disk packs weren't removed, <laughs> in terms of the spindles.

Gardner: That's interesting. We'll probably never know.

Saviers: Right. <laughs>

Gardner: A question, although I would share with you later statistics. I think Iomega ultimately shipped 12 to 15 cartridges per zip drive, and I know SyQuest was up around 10 cartridges per its drive.

Saviers: Yes. Well, that's a little different market, though. It's not kind of online data -- it's more like a backup application. It would make some sense. Right.

Gardner: In SyQuest's case, an awful lot of them were mailed. It was a U.S. Postal network.

Saviers: Right. Highest bandwidth networking you could have.

Gardner: For a time, and to this day, probably, a DVD by FedEx is still pretty good.

Saviers: Pretty good.

Gardner: No particular interesting stories about that last disk pack disk drive?

Saviers: Well, I know we didn't sell as many as we thought we'd sell, so it was clearly the end. <laughs> So there wasn't any motivation to do another one.

Gardner: Okay.

Saviers: And all the limitations of prior removable disk packs hit that one, too-- the cascading head-crash problems, et cetera, et cetera. So moving to fixed media increased system availability significantly.

Gardner: Now is probably the good time to go into the RA80. I think we've covered just about every other track we could possibly cover, and I believe you already talked about the need to get on IBM's areal density curve without IBM's components.

Saviers: Right, that was the objective, and that's what Mike Riggle really engineered, designed, architected, and put the resources together to accomplish that.
[2:15:00]
We wanted to build an HDA, the RA80, which would then go triple the density to the RA81, and that maybe had been done before, but it was a pretty big step at that time in the industry to go 3x. So, by the way, at the same time, we were starting to feel some heat from the Japanese. Fujitsu was out there with the product called Eagle, and they had a roadmap for that, and I went over and spent time with them, and became friends with Tatsuda-san, who was the guy running Fujitsu's disk drive operation. So we were pretty concerned about where things were going and that we needed to be much more competitive than we'd been in the past, so the plan was to take the thing in a step of one-- about 120 megabytes-- to a step of three-- 460 megabytes-- and that's what we did, -- just plugged in on the controller, worked fine; all the controllers-- UDA, QDA, BDA, HFC. And then that was very successful. Success generated a set of problems. We had our product recall because of a contamination problem from a supplier-- the things that bite you in the disk drive business when you're not looking and they're not looking hard enough. And then we did one more iteration of it in the RA82, and then that was the end of the 14-inch family. Now, at the time, we also came to another conclusion, which was we knew we couldn't buy IBM technology-level heads and media. The RA81 experience in particular convinced us we couldn't buy IBM reliability or durability-level head and media technology. So we said, "If we're going to make a hundred-thousand-hour MTBF disk drive, we're never going to do it by buying somebody's thin-film head and somebody's oxide disk." They just can't make them anywhere near the quality that IBM-- we'd buy an IBM thing and take it apart and test it, and-- pfff-- it was an order of magnitude better than the best that we could buy from the copycats. So we said, "Look, for the next generation, we've got to be in the head and media business. We may not need to push the envelope as far as IBM, but we'd like to push the envelope in terms of areal density, but we have to push the envelope as far as IBM in terms of quality and reliability of the heads and media." So we put together a plan which led into the RA90 series-- RA90, 92 series of disk drives-- and we went to the board of directors and asked them for $400 million. <laughs> This was like-- smoke came out of their ears, you know? Nobody had ever asked for anything more than like $50 million before for a project. So we said, "Here's what we need: $400 million to build a thin-film head capability, including

a manufacturing plant, a thin-film disk capability, including a manufacturing plant, and to go build this next generation of Digital Storage Architecture disks," which we did. And so we built-- we bought along the way, we were still expanding, planting plants down everywhere around the world. So I bought the Shrewsbury facility-- a couple hundred acres there-- and we sited that to have a thin-film head plant there. And then one of the DEC printed circuit board facilities in Tempe, Arizona, got repurposed into making thin-film disks. So we made plated disks because we were at nine-inch form factor, and you couldn't buy very cost-effective sputtering machines at nine inch, so we went plated. So we repurposed that facility and got into the heads and media business on our own.

Gardner: And this would be when?

Saviers: We shipped the RA90 in the summer of 1988[7] so the request to the board probably was made in the mid-80's..

Gardner: A good decision, bad decision?

Saviers: A very good decision, with the constraint that - and I'll take full credit and full blame here, the corporate objective at the time was we're going to build a VAX 9000 out of ECL, and it's going to go head-to-head in the IBM mainframe market. Now, it was a terrible idea. I was on two of the three committees that were convened to review this project, and every committee said no, but Ken said yes, so he did it anyway.
[2:20:00]
But anyway, we needed mainframe-capable I/O stuff, and that was pushing beyond the capability, or right at the limit of the capability, of what we could do with HSC. It didn't break the storage architecture-- DSA-- but it pushed on terms of throughput and access times. So we're looking at accesses per megabyte per second, and throughputs, et cetera, so it required some new subsystems technology, and so that was part of that plan, also.

Gardner: I was more interested in whether entering the head and disk business was a good decision or a bad decision, in retrospect.

Saviers: Oh I think that was a good decision, 'cause I don't think we could've built the products without it. The thing that-- you can argue this either way. 'Cause clearly IBM had gone the same way with large disks as opposed to 5¼s. If you look at a careful analysis of the RAS features-- reliability, availability, service ability-- and you look carefully at the performance issues, and look carefully at what you can get in terms of dollars per megabyte, it's crystal clear, it's cut and dried, build large form factor disks. Okay? If you look forward maybe five years and say, "What's a disk drive business going to look like? What's the

7 Editor's Note: First edition of RA90/92 Disk Drive Service Manual was dated June 1988.

disk drive market?" Of course, nobody's got that crystal ball. But if you could then you'd say, "Well 5¼s are going to be everywhere-- or maybe 3½s-- and the performance level is going to come up, and maybe if we beat the heck out of those guys they'll build a reliable one. And maybe if we beat the heck out of them, they'll build one with RAS features sufficient; and maybe RAID will get cheap enough that small box drives would've been the way to go. So we were both on that same cusp of one more generation of things that we knew were going to meet the customer requirements, versus a big gamble that 5¼ technology is going to be able to meet-- even with RAID-- is going to be able to meet the customer requirements. Interesting side story. I (DEC) paid for Randy Katz's work at Berkeley on RAID. So we had a real good idea of what a RAID box was going to cost in the late-'80s. And it's acres of printed circuit board; it's reams of code. And it's not easy. We did HSC, back in '81, which had RAID, striping and mirroring, and we knew how hard it was in terms of all the special cases you get into; and mirroring catch-up and partial failures and backtrack and all that stuff is really messy. And so we said, "Someday it's going to be interesting but not now." And that someday was more I think in the mid- and late-'90s when RAID started to really get capable. And I don't think it was really until say RAID 5; that real limitations that would be pounded into them at a data center, probably wasn't about 'til 10 years later.

Gardner: Well, you know EMC with only mirroring took 50% of IBM's market share by 1995.

Saviers: Yes, and mirroring was clearly the way to go.

Gardner: And I think we might agree that mirroring is probably the best way to go; and the RAIDs may not be as good an idea

Saviers: Yes, I think now we're probably beyond that point, where the other RAIDs are a better way to go. But not at that time, it was clear.

Gardner: I'd argue that even today.

Saviers: Yes, I tend to agree.

Gardner: The complexity of the RAID 5 or 6 or 50 or 60 may not be justified given the simplicity of a mirror.

Saviers: Right.

Gardner: But by '95, IBM lost half its market share.

Saviers: yes, right. DEC delivered reliable mirroring-- it didn't come out with Version 1 of the HSC. But I think by 1983 we had N-way mirroring. You could have 10 copies out there if you wanted. We only did two or three, but it had very general capabilities.

Gardner: We'll get more into that in the HSC section. RA90 was what diameter disk?

Saviers: Nine.

Gardner: Nine inch. So you were following the Eagle and the not all the way down to five inches.

Saviers: Right.

Gardner: And it was plated.

Saviers: Plated, going back to 1969, where we had had a process to build plated disks.

Gardner: Now the RA90 was a successful product?

[2:25:00]
Saviers: Yes. About this time is when I leave storage; 1989. I hire my successor. My boss wanted me to do something different in the company. He didn't anticipate that I wanted to do personal computers; but I did. And so I'm turning over the storage business to Charlie Christ, the guy I hired.

Gardner: We all know the ALAR story in plated media. I thought Ampex had given the plated media a totally destroyed reputation.

Saviers: Well there are earlier destroyers. NCR made disk packs in the early '70s which were electroless plated; didn't work out so well. Now they used-- see okay, I wouldn't go back and say, "My plated is better than your plated." But they used electrolysis nickel plating and we used electrolytic; and maybe a difference there. One's much more abrasive than the other. It's got phosphorous in it, it's hard; yada-yada. So anyway.

Gardner: Well ALAR was the great hope in the early-'80s. And you may be the only successful plated media company in the history of the industry.

Saviers: Oh I'm sure that's true. But the interesting thing is here's wafer size and semiconductors going up like this. You know? And you can just look at those early wafers- like one inch-- now we're talking about bigger than 12 inch-- and here's disk drive size coming down. So at some point the disk drive guys can hop on the economy of scale of semiconductor manufacturing equipment.

Gardner: I'll have to think about it. It's always been remarkable to me that disk storage has gone from 24-inch to 65 millimeter disks while the semi guys went in the other direction.

Saviers: Yes. But if you wanted to buy a machine to sputter 9-inch disks in production, in 1988, forget it. You could get one engineered from scratch.

Gardner: So is that an argument that we should be seeing nine-inch disks again because we measure them by sputter image and they sputter nine inches?

Saviers: I think volume swamps the other economies. It was-- and of course track densities are up so high that seek time would become unreasonable. So it's over.

Gardner: That was a facetious question.

Saviers: It's over.

Gardner: I knew the answer. I just couldn't help but asking it. So I think now is the time to talk about HSC. I think we've pretty much covered most everything that I have down. So the first one was the HSC50.

Saviers: Yes. There were HSC50s, then 70s, then 90s, and 40s; all kind of variations. Basically it was all the same chassis with a different number of connectors, and the blades got more capable. So the blades came out initially I think with four storage interconnect ports on them, and maybe four blades that you could put in; so you have 16 drives. And then they got 8-port blades and maybe six, or something like that, so you could have 48 drives. And so- and it was all serial, it was all radial, and the cables were nice and small; about Ethernet size. So it was kind of a great implementation for computer rooms. Cables can be long.

Gardner: Well this is all back to the advantages of--

Saviers: Storage interconnect, the serial interconnect.

Gardner: And the storage interconnect really has nothing to do with the HSC

Saviers: SI wasn't unique to the HSC. But that capability becomes important when you have disk drives spread over two floors. Okay? And you weren't going to have that with a Unibus or a Q-bus system. So that was one of the important capabilities. And the fact that it could do backups; it could do tape. Autonomous backup was never really supported by the operating systems but it was there. You could do disk to tape copy, online and offline. It could interleave that with processing normal requests. It could do-- it did all the bad block stuff, and error correction on the fly; no problem. It did all kinds of optimizations.

Gardner: It did it, or it was done in the drive itself?

Saviers: No it was done in the controller. The data separator was in the drive. But the Reed-Solomon was in the controller.
[2:30:00]
Okay? So the drive sent decoded data, 1s and 0s, across the serial line. So the data separator was in the drive with bad block correction-- error detection and correction in the controller.

Gardner: Oh, that's sort of a different place than the line was drawn in Massbus?

Saviers: That's right. Well Massbus didn't-- it had error detection polynomials, which if you stood on your head you might be able to correct an error. But it was primarily error detection. So Massbus would say "Bad Block"; and then the operating system would decide what to do. It wasn't done inside the Massbus. HSC had all kinds of optimizations. You could start a transfer in the middle of the transfer to increase throughput. The seek reordering; it did request reordering-- everything you could possibly think of in terms of getting some more performance out of it.

Gardner: And so these are all on the blades in the HSC; or within the processing power, the separate control process?

Saviers: Yes, the control processor did what we call the policy decisions and issued commands for the blades to do something: Read these blocks, write these blocks. There was some buffering. We looked at putting a cache in it; several megabytes. And I think that later-- after I left I think some cache was stuck in it. But it did control tapes; it did control fixed and removable drives; it did control 14-inch and 9-inch and 5¼-inch, and ultimately 3½-inch. It did control electronic storage. We built a box of CMOS memory that was very-- basically a replacement for a head per track disk, but very, very fast.

Gardner: So at the DSA level, and the SI level, it was not a SCSI-type contiguous blocks of good data. The bad-block recovery was on the blade side or the-- let's say in a Unibus machine where you didn't have an HSC.

Saviers: No. The DSA architecture, and the disk layout, guaranteed that no matter what-- unless the head crashed-- you've got N sequential perfect blocks on this disk.

Gardner: But the logic to do that was-- say in a Unibus machine-- on the blade side of the SI, not in the device.

Saviers: Yes. Right. If a re-vectoring was called for-- "Here's a bad block; we need to mark that as a bad block"-- the controller would update the tables on the disk drive to not use that, and find the appropriate spare bad-block somewhere. We had spare blocks per track; we had spare tracks per platter; the whole enchilada-- spare cylinders, the whole-- lots of-- I don't know, it was probably 5 percent of the disk was spares.

Gardner: As they are today.

Saviers: Yes.

Gardner: So this is the way we do the mirroring we talked about.

Saviers: Yes.

Gardner: That was again implemented in the HSC. There was some sort of a method of telling it to do one or two or-- how many mirrors?

Saviers: Yes, use some of the commands: I declare the following volumes to be a mirror set; and here's the one that's the one that's the lead. And HSC copied the data, kept it up to date. If a disk crashed, you put a new disk in its place, it copied it back. So it did everything that a RAID system does, in terms of I can unplug it and I can plug a new drive in; and it figures it out.

Gardner: And when was this implemented? Whose idea was it?

Saviers: Well it was part of the Digital Storage Architecture. It was specifically implemented in the HSC. And all-- HSC40 through 90 just basically was an expansion of the chassis; the bigger blades. But anyway, Barry Rubinson, Mike Riggle, Bob Bean, Richie Lary, Pete McLean, Ralph Platz; and all the guys that were involved in architecting and building DSA controllers.

Gardner: Right from the beginning, that was part of the architecture?

Saviers: Yes I think so. I couldn't point to the point where somebody said, "We're going to do this." But we knew that we wanted logical block storage, block addressable, perfect, so the operating system guys never had to worry about all that stuff.

[2:35:10]

Gardner: Now the HSC collectively was a fairly successful product.

Saviers: Well by my accounting-- and by some other folks who were involved with it; so we may be biased-- it was the longest shipping product in DEC's history; it was the highest revenue product in DEC's history; and it was the highest profit product in DEC's history.

Gardner: And we're talking about the HSCs themselves, not the associated peripherals.

Saviers: Correct, just the controller. We think there was a billion dollars of net profit. That's a pretty good number.

Gardner: And digital DEC's typical margins? That would be three or four billion dollars in revenue.

Saviers: Yes. But we made a lot bit more on the HSC; because it was a good product, and we made more on controllers to start with. It lowered the revenue number a little bit I should say; a higher profit percentage a lower revenue number.

Gardner: Is there anything else you would like to tell us about the HSC?

Saviers: Well of course it was a super strain to get all the software to work. The hardware was relatively easy. But I remember in a later career step I made, people came to me and said, "Oh we're going to build this RAID box." And I said, "Well tell me about the software." "Well that's easy." You have no idea how hard it is to build commercial reliability RAID.

Gardner: Somebody once told me micro programs are free but microprogramming isn't.

Saviers: That's right. So it was a phased release; first with basic disk functionality. Then the tape drive functionality came in. Then I think the mirroring capability came in. So over a period of like two years after we first shipped the first one, everything got added. In the beginning the VMS guys couldn't support everything. So over time the number of disk drives, the number of tape drives, the number of HSC's per cluster, all continued to grow. Because I think when we first shipped it, you could have four VAX's and two HSC's; and by the end, at least when I left, it was I think as high as 24 of anything.

Gardner: Of anything?

Saviers: More or less, yes.

Gardner: And at essentially a fairly long distance.

Saviers: Yes, I forget the actual cable length of CI, but 140 feet sticks in my mind as-- so from one radial point you can draw a 140-foot diameter circle, and that's what you could spread out to.[8]

Gardner: Before I segue into tape I would like to ask about interesting experiences on both sides of litigation, being sued and suing.

Saviers: Yes, there were a number of times in which I got to make depositions or show up in Federal District Court. So I was kind of the hard-ass around DEC for protecting our IP. And so we basically took the strategy that we wanted most of the business in the storage world; very profitable, a big chunk of revenue; invested a lot in terms of these unique architectures; patented the architectures; patented the details. So we said, "Why not? We want to protect that stuff." So some customers really didn't like that. But by and large it was successful. So over the years I think we sued Emulex and Systems Industries and EMC; at least once each, if not more than once. And every time we prevailed in terms of getting them to stop doing whatever they were doing in terms of a Xerox copy of our product or infringing on some of the of the patents. Occasionally there was a theft - it went a little further; the guy going to work for one of the competitors with a trunk full of mag tapes with all the data on it about a certain product. So that was fun -- we caught some competitors red-handed. But that was the exception not the rule.

Gardner: Well our good friend Al Shugart had the opposite philosophy. He basically used his intellectual property to cross-license everything.

Saviers: Yes.

[8] Editor's Note: According to Guidelines for VAXcluster System Configurations. September 1992, the maximum length of a CI cable from a node to a star coupler is 45 meters (148 ft).

Gardner: And at least it's a different approach. In retrospect do you think Al had a better idea or you had a better idea?

Saviers: Well it's an interesting commentary, more about- probably about DEC.
[2:40:00]
Because after I left, and Clayton Christensen published his book about the innovator's dilemma-- he crystallized one of the things that I knew was really going wrong at DEC. And it's not the not investing in new technology; because DEC invested in new technology like gangbusters. But it was the notion of under-served, served and over-served markets. And so DEC created the minicomputer business, and served that market pretty much as it wanted, as that market wanted, through 20 years of evolution; at least during the 25 years that I was there. The big computer market was clearly an under-served market, and when Ken decided he wanted to be in that market with the VAX 9000 none of the ducks were in order to really compete solidly in the glass house, toe to toe with IBM's biggest mainframes. We just didn't have the software. We didn't have the support. We didn't have a reliable computer. The VAX 9000 was a disaster in terms of its design and manufacturing. At the low end, where we had grown-- because we were the guy serving what other people were over-serving.; DEC had less. And that's what people wanted, they wanted less. They didn't want a full IBM handshake and full service. So the minicomputer existed because of that. But as we became capable of solving the whole problem for the customer, at least in that mid-range, we walked away in a pricing and business practices from that low end. So if you would say was my strategy good for that under-served, or over-served, market?- okay?-- down there where Sun found the niche and gave us real grief, it was probably the wrong strategy. But for the mainstream strategy where DEC was, it was the right strategy. We managed to maintain something like 87% of the dollars of storage customers spent, spending them with DEC. So when it's a very profitable product, that's a very good business decision. So as a business decision I think completely vindicated in the markets that DEC was in. If you would say, "If you're going to really attack the workstation market credibly"-- which we never did-- "was it the right strategy?" We probably should've had it opened up that's kind of SCSI; if you want to buy the SCSI from XYZ, that's okay. It's interesting, in retrospect, I saw some statistics from Sun, and most of the drives that are connected with Sun are sold by Sun; the Sun SCSI drives. So you could say, "Well what's the industry's commodity." Well when Sun puts its name on it, tests it and verifies it, and makes sure the software supports it, customers are willing to pay for that. Well they might not get 87% penetration, but they still have a good storage business as a result.

Gardner: At most it might be 51%.

Saviers: It could be, right. But that's the nature of that market; and if you want to serve that market, then that's the right way to run the business.

Gardner: And the challenge is-- that's Clayton Christensen's innovative dilemma challenge is serving that-- I guess the term you used is underserved.

Saviers: Well, over-served at the low end. We have more than the customer wants to buy.

Gardner: Well you always do. But if you're used to over-serving, how do you develop an underserved capability to keep your market position?

Saviers: Right, get the impedance match between what you've got and the customer wants; so there's the transaction.

Gardner: Well without-- it is a dilemma. I'm not sure it's an innovator's dilemma. It may be more of a business practices dilemma.

Saviers: Yes. Let me tell you a little story about the litigation. Now Systems Industries kind of went away; that was wonderful. MTI came along in their place; that wasn't so good. We sued EMC a few times. I got to know Dick Egan reasonably well, and after I left DEC he took EMC a different way altogether. And Dick and I would get together at the various investor conferences and have a drink; and so we got to know each other pretty well. So at one of the investor conferences he says, "Grant, I really want to thank you for making me a multi-hundred millionaire." I said, "Why's that?" And he said, "You chased me out of the DEC market, and I had to go after IBM."

Gardner: That's an interesting story. I'll have to go back and think about the EMC history. Well you got sued. You got sued by IBM.

Saviers: Well IBM sued us. It was one of these whoops. We were kind of wandering in the woods of creating our technology base of disk drives for disk drives in maybe '73, '74. And I said, "We've got to make a leap here."

[2:45:00]

So Telex got their pants sued off by IBM for appropriating the 3330 controller design; that's what the basis of the IBM-Telex suit was. So we're talking to Telex -- and I don't know how we got connected, but anyway we got connected. And we knew they had a 3330 drive in development, kind of in prototype form. So we started talking to them about buying their 3330 design. And it turned out the general counsel (Ed Schwartz) of DEC went to school with Nick Katzenbach, the general counsel of IBM. So he calls up Nick and says, "We're thinking about doing this, and want a read from you, because we don't want to step into this mess between IBM and the plug compatible guys. Do you care?" Nick says, "I don't care. Yes, looks okay to me; shouldn't be a problem." So we make this deal with Telex actually. We buy the 3330 design for the princely sum of $160,000.00, and start to work on absorbing the technology. And the next thing, a suit from IBM arrives saying, "You're misappropriating our technology." Well it turned out Katzenbach didn't talk to the guys down the chain somewhere in IBM; and whoops. So we had this lawsuit hanging over us for about two weeks before it went away.

Gardner: Never used any of that technology.

Saviers: Well it was a good education. We tried to build a 3330 short stack out of it. But we didn't-- we couldn't get there. We just didn't have the horsepower to do that. We couldn't design a good enough base plate to save our life.

Gardner: It is a fine art. That's what CDC did-- right?-- a short stack 3330.

Saviers: Yes right.

Gardner: Which became SMD's.

Saviers: SMD.

Gardner: Right ideas and ability to execute.

Saviers: Yup.

Gardner: I think that pretty much finishes the disk side and subsystems side.

Saviers: Uh-hum, that's pretty good.

Gardner: Did I miss anything?

Saviers: No I think-- well I think you've seen the packaging up that we did, in StorageWorks to put more capacity in a single line item the customer could buy; or mixed capacity, high performance 5¼s with-- sorry, high performance 5¼s with high capacity 9s or 14s; and so on.

Gardner: Help me better understand what StorageWorks is or was.

Saviers: Well mostly marketing. But it was packaging. So a customer said, " I need some fast storage because I got a lot of transactions per second. I want smaller capacity, faster access spindles, higher RPM. And then I need a lot of storage, and I want it cheap." So you go to the big drive, the 14 or the 9-inch for the big capacity, cheap; and you go to the high performance 5¼, 7200 RPM, fast seek time, for the fast- high throughput, high transactions per second.

Gardner: So StorageWorks was a DSA product. It had an SI and it was a cabinet into which you could mix and match?

Saviers: Well we preconfigured them. So you got X of this and Y of that. And yes; and then it evolved over time to have-- later, after I left-- it was called StorageWorks. But the basic idea of multi-technology cabinets came in I think probably the mid-'80s, after we shipped HSC for awhile.

Gardner: This was fundamentally going back to the DSA concept, which is mixing them up.

Saviers: Right.

Gardner: Did the SI have to evolve?

Saviers: No, not in my time -- through late-1989 it didn't change.

Gardner: So it had sufficient bandwidth. SCSI arguably, or fiber channel, or all of the current interfaces, about every two years they get an evolution. Right?

Saviers: Yes, Low Voltage Differential, Ultra-SCSI; Super-Ultra-SCSI, etcetera. Right.

Gardner: Right, and Wide SCSI; Very Wide SCSI.

Saviers: Right, yes, yes.

Gardner: But the DSA architecture survived --

Saviers: For 10 years.

Gardner: Ten years?

Saviers: Well yes about-- almost 10 years. And then it needed to change, and I think-- I think the way it went after I got out of storage, it went to SCSI, as the way to connect things together and fan them out. But that had some other set of problems.

Gardner: That's a subject of another interview.

Saviers: Right.

Gardner: Which you have a lot of experience to share with us. I guess it's time now to switch into tape. I'm going to jump way back in time, and tell us about the origin of DECtape.

[2:50:00]
Saviers: Well DECtape was a version of LINCtape, which was developed by Wesley Clark at Lincoln Laboratories when they developed a laboratory instrumentation computer. LINC doesn't stand for Lincoln Labs, it stands for Laboratory Instrumentation Computer. And the idea was how do we get researchers to use a computer rather than wiring up AND gates and OR gates and so on, and relays, to make something happen? And connect analog to digital converters and plotters and so on. So they built a little computer that had a tape drive, which was a pocketsize reel; and it only was a single reel. And it had multiple tracks on it; it had three data tracks, but they were replicated. They were analog added together, so the tape errors didn't happen very often. And it was block addressable. So you could say, "Go to block 32 and read a fixed block." So it was fixed block architecture, and it was block addressable, and the tape controller was smart enough to go find that block. So it was an easy to program, reliable and very convenient personal storage. I think it held around 300 kilobytes maybe-- I'm a little fuzzy on this-- to convert 12-bit words or 16-bit words or 18-bit words-- to kilobytes. But anyway, Tom Stockebrand was involved in it, at Lincoln Labs I believe, and came to DEC and did the first DECtape; which was the same basic idea but a little different. Reel to reel, no capstan; just two motors and you put a little drag on one motor and drove the other motor. And it was variable speed; self-clocked. So it all worked out. And it was the only personal storage of any reliability and decent performance on small computers. So that was there, when I arrived at DEC. And Tom Stockebrand was there. And when I got involved with this first disk project that was in deep trouble, Tom was one of my advisors; because he knew something about magnetic recording. And Roger Cady-- I don't know quite why Roger. But he was the project manager for the PDP-11/20, the first PDP-11. And Tom was doing all the manufacturing automation for DEC He was manufacturing engineering manager for the company. So there would be occasional hallway conversations, that was the degree of coaching I got . But anyway, that led to the RK05, which became the personal media. So DECtape kind of fades into the background at that time. There were maybe one or two generations of cheaper transports of DECtape which weren't my responsibility but were done while I was there. But then the ST 506 comes along down the road-- I'm talking about 1980 now-- and we're looking at larger and larger 5¼ inch disk drives happening. And there's no way to get the software to the customer-- VMS was not a small distribution-- and there's no way to back up when you get any of the 100 megabytes or 50 megabytes of 5¼ inch disk drives. So we thought about that problem, and said, "Wow, we got this 3M rubber band tape drive, but it doesn't look so good. We got helical scan." I guess Exabyte was out there. But helical scan and computers just don't go together too well, in terms of durability. So--

Gardner: You must've been OEMing the conventional 1600 bpi 10½ inch reel drives.

Saviers: Oh yes, we had a family of 25, 45, 75-inch per second-- later 125 inch per second-- tape drives, we bought from Pertec or later IBM or supplier du jour. I wasn't responsible for tape drives until 1979. So some of this is an aside during this period. But anyway Fred Hertrich finishes up the-- so now we're getting into the late-'70s-- Fred Hertrich finishes up the RL01 and RL02, and is kind of knocking around. He was a very creative, very smart guy. And I said, "Fred, can you design us a 1600 BPI IBM tape equivalent in a 5¼ inch form factor?" That was the goal; it has to have data processing reliability and data integrity and performance. And he says, "Yes, I'll give it a shot."

Gardner: Now Fred was a disk guy. This his first tape program?

Saviers: Fred was an everything storage guy. He solved the problem on the data cell, of crashing heads.
[2:55:00]
His Ph.D. was in high-speed mining elevators; cables and thin shell mathematics. And I think his last job at IBM was called Director of Magnetic Recording Technology. IBM used to say on the punch cards, "Do not fold, spindle or mutilate." But IBM has taken every conceivable form factor of magnetic media and folded it, spindled it or mutilated it.

Gardner: When you said "data cell," you meant the 2321?

Saviers: The 2321. His name is on the paper about-- I love this-- Anticlastic Curvature of the tape edge, and how to solve that problem. So anyway-- and he had IBM Outstanding Invention Awards, several of them, by then. But anyway, he's a bright guy; German, spoke German; liked to smoke and liked to drink martinis, played Bach on the piano from memory, and we had a lot of fun together. But anyway, he said, "Okay I'll give it a try." So the idea was-- we called it the TK50-- the idea was about, 50, 60 megabytes in a reel, in a 5¼ inch form factor, that had data reliability equal to a 45 inch per second 1600 bpi IBM compatible tape drive. So he quickly came to the conclusion that serpentine recording was the way to go, and embedded track following was the way to go. Gee, that's interesting; track following is what he did on the RL01. So he designs the TK50. And it was a front loader; cartridge had to go in the front. And the big problem was picking up the leader; a mechanical problem of how do you get this thing to thread? But it worked. And the first generation was a bit of learning in terms of human factors because we had the-- an Olsen door story again; how do you not let people-- how do you build a door strong enough that people don't rip it off with the tape spooled halfway? And how do you get them trained to know, "You can't take the tape out now." If you were playing the Philips cassette, or a VHS, or a Betamax, then just push the button and out it comes at any time. Well you push the button on the TK50, or any single cartridge, single reel, you have to wait. So the first one had some human factors problems; and then for V2 we went into big human factors-- the proper lights, the proper instruction, the proper buttons. And that made the product good. And then we did the TK70. I think it was three or four times the density, Then the TK90, another kick up in density; and now, what's the biggest a DLT ever got to? Eighty gigabytes for

DLT[9] and some humongous number for later versions. So the basic idea: track following servo, serpentine recording. He did the studies, and this is the way you get the maximum amount of data per cubic inch. And it was dead nuts on. It was a great product. At the time, we said, "Gee, tape has more value if it's interchangeable." And at the time-- now we're talking about - maybe '84-- it was clear that if DEC didn't branch out and start selling more volume, we weren't going to be able to support this mass storage business; that we needed to be in the OEM business. And the tape drive was the obvious candidate. Here's value to other people, because it's the same as ours; but in a way it's not the company jewels. So I went around to the Operations Committee, DEC's executive committee, and said, "I'd like to do this." And everybody said, "Oh yes, great idea, let's go do it." So I went to Ken Olsen, and he said, "No." Okay, waited about two months, three months. Go back to Ken Olsen and say, "Ken, I actually have a deal with Hewlett-Packard, and they'll get out of the DAT business if we license them the TK technology with the TK70-- they're going to buy it for a while OEM too." Another "No." So after two no's, I thought maybe it was somebody else's turn. So I took Dave Brown, who was the tape engineering manager, in tow; and I said, "Dave, it's your turn to make the presentation." And that led to a very aggressive confrontation with Ken, where I thought he was going to take a swing at Dave. So we gave up the proposal of selling OEM at that point. History had its revenge. So sometime thereafter-- not in my tenure-- DEC bought some DAT tapes from Hewlett-Packard.

Gardner: That's a good one. And that was the time when we had the reel to reel big machines with vacuum columns.

[3:00:00]
Saviers: Right.

Gardner: But the world was looking for a smaller, reliable backup medium.

Saviers: Yes right.

Gardner: I guess the big contender was 4mm DAT?

Saviers: Yes. Well Exabyte was there.

Gardner: That was 8 millimeter and you couldn't fit it in a 5¼ form factor because it was just too big.

Saviers: Too big, right.

---

[9] DLTtape IV with 2:1 compression, see; Quantum DLTtape® III/IIIxt/IV Data Sheet, Oct 2006.

Gardner: And DAT was 4 mm?

Saviers: Yup. I think it went in, yes. We never-- we kept looking at the helical scan stuff and said, "This is really not too reliable for computers -- it's consumer stuff." And you can make it a little better but you're not going to make it equal to the objective which we had, which was the PE 1600 BPI.. refrigerator class tape drive.

Gardner: The TK70 becomes DLT; digital linear tape.

Saviers: Yes right, right. And after I left DEC and the storage business was parceled out from DEC, because they were in so much trouble, it was given to Quantum; literally given to them. And I think DEC was selling a few hundred-million dollars, at system prices, of whatever that generation was at the time. Dimitrios Lignos was the project manager-- he moved there from RP disk drives. . He actually managed the program with Fred. And then a guy named George Saliba , who we hired from STK , was the guy who drove probably from Generation 4 onward of what became DLT.

Gardner: Did Digital develop all media or did they license people to do the media? How did the media side of the TK/DLT tape business evolve?

Saviers: My recollection is we started with 3M, and after some period of time they had rights to sell it on their own. So I think there were multiple companies in the media business over time.

Gardner: Well its proprietary DEC tape drives but the DEC customers could buy media from multiplier suppliers.

Saviers: Yes.

Gardner: So that created a media market, and then Quantum could sell OEM drives?

Saviers: Right. My understanding was that within a year-- maybe a little longer-- Quantum was selling a billion dollars worth of DLT’s at OEM prices. So they multiplied the volume of the product by five or six X. So a big opportunity was missed by not licensing/selling to HP.

Gardner: Actually I think that's about as far as I want to go this session. We got a whole other part of your life to talk about. But are there things you'd like to share now about the storage industry; what you've learned over the 40 years?

Saviers: Well it's always interesting to see the things which were done once come back again as revolutionary ideas. You know? So there seems to be this reinvention that occurs; or the guys who are now the prime of their youth engineers don't look back on history and see what was done before. So they reinvent it. So that's always interesting. Who would've predicted-- nobody would've predicted that the disk drive industry would be making about a million disk drives a day. That's just kind of like-- kind of fries the brain when you think about it. We thought a couple of thousand a month was really a big deal. So what a change. And of course, the whole lesson about reliability-- it's an interesting lesson in that the more you make of something, the better it has to be, or you're not going to be able to handle the ones that come back. So you think about, let's see, we've got a one percent failure rate and we're making a million day, that's a lot of trucks showing up at your door with bad product. You can't survive that—you've got to be way out there on the reliability curve. I brought Six Sigma to DEC; that was one of my corporate citizenship things. And we were the point of the spear in terms of customer noticeable reliability problems in the field, the disk drives. And so I asked Bob Galvin-- he came to one of my meetings and we talked about Six Sigma, and he gave a presentation and told the Motorola story; very compelling. And so that was the launch of Six Sigma within storage; which then the company picked up as a company-wide program. So that was interesting.

[3:05:00]

Gardner: We have a copy of an '82 speech you gave to Dataquest. And two of the predictions you made there was one, that semiconductors would not pass storage in the 20th Century. And so far that prediction's pretty good.

Saviers: Yes.

Gardner: You were a little bit tough on small disk drives in your '82 presentation. You would've been surprised to see perhaps the 3½ inch dying, and the only thing we were making was 2½.

Saviers: Well I think the visibility at that point-- form factor-- sizes are drives the industry? So okay, it's laptops now. It's not desk size. In that speech-- I've looked at it again-- it never really considered that form factor is much of an innovation. It's not technology. It's sticking it in a different size box. Sometimes you get some pretty challenging problems, like this tiny little motor or this tiny little actuator you need to build. But I think that that speech was about the rate at which the small disk drive guys were consuming the free technology from IBM. And at some point, here's the IBM curve going like 30% per year, and here's the small disk curve going twice or three times that; and unless you start inventing your own technology at some point it doesn't work. So DEC had already made the decision that we were going to invent our own technology. But Seagate hadn't, Connor hadn't, Maxtor hadn't. So it turned out that there was enough volume that the component suppliers could advance the technology -- and the drive makers themselves started to spend some serious bucks on technology. In preparation for this, I found a document from Mike Leis, who was one of the chip designers on the original Reed-Solomon at DEC, he did the same work at Quantum, and then at Maxtor. Well he retired before the next thing happened. But you look at his chip history, and clearly a lot of technology had to go into those chips over the time in order to accelerate the bending of the curve upwards; which did occur in the disk drive industry. My point of that speech was to give everybody a kick in the butt. Whether it made any difference or not, who knows?

Gardner: Just like you give the industry a kick in the butt today.

Saviers: Well I'm so far from it now. So much has gone on - so much water over the dam. We did a lot of interesting things at DEC; very innovative things, which in some cases are still being reinvented. But I went away from drive design and recording stuff in '89/'90; dabbled a bit in it, in some investments. But in terms of knowing what's going on, 20 years later I'm obsolete as I can be.

Gardner: Another side perspective on Grant Saviers is you're on the board of trustees at the Computer History Museum. Would you like to tell us a little bit about how the Computer History Museum came about?

Saviers: Oh yes. My life with the Computer History Museum began in Gordon Bell's living room in Lincoln, Massachusetts. I got hooked; because he had a wonderful collection of calculators and slide-rules and early stuff, in his living room. I had a little experience with the Univac I in college. Gordon's living room was one beginning of the DEC Computer Museum. Then he and Ken started saving DEC stuff, and then they saved parts of Whirlwind. And DEC Computer Museum was in Marlborough, the old RCA facility. And then it became its own entity, The Computer Museum in downtown Boston. And then it got into trouble and the collection came out to Moffet before Shoreline. So I've been involved oh--whatever it is; 35 years or something like that. And this is the ultimate- this is so far the ultimate expression of it; extremely well done. We've raised a lot of money and managed it very well. I think the exhibit speaks for itself. And it's a tremendous place to tell stories and meet the real creators of the computer industry, and share perspectives on what's happened and what might happen.

[3:09:23]
END OF INTERVIEW